



HARVARD | BUSINESS | SCHOOL

DOCTORAL PROGRAMS |

DISSERTATION ACCEPTANCE CERTIFICATE  
*(To be placed in First Library Copy)*

The undersigned, approved by the Chair of the Doctoral Programs,  
have examined a dissertation entitled

**Essays on Online and Multi-Channel Marketing**

presented by **Lingling Zhang**

candidate for the degree of Doctor of Business Administration  
and hereby certify that it is worthy of acceptance.

Signature.....  
*Doug Chung, Co-Chair*

Signature.....  
*Anita Elberse, Co-Chair*

Signature.....  
*Sunil Gupta*

Date.....*4/14/2016*.....

**Essays on Online and Multi-Channel Marketing**

A dissertation presented

by

Lingling Zhang

to

The Marketing Unit at Harvard Business School

in partial fulfillment of the requirements

for the degree of

Doctor of Business Administration

in the subject of

Marketing

Harvard University

Cambridge, Massachusetts

May 2016

© 2016 Lingling Zhang

All rights reserved.

Essays on Online and Multi-Channel Marketing

ABSTRACT

Firms increasingly adopt online and multi-channel marketing strategies to reach and persuade consumers. Therefore, designing an effective marketing mix is critical to their success. The aim of my dissertation is to understand the strategy behind firms' channel choices and assess marketing effectiveness. It consists of three large-scale empirical studies examining several important aspects of online and multi-channel marketing.

My first essay focuses on the business-to-business (B2B) interactions involving online platforms, which serve as new channels for traditional merchants to reach consumers and grow business. Using data from the daily deal market, we specify a structural model that examines consumer choices on the demand side and firm strategies on the supply side. In particular, we incorporate merchant heterogeneity and allow prices to be jointly determined by merchants and platforms through negotiation; both of these match the real-world complexity but are challenging

to be modeled theoretically. Our results show how platform size, commission rate, and the allocation of price-bargaining power jointly determine the price setting and the platform differentiation among merchants.

Essay two studies to what extent marketers' actions can affect the reach of video advertising campaigns through influencing the amount of user-generated content. To do so, we compile a unique and comprehensive data set on ad campaigns conducted on video sharing sites such as YouTube. We find that several instruments under the control of advertisers can influence how much the reach of a campaign benefits from user-generated content. Our results underscore that, with the right strategy, advertisers can substantially increase the number of impressions that their online video campaigns yield.

Essay three assesses the effect of advertising and personal selling in the U.S. presidential elections, where advertising involves both candidate campaign ads and those sponsored by outside political interest groups and personal selling takes the form of field operations. We set up a structural model that treats campaign allocation as endogenous and also allows the campaign effect to vary across individuals. Among the many findings, we show that field operations are more effective for partisan voters whereas candidate campaign ads are effective for non-partisans. Interestingly, ads from outside political groups are more effective for partisans than for non-partisans. Our counterfactual results indicate that field operations play a critical role in the

2008 and 2012 elections while the importance of ads is only substantial in a close competition like the 2004 election.

## TABLE OF CONTENTS

Acknowledgements.....	vii
<i>Essay 1</i>   Strategic Channel Selection with Online Platforms: An Empirical Analysis of the Daily Deal Market .....	1
<i>Essay 2</i>   Viral Videos: The Dynamics of Online Video Advertising Campaigns .....	56
<i>Essay 3</i>   The Air War versus the Ground Game: An Analysis of Multi-Channel Marketing in U.S. Presidential Elections.....	95
References.....	147

## ACKNOWLEDGEMENTS

I started my doctoral training after having worked as a statistical consultant for several years. Life then was easy but I wasn't content. The desire for learning made me decide to go back to school. I am truly grateful for having had the opportunity to study and research at HBS; during my stay here I constantly felt challenged and motivated. As I reflect on the past six years, I am full of gratitude and appreciation. I wouldn't have survived to write a dissertation were it not for the invaluable advice, support, and friendship from so many people over the years.

First of all, I am deeply indebted and forever grateful for my two beloved advisors. Doug has been and will continue to be my source of inspiration. I cannot possibly count all the things that I learned from him. To say the minimum, he is the one who taught me to enjoy hard working and take pride in conducting high-quality research. His passion, curiosity, and persistence have inspired me profoundly. I cannot thank him enough for never giving up on me, for holding high expectations for me, and for selflessly investing time in me. Without his guidance, support, and encouragement, I would not have been who I am as a researcher. I am also indebted to Anita. She took me under her wings when I had little idea what research was. She encouraged me to aim high, to be independent and persistent, most importantly, to be confident in myself. She is my role model for commitment to true passion of work.

I am also deeply grateful for Sunil, who has always encouraged me to think big and think deep. He cares about students and generously offers me time and advice whenever I need them.



He sets an example for what a great educator and mentor ought to be. I am fortunate to have received his guidance and will always keep it in mind in my career.

There are many other former and current faculty members in the Marketing unit who have helped me tremendously—Donald, Das, Thales, Tom, Ayelet—and the list goes on. I will miss the wonderful group of Marketing students, whose friendship means a lot to me. Of course, I am the saddest to say goodbye to my officemate, Bhavya, who has stood by me and accompanied me through all the ups and downs over the years. We are each other's most adamant supporter and cheerleader. In many ways, she is part of the family.

I am also fortunate to be a member of the 2010 DBA cohort. I love all of them. Shelley, my dear friend and sister, is a most eager sounding board for my preliminary ideas. She has so many qualities that I admire; I consider it a privilege to have her friendship. I will greatly miss the numerous conversations we had at lunches and walks.

My friends—in particular, Yang, Ge, Yuan, Yuanyuan, and Xiaomu—contribute to my life in a way beyond words. They cheer me up when I feel down and keep me grounded when it is due. They are always there for me, willing to help in whichever way necessary. I am forever grateful for their loyalty and friendship.

Lastly, I want to thank my family. Thank you, mom and dad, for loving me as who I am! Being your daughter is the luckiest thing in the world. My sister, who has grown to be a best friend, takes on a lot of family responsibilities so that I can focus on work. My husband,

Guoqing, is truly the one who leads me onto the path of research. He is the best husband and father, and is the anchor of the family. Life would not be the same without him. Finally, my dissertation would not have been remotely possible without the cooperation of my daughter Annelie. I fail to remember the countless moments when she, voluntarily or not, left me alone to work. Even with the limited attention from mom, she turned out to be a caring and cheerful kid. For this, I am forever grateful.

Strategic Channel Selection with Online Platforms:  
An Empirical Analysis of the Daily Deal Market

LINGLING ZHANG

DOUG J. CHUNG

Abstract

The prevalence of platforms opens new doors to traditional businesses for customer reach and revenue growth. To take advantage of this opportunity, it is critical to understand the dynamics of platform choice. Theoretical predictions may not be enough to guide managerial practice, because the stylized models often fail to incorporate the real-world complexity, where merchants are heterogeneous and prices are determined through merchants negotiating with platforms. In this research, we set out to empirically understand price setting and profit sharing in an online platform market. To do so, we compile a unique and comprehensive dataset from the U.S. daily deal market. We specify a two-stage structural model based on Nash bargaining solutions, and conduct counterfactual analyses using parameter estimates. Our results shed light on how the size of platforms, the commission rate, and the allocation of bargaining power jointly determine the price setting and the platforms' differentiation among different types of merchants. We find that merchants' price-bargaining power vary and that larger and chain merchants have higher influence on price setting than smaller and independent merchants. We also find that merchants pay lower transaction cost and have higher bargaining power on the smaller platform,

suggesting that merchants can take advantage of the smaller platform's lack of market power despite its smaller customer base. Our counterfactual results show that larger and chain merchants are more incentivized to use the larger platform when the sales there can compensate the higher transaction cost and lower price bargaining power.

## 1 Introduction

Platform companies such as Amazon, eBay, and Groupon have gained notable growth momentum during recent years and attracted much attention among researchers and practitioners. Numerous start-ups operate on the platform concept and existing companies are looking for ways to become platforms. The prevalence of platforms has opened new doors for even traditional businesses to take advantage of what platforms can offer. However, with new opportunities come new challenges: understanding the strategy for choosing platforms and the dynamics of working with them has become more urgent than ever.

A platform business simultaneously serves end consumers and business users, often referred to as “sides”<sup>1</sup>. It has two salient characteristics: (a) the business exhibits network externalities, that is, the benefit to one side from using the platform increases with the size of other sides; and (b) the growth of the platform depends on the relative prices charged to all sides. Therefore, when it comes to the question of why one platform is chosen over another, theoretical research often focuses on the size of the platform and its pricing strategy (e.g., Armstrong 2006; Rochet and Tirole 2003, 2006). Although we gain insights from theoretical predictions, it is unclear to what extent those results can guide practices in the real world, which is much more complicated than the highly stylized models used in the theoretical literature. In particular, two things are often missing. First, theoretical models often focus on how a “typical” business user chooses platforms, while the reality witnesses tremendous heterogeneity among businesses. For example, business may vary in their attractiveness to consumers and price elasticity; therefore, it is

---

<sup>1</sup> Some platforms serve different groups of consumers, rather than businesses. Matching platforms are one example. In this paper, we focus on the type of platforms that bridge end consumers and businesses (such as local businesses, content providers). This type has been the center of concern for theoretical papers as well.

reasonable to conjecture that businesses with varying level of quality may adopt different strategies for choosing platforms. Second, most extant literature assumes that platforms have full pricing discretion and make take-it-or-leave-it offers to the sides, although in many industries prices are determined through a bargaining process between business partners. After all, price bargaining is pervasive in business-to-business (B2B) contexts and the interactions between platforms and business users are no exception. An eye-catching example is that Amazon and Hachette, the fourth-largest publisher in the U.S., settled a much-debated dispute in 2014 and signed a contract concerning pricing and profit split for e-books.<sup>2</sup> With the presence of price bargaining, extant theoretical insights wait to be validated in real-world settings.

In this paper, we empirically study the competition in a two-sided market. We ask two questions: (1) what are the determinants for price setting and profits splitting between platforms and their business users; and (2) how do the dynamics of the pricing decision determine the platform differentiation among business users. In order to answer those, we allow both sides of the platform to be heterogeneous and do not restrict the sides to single-home on one platform. In addition, we incorporate bargaining in the pricing decision and explicitly estimate the distribution of price-bargaining power across different types of business users. All of these map the complexity of real-world business settings, but are known to be challenging to model theoretically, which makes it an exciting opportunity for an empirical study.

We set out to answer the research questions using data from the U.S. daily deal market, where deal platforms connect local merchants and consumers by selling daily assortments of discounted goods and services. Merchants use deal sites primarily to bring in potential consumers as well as generate revenues by selling deals. We choose this empirical setting for several reasons. First, the daily deal market is a representative platform business and price

---

<sup>2</sup> Streitfeld, David. "Amazon and Hachette Resolve Dispute". *New York Times*, November 13, 2014.

bargaining is an important element in transactions between platforms and merchants. Second, the daily assortment of deals provides much-needed data variation in the number and variety of merchants within a short period of time, helping model estimation. Third, the market is largely a duopoly competition between two deal sites—Groupon and LivingSocial; therefore, it is relatively straightforward to examine merchants’ tradeoffs in their platform choices. Fourth, the daily deal business in 2014 had \$3 billion in revenue in the U.S. market alone and even more in developing economies, making it an important market to study in its own right.<sup>3</sup>

Our research setting poses several modeling challenges. First, the size of the consumer base for a deal platform is endogenously determined by the size and composition of the other side—the merchants. Furthermore, as pointed out in the platform literature, a consumer’s decision should condition on her expectation of the other side. Second, the pricing process in this setting involves many interrelated components. When dealing with merchants, platforms can set their commission rate but cannot fully decide on the deal’s final price, which is jointly determined by negotiations between platforms and merchants. Therefore, both platforms and merchants act strategically on price setting and merchants internalize their bargaining power in the decision to choose platforms. Third, during the pricing decision, the platform considers not only how much revenues would be generated from selling the focal deal, but also how much the deal can attract customers and thus help sell other deals. At the same time, the merchant calculates not only the current deal sales but also the future payoffs from retaining customers acquired via the deal promotion. Note that all of those challenges, perhaps except the future payoff consideration for merchants, are applicable to empirical settings beyond the daily deal market. Therefore, the results from this study can be generalized to other two-sided markets when merchants and platforms split the pricing discretion.

---

<sup>3</sup> *IBISWorld*. “Daily deals sites in the US: Market Research Report.” December 2014.

Taking those challenges into account, we specify a two-stage structural model using unique and comprehensive data on the demand and supply of the daily deal market. In the first stage, deal platforms and merchants negotiate—through an independent bargaining process—the price charged to consumers. Despite being prominent in B2B markets, price bargaining has only recently been examined empirically. We model the outcomes of the platform-merchant negotiations following the Nash bargaining solution pioneered by Horn and Wolinsky (1988) and further developed by Crawford and Yurukoglu (2012). The solution of this supply model specifies the prices that solve the Nash bargaining problem between a platform and a merchant, conditional on other observed prices. This model is flexible enough to nest the scenario in which platforms have all the pricing power. Thus, it essentially becomes an empirical question to estimate the allocation of price bargaining power across deals.

In the second stage, we examine consumers' decisions to choose platforms and deals given the prices determined in the first stage. A consumer first needs to decide which platform(s) to use. Conditional on that choice, she decides which deal to buy. We formulate that a consumer's choice of a platform is consistent with her expectation of its value, which depends on the quantity and quality of the deals offered on the platform. By this specification, we endogenize the network effect such that the size of the customer base depends on the assortment of merchants on the other side. In contrast to many two-sided-market papers that simply specify the number of consumers as a function of the number of merchants, our approach explicitly takes the composition of the merchants into consideration.

Our demand specification incorporates the heterogeneity of consumers' price sensitivity, but the distribution of price elasticity needs to be estimated based on aggregate sales data rather than



on individual-level data that are hard to obtain on a large scale.<sup>4</sup> To address this challenge, we cast our estimation using the “BLP” model—the random-coefficient aggregate discrete-choice model developed by Berry, Levinsohn, and Pakes (1995). We modify the BLP method to allow the deal decisions to nest under the platform choices. We also adopt a technique to accelerate the computation speed, because the original BLP estimation performs extremely slowly for such a large sample size as ours. We use the squared polynomial extrapolation method (SQUAREM) to speed up the convergence. To address the concern that deal prices are endogenous to unobservable demand shocks, we use instruments to estimate the price coefficients.

Our estimated distribution of price elasticity indicates that customers in the daily deal market are price sensitive and that there is large variation in price elasticity across deal categories. Consumers are the most price sensitive to deals on life skill classes (for example, computer training), with an average estimated elasticity of -10.45. They are also fairly price sensitive to beauty deals (-10.19) and personal care deals (-8.83), and the least sensitive to deals on live events such as concerts and performances (-2.07). We also find significant heterogeneity among consumers: those who are older, have higher incomes, or are from a larger household tend to be less price elastic. Furthermore, our results reveal varying consumer preferences for different deal categories. Beauty deals and home and automobile services are the top two categories that help platforms grow their customer base. Deals on life skill classes, live events, outdoor activities, personal care, restaurants, sports, and travel activities are effective as well.

Results from our supply model shed light on the primitives that determine the pricing outcomes and profit split between platforms and merchants. We find that Groupon, the larger

---

<sup>4</sup> When modeling competition in a two-sided market, it is important to incorporate the full scale of both sides in order to account for the network effect. Using individual-level data from a representative sample is sufficient.

platform, charges a higher commission rate than LivingSocial, and that merchants have lower bargaining power on Groupon than they would have on LivingSocial, both of which are perhaps rationalized by Groupon's larger customer base. We also find evidence that merchants exert influence on price bargaining and that different types of merchants vary in their bargaining power. In particular, chain merchants and those with more employees tend to have higher bargaining power than independent and smaller merchants. When comparing the relative bargaining power between the merchant and the platform, we find that, interestingly, an average-sized independent merchant has lower bargaining power than the platform but an average-sized chain merchant has higher bargaining power than the platform. In other words, chain merchants can dominate the pricing decision for their deals but independent merchants often cannot.

Based on the parameter estimates we then conduct counterfactual analyses to disentangle the effect of price bargaining power, customer retention rate, and platforms' commission rate. We first look at the effect of merchant's price bargaining power. With higher bargaining power, merchants are able to shift the final price closer to their most-preferred level, yielding an increase in profits. With a 10% and 30% increase in the bargaining power, merchants would end up with an 4.2% and 11.6% increase in profits, respectively. Second, we study how much merchant benefit from retaining the deal consumers as regular customers. If they are able to increase the retention rate from 25% to 30%, the merchants would see a profit increase of 45.9%. The boost is as high as 95.7% when the retention rate is 35%. Both counterfactual analyses provide managerial implications on how much merchants are willing to invest to increase their bargaining ability and retention ability, respectively. Our third counterfactual analysis directly speaks to the question of how platform size, commission rate, and bargaining power jointly determine how platforms are differentiated among merchants. Our results suggest that merchants

work with LivingSocial to take advantage of its lack of market power, and that they would use Groupon when the sales there can compensate the higher transaction cost and their lower bargaining power on Groupon. This is more so for larger and chain merchants because they have a higher chance of selling more deals than smaller and independent merchants.

Our research makes several contributions. First, it empirically studies an important question of price setting and platform competition in a two-sided market. We incorporate three types of merchant heterogeneity: (a) price elasticity of their goods and services, (b) their ability to attract consumers to platforms, and (c) their ability to influence the pricing decision when dealing with platforms. Results from this paper can help understand the strategic interactions in this domain and generate managerial implications more targeted to merchants of different types. Second, our research also contributes to the empirical work of price bargaining. To the best of our knowledge, this is one of the first empirical marketing papers that examine price bargaining in a competition with network externalities. While two-sided markets have attracted marketing researchers in recent years (e.g., Dubé et al. 2010; Shankar and Bayus 2003; Wilbur 2008), price bargaining is either assumed away or inapplicable; so little is known about this important business practice. We bridge this gap and believe that our approach offers a good modeling framework to study platform competition where platforms do not have full control over pricing. Finally, this paper is also related to the stream of research on daily deal market (e.g., Subramanian and Rao 2016), which has become an interesting subject of study due to its increasing presence among consumers.

The rest of the paper proceeds as follows. Section 2 reviews the related literature. Section 3 describes the empirical setting, reports summary statistics, and provides some model-free evidence. Section 4 specifies the model, and Section 5 presents the estimation and identification

strategy. Section 6 presents the parameter estimates and the counterfactual results. We conclude and discuss future research directions in Section 7.

## **2 Literature Review**

This paper builds on two streams of research. First, it extends the rich literature on platform competition. Ever since the pioneer paper by Katz and Shapiro (1994) that highlighted the importance of network externalities, a series of theoretical papers have studied the role of the pricing structure on platform competition (Rochet and Tirole 2006; Rochet and Tirole 2003; Armstrong 2006; Caillaud and Jullien 2003). The key insight is that the price charged to a side is inversely related to its price-elasticity adjusted by its strength of network externality on the other side. A stream of platform research that is directly relevant to this study is how platforms are vertically differentiated (Caillaud and Jullien 2003; Jullien 2005). In other words, under what circumstances do merchants choose one platform over another? In the attempts to endogenize the decision, researchers have attributed the platform differentiation to three possible reasons: platform size, transaction cost, and consumer heterogeneity. Jullien (2005) notes that, with the presence of two active platforms, merchants use the one with higher transaction cost only when they cannot achieve the same payoff on the platform with lower transaction cost. Consumer heterogeneity can also lead to merchants distributing on different platforms, if they aim to reach a particular type of consumers only available on certain platforms. As aforementioned, this line of research has largely been silent on merchant heterogeneity; however, it is reasonable to conjecture that different types of merchants may adopt different strategies to choose platforms, which in turn would cause vertical differentiation between platforms.

Merchant heterogeneity becomes perhaps even more important when the business decisions between merchants and platforms are negotiated. In that case, merchants may also differ in their

bargaining power to shape the outcome of the strategic interactions with platforms. In particular, in many markets merchants and platforms split the control over price setting; however, extant literature has almost always assumed that platforms set the price. Some recent papers relax this assumption. Hagiu and Lee (2011) study the pricing control between content distributors (platforms) and providers. By examining two extreme conditions—in which either the platforms or the content providers set the price, they find that how pricing control is distributed between platforms and content providers may determine the extent to which content providers are willing to be exclusive on one platform. One obvious restriction of this paper is that it studies the extreme cases and neglects the fact that prices are jointly determined through negotiations in many B2B settings. Shao (2015) studies more flexible price negotiations between platforms and content providers, and finds that an entrant platform’s greater bargaining power would make content providers more willing to work exclusively with the incumbent.

Our paper is similar to Hagiu and Lee (2011) and Shao (2015) in the sense that we examine the extent to which control over pricing (that is, the allocation of price-bargaining power) affects price and market outcomes. We differ by empirically examining the phenomenon using transaction data from a two-side market. In addition, we incorporate heterogeneity among consumers and merchants. Allowing merchant heterogeneity enables us to attribute the relative bargaining power to merchant characteristics, such that we can provide managerial implications more targeted to different types of merchants.

Our research also contributes to the empirical work of price bargaining between firms and suppliers. Despite the pervasiveness of bargaining in B2B environments, empirical treatment of this subject has only recently gained traction. The bilateral Nash bargaining model proposed by Horn and Wolinsky (1988) is advanced by Crawford and Yurukoglu (2012) to study pricing

decisions between content distributors and conglomerates in the cable television industry. This Nash solution has since become the workhorse bargaining model for predicting the payoff split during B2B transactions in many applied settings. Grennan (2013) examines the role of bargaining power in price discrimination among hospitals in a medical device market. Gowrisankaran et al. (2015) estimate a bargaining model of competition between hospitals and managed care organizations. Different from those papers emphasizing the effect of bargaining on social welfare, our research aims to provide managerial implications for merchants on when to use which platforms.

Empirical marketing studies of price bargaining are sparse. The closest to our paper are Misra and Mohanty (2008) and Draganska et al. (2010), both of which examine bargaining in the retailing setting. The former estimated the relative bargaining power of manufacturers supplying to a single retailer and the latter extended the framework to include competition between retailers. In contrast to their works, the current research focuses on bargaining in a two-sided market, which distinguishes itself from retailing in two critical ways. First, the network effect between merchants and consumers is more prominent in a two-sided market than in retailing. Hence, it is critical to capture the externality value of a merchant to a platform and further allow that to enter the pricing decision. Second, the two settings also differ in terms of where the strategic actions may posit. In the retailing setting, retailers typically are the only strategic players when competing for consumers after they already contract with manufactures. However, platforms often facilitate the transactions between merchants and consumers, and hence both platforms and merchants may be strategic in setting the prices charged to consumers (Hagiu and Lee 2011). In other words, the bargaining outcomes may have a more direct impact on consumers in a two-sided market than in retailing.

### **3 Data and Model-free Evidence**

#### **3.1 Empirical setting**

Daily deal sites emerged around 2008 as a marketplace that connected merchants to consumers by offering discounts. This business model experienced skyrocketing revenue growth for several years. In 2010, the Chicago-based market leader, Groupon, became the “fastest growing company in history.”<sup>5</sup> Several factors may have accounted for such growth: consumers enjoyed a wide variety of deep discounts while merchants could use the deal platforms’ large customer bases to build awareness and generate extra revenue. Even though growth has slowed in recent years, the daily deal business remains a multibillion-dollar market.

The name of the business—daily deals—refers to the fact that the sites in their early years typically featured one deal per day. This quickly evolved to multiple deals a day. A platform’s customers now have access to a searchable inventory of available deals and typically learn about new deals through email alerts or mobile app notifications or by visiting the platform’s website. The vast majority of the deals are from local businesses, although platforms do occasionally promote deals offered by national merchants to build awareness, acquire new customers, and generate additional revenue.

The business model attracted a number of competitors, ranging from small local deal aggregators to large companies that offer deals as a sideline; Google Offers and Amazon Local are prominent examples. By and large, the market is dominated by two sites—in 2013, Groupon and LivingSocial earned roughly 59.1% and 16.6% of the U.S. market’s revenue, respectively.<sup>6</sup>

---

<sup>5</sup> CNBC, December 2010, <http://www.cnbc.com/id/40454493>, accessed February 10, 2015.

<sup>6</sup> Statista 2015. Retrieved from <http://www.statista.com/statistics/322293/groupon-market-share-us/> on February 15, 2015.

We compiled a comprehensive dataset from these two market leaders. It has four components: (1) deal data including sales, price, and other deal-level characteristics; (2) platform-level market shares; (3) the distribution of consumer characteristics; and (4) merchant characteristics. Figure 1.1 illustrates the data components and the sources.

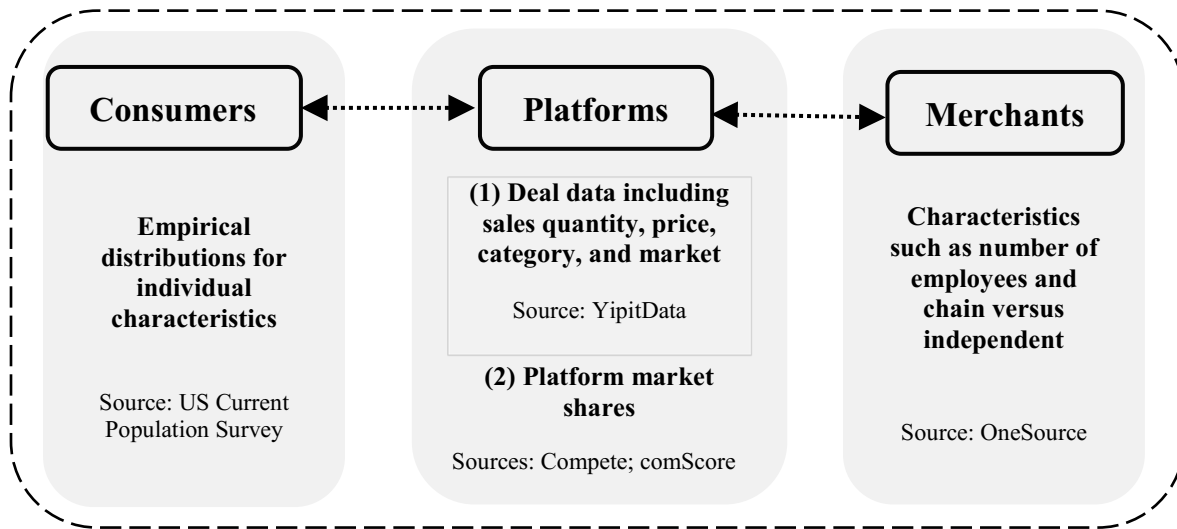


Figure 1.1: Illustration of data components.

### 3.2 Deal sales

We acquired deal sales from YipitData, a premium business database. Our data include all the deals offered by Groupon and LivingSocial in 2012. Each observation is a sales record for which we know the deal title, price, sales quantity, discount depth, face value, starting date, ending date, category, city, and merchant information. For example, Groupon featured a restaurant deal titled “\$79 for an Italian Steak-House Prix Fixe Dinner for Two with Wine at Padre Figlio (Up to \$189 Value)” from June 27 to July 3 in New York. In this case, the price is



\$79, the original face value of the voucher \$189, and the discount depth 63%. We also know the sales quantity for each deal.

Table 1.1 presents summary statistics for the deal data. In 2012, Groupon promoted a total of roughly 129,000 deals, with an average price of \$59.26 (SD=\$61.15) and an average sales quantity of 244.2 (SD=886.0). Deals were evenly distributed over the months with slightly more offered in the third quarter. LivingSocial offered approximately 69,000 deals. The average price was \$48.3 (SD=\$48.1) and the average sales quantity was 274.45 (SD=1,259.8).

Table 1.1: Summary statistics for deal characteristics and sales

	Platform	Mean	SD	Min	Median	Max
Groupon (N=128,749)	Sales	244.19	885.97	1	90	100,000
	Price	59.26	61.15	1	39	400
	Discount	58.70	12.21	0	53	99
	Face value	196.36	317.16	2	100	9,600
LivingSocial (N=69,340)	Sales	274.41	1,259.82	1	92	94,226
	Price	48.29	48.11	1	35	400
	Discount	57.39	11.67	0	51	100
	Face value	136.23	173.55	4	85	5,950
All deals (N=198,089)	Small market	0.18	0.39	0	0	1
	Medium market	0.33	0.47	0	0	1
	Large market	0.49	0.50	0	0	1
	January	0.07	0.25	0	0	1
	February	0.07	0.25	0	0	1
	March	0.07	0.26	0	0	1
	April	0.07	0.26	0	0	1
	May	0.08	0.27	0	0	1
	June	0.08	0.27	0	0	1
	July	0.08	0.27	0	0	1
	August	0.09	0.29	0	0	1
	September	0.09	0.29	0	0	1
	October	0.10	0.31	0	0	1
November	0.10	0.30	0	0	1	
December	0.10	0.30	0	0	1	

A deal belongs to one of twelve categories. Table 1.2 presents the size of each category and the summary statistics for price and sales by category. Across both platforms, the largest category is beauty followed by home and automobile services deals and restaurant deals. The relative sizes of the categories are largely comparable across platforms except that Groupon offered more goods deals than LivingSocial but the latter had more family deals and fitness deals.

Deal prices vary substantially across categories. In general, the average deal price for a category was higher on Groupon than on LivingSocial except for live events deals, which had a higher average deal price on LivingSocial. Sales varied across categories and platforms. Groupon had higher average sales than LivingSocial for family, fitness, live events, and restaurants categories; LivingSocial had higher average sales for the others. We depict the number of deals and the average sales per category in Figure 1.2.

Table 1.2: Deal categories on platforms

	N	%	Price		Sales	
			Mean	SD	Mean	SD
<i>Groupon</i>						
Beauty	24,657	19.2%	91.7	78.2	135.7	448.5
Family activities	4,700	3.7%	57.8	66.8	222.2	561.2
Fitness	8,377	6.5%	48.0	30.3	139.4	199.9
Goods	14,994	11.6%	40.9	48.1	394.9	1,434.0
Home and automobile services	16,830	13.1%	65.4	55.4	144.7	485.1
Life skill classes	7,262	5.6%	69.9	51.3	97.1	206.0
Live events	6,190	4.8%	29.4	25.5	419.5	2,783.4
Outdoor activities	9,083	7.1%	67.7	64.1	226.2	476.8
Personal care	8,838	6.9%	58.9	38.2	151.8	215.3
Restaurants	20,226	15.7%	22.5	24.0	456.7	535.2
Sports	3,371	2.6%	55.1	53.2	224.3	275.2
Travel	4,221	3.3%	121.9	94.4	197.8	334.5
<i>LivingSocial</i>						
Beauty	12,562	18.1%	67.9	54.9	144.4	992.0
Family activities	7,927	11.4%	56.9	52.1	150.1	879.9
Fitness	7,524	10.9%	36.5	22.4	123.1	225.1

Goods	3,292	4.8%	31.9	50.8	1,577.1	26,580.4
Home and automobile services	9,597	13.8%	60.9	51.5	163.2	667.5
Life skill classes	3,893	5.6%	54.0	46.5	152.2	274.9
Live events	4,148	6.0%	36.0	39.7	367.0	967.6
Outdoor activities	3,270	4.7%	56.1	59.1	460.6	1,079.0
Personal care	4,288	6.2%	51.0	30.9	179.7	273.9
Restaurants	10,763	15.5%	20.3	28.2	451.0	665.3
Sports	1,249	1.8%	42.5	40.7	267.0	419.2
Travel	827	1.2%	56.0	69.2	439.4	793.8

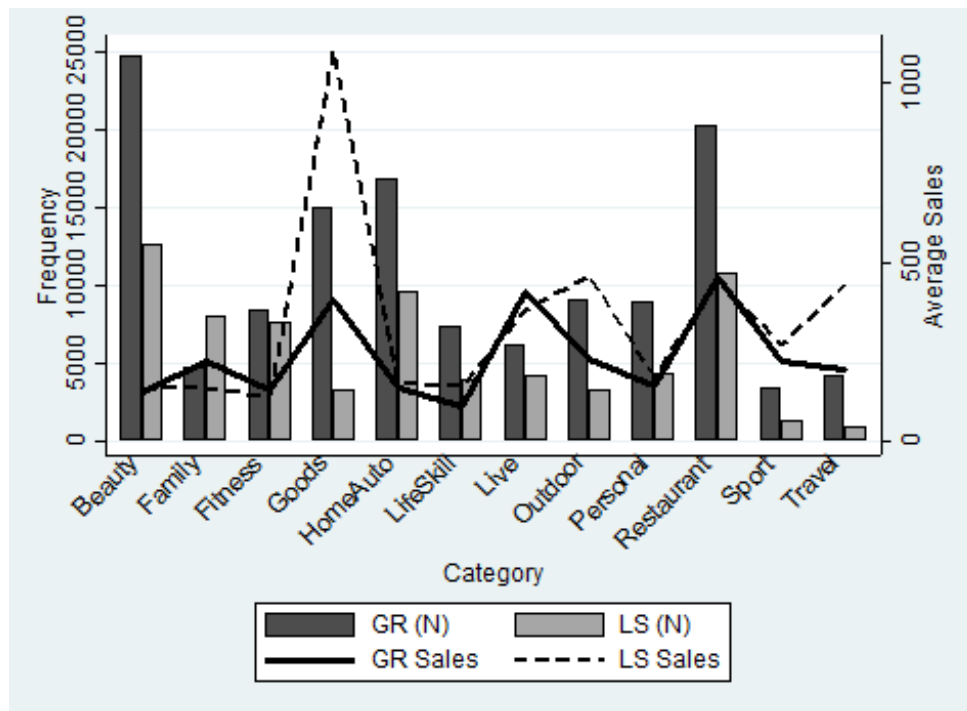


Figure 1.2: Quantity and sales of deals on platforms. The bars correspond to the number of deals by category for Groupon (“GR”) and LivingSocial (“LS”). The lines depict the average sales quantity per category on each platform.

### 3.3 Market definition and platform shares

We acquired platform market usage data from two premium data sources that capture Web-browsing behaviors for Internet users across the U.S. From Compete—the industry’s largest

consumer behavior database that updates daily clickstream data based on a panel of 2.3 million U.S. consumers—we obtained the number of unique visitors to [www.Groupon.com](http://www.Groupon.com) and [www.LivingSocial.com](http://www.LivingSocial.com) for each month of 2012.<sup>7</sup> Compete data also provide the number of unique visitors who visited both sites, which was important for this study. From the comScore Media Metric database, which has a representative U.S. consumer panel of roughly 47,000 members, we retrieve the geographical distribution of active users of Groupon, LivingSocial, and both. Combining these two data components, we computed the number of active users for each platform per market per month. We used these numbers to define the aggregate platform choices in the subsequent analysis.

Groupon and LivingSocial divide the U.S. market into so-called “divisions” which largely correspond to the metropolitan statistical areas (MSAs) defined by the Office of Management and Budget.<sup>8</sup> A typical MSA is centered around a large city that has economic influence over a region. For example, the “Chicago-Naperville-Joliet, IL-IN-WI” MSA surrounds Chicago and includes areas in Indiana and Wisconsin. In the context of our data, Groupon served 156 markets and LivingSocial served 166, with 131 served by both.

For each market, our analysis requires the “market size” for platform choices; that is, the total number of users who could possibly use one or both deal platforms. Potentially, any user with Internet access can use a deal site. Therefore, we use the number of Internet users to define the size of each market. The data are retrieved from the “October 2012 School Enrollment and

---

<sup>7</sup> In our data-collection period, mobile usage was very limited for daily deal business, though it has since become an important channel. In 2014, more than 50% of the transactions on Groupon were completed on mobile devices (*Groupon 10-K form 2014*).

<sup>8</sup> The Office of Management and Budget divides the US into 388 MSAs.

Internet Use Survey,” a supplement of the Current Population Survey (CPS) by the U.S. Census Bureau.

Groupon’s and LivingSocial’s market shares are computed by combining data from the monthly platform-level usage data and the distribution of users across regions. Table 1.3 summarizes these data. To construct the measures, we make the following calculation: (1) From comScore data we obtain the distribution of active users across census regions for each platform. For example, roughly 15.4% of Groupon users are from the mid-Atlantic region. (2) Within each region we assume that the number of users of a particular deal platform is proportional to the number of Internet users. For example, because Internet users in New York City make up 17.3% of the mid-Atlantic region total, the number of Groupon users in New York City is calculated as 17.3% of the number of Groupon users in the mid-Atlantic region. (3) Combining the distributions from steps (1) and (2) with the number of active Groupon users in a given month—e.g., 18 million—we calculate the number of active Groupon users in New York City in that month as  $17.3\% \times 15.4\% \times 18 \text{ million} \approx 480,000$ . Dividing these numbers by the market size gives us the market share for each platform choice in a market.

Table 1.3: Platform Shares by Month and Census region.

	User distribution across platform choices per month			User distribution across platform choices per region			
	Groupon only	LivingSocial only	Multi- homing	Groupon only	LivingSocial only	Multi- homing	
January	70.9%	14.7%	14.4%	region 1: New England	5.3%	11.1%	10.4%
February	54.6%	16.3%	13.4%	region 2: Mid- Atlantic	15.4%	13.9%	12.5%
March	57.9%	19.2%	15.5%	region 3: East North Central	15.5%	11.1%	12.5%
April	54.6%	21.6%	15.1%	region 4: West North Central	6.1%	4.2%	6.3%

May	55.6%	24.2%	16.0%	region 5: South Atlantic	23.0%	19.4%	14.6%
June	56.1%	20.2%	15.5%	region 6: East South Central	2.3%	5.6%	6.3%
July	60.3%	22.5%	15.4%	region 7: West South Central	7.7%	8.3%	12.5%
August	57.8%	20.1%	14.6%	region 8: Mountain	6.4%	9.7%	12.5%
September	46.2%	20.5%	14.6%	region 9: Pacific	18.2%	16.7%	12.5%
October	50.5%	21.4%	14.4%	<i>Total</i>	100%	100%	100%
November	45.8%	26.5%	13.1%				
December	45.3%	21.4%	13.4%				
<i>Average</i>	54.6%	20.7%	14.6%				

Note: The number of unique visitors for Groupon, LivingSocial, and both sites were acquired from Compete, Inc. The numbers in the top panel are the percentages of active users in each month of 2012. The numbers in the bottom panel are the percentages of active users across US census regions for each platform choice. We acquired these data from comScore, Inc.

During our data-collection period, approximately 6.5% of the Internet users exclusively used Groupon, 2.5% exclusively used LivingSocial, and 1.7% used both. The remaining 89.3% chose the “outside option”: either they purchased daily deals from other platforms or they did not participate in this market.

It is noteworthy that our platform market shares are based on “active users”—visitors to one or both platforms during our data-collection period—who may be a subset of the subscribers who have signed up to receive email alerts. We consider active users to be a better measure of platform size than subscribers because the former better represents the pool of users who actively consider deal offers. A subscriber may use an inactive email account to sign up and not truly be a

platform user. Indeed, Groupon’s 2012 annual report stated that retaining active users was its strategic emphasis<sup>9</sup>.

### **3.4 Other variables**

To estimate consumer heterogeneity in price sensitivity, we collected data on consumer characteristics. These data are also from the October 2012 CPS, which provides the empirical distributions of demographic and socioeconomic variables (such as income and household size) for Internet users in each market.

We obtained the merchant profile data from OneSource, one of the most comprehensive providers of business and company data. For each merchant, we know the number of employees, the annual sales, and whether it belongs to a chain.

### **3.5 Model-free evidence**

Here we provide some model-free evidence that motivates our model predictions. In this setting, the key ingredients of the pricing decision—the transferring of payment from the platform to the merchant as well as the allocation of bargaining power between them—are unobservable to researchers. Hence, it is challenging to find direct evidence to indicate their effect size. However, we can still look into how merchants and deals differ between Groupon and LivingSocial, and then use the observed differences to conjecture what might be happening in the background decision making. Our next step is then to specify a structural model that rationalizes the observed behavior in a way consistent with the data.

We run some thought experiments to predict what platforms may work better for what kind of merchants. First of all, if the bigger and smaller platforms only differ in their size, all merchants would want to work with the larger platform. The fact that the smaller one also attracts a decent number of merchants suggests that LivingSocial must be more preferable than

---

<sup>9</sup> *Groupon 10-K form 2012*

Groupon in some characteristics. One immediate explanation would be that LivingSocial may charge a lower commission rate; in other words, the transaction cost of using LivingSocial is lower from merchants’ perspective. If this happens, we would expect merchants to use Groupon more when they are able to generate more sales than on LivingSocial. And it would be more so for larger merchants, because they are more likely to sell more than smaller merchants.

Indeed, this is what we see in the data. We obtain a subset of data for which merchants’ characteristics are available (N=17,470). We focus on three merchant characteristics—size as measured in the number of employees, whether the merchant belongs to a chain business, and the amount of annual sales for the merchant. Table 1.4 provides the summary statistics. We see that merchants that are larger, belong to a chain, and have higher annual sales tend to work more with Groupon than with LivingSocial. The average discount level is higher on Groupon (0.61) than on LivingSocial (0.59), also suggesting that higher sales are necessary to make it worthwhile to endure the higher transaction cost on Groupon.

Table 1.4: Merchant characteristics on Groupon and LivingSocial.

	Groupon (N=11,158)		LivingSocial (N=6,312)		ttest
	Mean	SD	Mean	SD	
Merchant size	2.1	1.3	1.9	1.3	<.001
Chain	0.11	0.31	0.08	0.28	<.001
Annual Sales (Million \$)	2.8	15.8	2.6	25.6	<.001
Price (\$)	58.8	60.0	46.5	46.0	<.001
Value (\$)	212.4	355.2	145.2	192.8	<.001
Discount	0.61	0.13	0.59	0.13	<.001

Note: The variables *merchant size* and *chain* measure the number of employees per merchant and whether it belongs to a chain, respectively. Currently, the merchant-level dataset has 17,470 observations. The smaller sample size is due to data-processing constraints; we are working on data cleaning and matching to increase the sample size.



The data pattern seems to suggest that smaller merchants work with LivingSocial to take advantage of its lack of market power, but larger merchants still benefit more from Groupon due to its platform size. However, ex ante, larger merchants can still prefer LivingSocial if they could have a higher leverage on price setting, either because they have higher bargaining power there or because they are more valuable to LivingSocial than to Groupon in growing the platform size. If this happens, the platform size and transaction cost alone would not be sufficient to explain the differentiation between Groupon and LivingSocial. The bargaining power and the merchants' value to platforms, however, are not directly observable to researchers. Therefore, we need a structural model to examine how various factors determine prices and profit split in this setting.

## **4 Model**

In this section, we model consumer choices and the price decision for deals. Because estimating the supply-side parameters takes the demand-side parameters as input, we first describe our demand specification and then present the supply-side model.

### **4.1 Demand**

In the daily deals setting, a consumer follows a two-stage process: first, she chooses which platform(s) to use; second, given the choice of the platform(s), she considers which deals to purchase. This nested structure is similar to how consumers choose intermediaries in vertical markets, such as choosing an insurance policy and then selecting a health care provider in the network (Ho 2006). We present the model for deal demand followed by that for platform choices.

#### **4.1.1 Deal demand**

Consumer  $i$  derives utility from deal  $j$  that belongs to category  $c$  on platform  $k$  in market  $m$  during time  $t$ . Her utility is specified as

$$u_{ijkmt} = \alpha_i + \alpha_i^p p_{jkmt} + \beta x_{jkmt} + \xi_{jkmt} + \varepsilon_{ijkmt} , \quad (1)$$

where  $\alpha_i$  and  $\alpha_i^p$  are individual-specific deal preference and price parameter;  $x_{jkmt}$  is observable deal characteristics;  $\xi_{jkmt}$  is deal-specific shocks unobservable to the econometrician but observable to the consumer, platform, and merchant, and  $\varepsilon_{ijkmt}$  is the idiosyncratic utility shock. For ease of exposition, we omit the subscripts for categories, markets, and time, and use only the primary subscripts to index deals.

The taste parameters for overall deal preference and prices are allowed to be individual-specific and specified as a function of observable and unobservable individual characteristics:

$$\begin{aligned} \alpha_i &= \alpha_c + \varphi D_i + \sigma v_i \\ \alpha_i^p &= \alpha_c^p + \varphi^p D_i + \sigma^p v_i^p \end{aligned} , \quad (2)$$

where  $\alpha_c$  and  $\alpha_c^p$  are the grand means for the category-specific preference and price sensitivity, respectively. Note that both of these are specified to be category specific, so that we capture the differences in demand and elasticity among categories.  $D_i$  are the observable individual-level socio-demographic variables and are assumed to follow an empirical distribution  $D_i \sim F_i(D_i)$ ;  $\varphi$  are the deviation from the mean preference that is attributable to  $D_i$ ;  $v_i$ 's are individual-specific idiosyncratic shocks and are assumed to follow a multivariate normal distribution  $N(\mathbf{0}, \mathbf{I})$ ; and  $\sigma$ 's capture the degree of preference heterogeneity related to  $v_i$ .

Plugging the individual-specific parameters into the deal utility, we get

$$\begin{aligned} u_{ijk} &= \alpha_c + \alpha_c^p p_{jk} + \beta x_j + \xi_j \\ &\quad + (\varphi D_i + \sigma v_i) + (\varphi^p D_i + \sigma^p v_i^p) p_{jk} + \varepsilon_{ijk} . \end{aligned} \quad (3)$$

We rewrite the utility as the sum of three components: the grand mean utility across all individuals,  $\delta_{jk} = \alpha_c + \alpha_c^p p_{jk} + \beta x_j + \xi_j$ ; the individual deviation from the grand mean,

$\mu_{ijk} = (\varphi D_i + \sigma v_i) + (\varphi^p D_i + \sigma^p v_i^p) p_{jk}$  ; and an idiosyncratic shock  $\varepsilon_{ijk}$  . Our deal-demand specification closely follows the aggregate random-coefficient discrete choice model developed by Berry, Levinsohn, and Pakes (1995, i.e., the “BLP model). Following its notation, we refer to  $\theta_1 = (\alpha_c, \alpha_c^p, \beta)$  as the linear parameters and  $\theta_2 = (\varphi, \varphi^p, \sigma, \sigma^p)$  as the nonlinear parameters.

We make two assumptions concerning deal choices. First, an individual chooses up to one deal per category during each month from the platform(s) on which she is active<sup>10</sup>; second, she treats different categories independently. These assumptions help capture the competition of deals within the same category but avoid assuming different categories as complements or substitutes. For example, during a particular month,  $i$  purchases one deal from the restaurant category and one deal from the auto service category. The two purchases are assumed independent of each other. For the other ten categories, she chooses not to purchase any deals, yielding the outside option, which can be understood as the best alternative to purchasing the deal. The utility of the outside option is defined as  $u_{i0ckmt} = \delta_0 + \varepsilon_{i0ckmt}$ , where  $\delta_0$  is a constant that sets the utility scale.

We assume that  $\varepsilon_{ijckmt}$  are independently and identically distributed (i.i.d.) from a type I extreme value distribution. With our assumptions, the share of consumers purchasing deal  $j$  is given as

$$s_{jk} = \int \frac{\exp(\delta_{jk} + \mu_{ijk})}{\sum_{j' \in J_{ckmt}} \exp(\delta_{j'k} + \mu_{ij'k}) + \exp(\delta_0)} dF(D_i, v_i), \quad (4)$$

---

<sup>10</sup> The average sales per deal are rather small relative to the platform’s user base. Hence we consider it innocuous to assume that a consumer buys at most one deal per category per month.

where  $J_{ckmt}$  is the collection of all deals belonging to category  $c$  in market  $m$  in time  $t$ . The deal-level market shares correspond to our observed data and hence form the basis for estimation.

Note that deal prices are likely to be determined endogenously, as deals with positive (negative) demand shocks may sell at a higher (lower) price. In our specification, this means that  $p_{jk}$  and  $\xi_j$  are not independent; we therefore need instruments for identification, which we discuss in Section 5.1.1.

#### **4.1.2 Platform choices**

Next, we model a consumer's decision to choose platform(s). Three main considerations underline our model formulation.

First, we assume that consumers choose which platform they want to be active on at the beginning of each month. This assumption is necessary because our data on platform shares are observed at a monthly level.

Second, at the moment of choosing platform(s), a consumer has not realized the idiosyncratic demand shocks for deals. Therefore, she forms an expectation on the utility that may be derived from each platform.

Third, consumers can single-home or multi-home: in our empirical setting, some only used Groupon, some only used LivingSocial, and some used both. Our model incorporates this flexibility and does not treat platforms as mutually exclusive options. Instead, we regroup platform choices so that each consumer may fall into one and only one of these four groups: Groupon only, LivingSocial only, multi-homing, and neither (the outside option). This coding scheme allows us to cast the platform decision under the discrete-choice model framework and take advantage of the closed-form formulation that such a model entails.

A consumer's ex-ante expected utility for a category equals the expected maximum utility across all the deals in that category, given by  $EU_{ickmt} = E_{\varepsilon}(\max_{j \in J_{ckmt}} (u_{ijk}))$ . Assuming i.i.d. type I extreme-value distribution for  $\varepsilon_{ijkmt}$ , the expected utility becomes

$$EU_{ickmt} = \log \left( \sum_{j \in J_{ckmt}} \exp(\delta_{jk} + \mu_{ijk}) \right), \quad (5)$$

where the log-sum form is the logit inclusive value of category  $c$  and represents the expected utility for the choice of deals within that category as opposed to holding the outside option.

As aforementioned, we assume that a consumer can choose either, both, or neither of the platforms. Let  $r \in R \equiv \{g, l, gl, 0\}$  denote the set of platform choices. A consumer's choice is coded as  $r = g$  if she uses only Groupon,  $r = l$  if only LivingSocial,  $r = gl$  if both, and  $r = 0$  if neither. The utility for platform choice  $r$  in market  $m$  during month  $t$  is

$$u_{irmt}^{pf} = \sum_{k \in r} \left( \sum_{c \in C_{kmt}} \gamma_c EU_{ickmt} \right) + \omega_t + \eta_{rm} + \Delta \eta_{rmt} + \varepsilon_{irmt}^{pf}. \quad (6)$$

The first term in Equation (6) captures the total expected utility across all available categories for platform set  $r$ , where  $\gamma_c$  is the taste parameter for deal category  $c$ .  $\omega_t$  is the fixed effect for month  $t$  that captures the time-specific shocks at the industry level (for example, mass media may broadcast stories on daily deals that boost (or diminish) consumers' overall interest in this market);  $\eta_{rm}$  represents the time-invariant fixed effects that capture the overall preference towards option  $r$  across consumers in market  $m$ ;  $\Delta \eta_{rmt}$  is the time-specific deviation from  $\eta_{rm}$ ; and  $\varepsilon_{irmt}^{pf}$  represents the idiosyncratic demand shocks specific to individual, platform, market, and time. A consumer chooses whichever set  $r$  maximizes her utility. We define the outside option

as an individual choosing neither platform ( $r = 0$ ) and scale the utility by restricting the outside option utility as  $u_{i0mt}^{pf} = 0 + \varepsilon_{i0mt}^{pf}$ .

The fixed effects for platform and market,  $\eta_{rm}$ , represent the market-level mean time-invariant set-specific value (independent of the deals being offered) net of the cost associated with using the platform(s) in set  $r$ . This value could be a manifest of things, including but not restricted to each platform's reputation and the quality of its customer services, such as shipping speed and return policy. There could also be search cost or other nonmonetary costs of using deal platforms; for example, the disutility of having to deal with the multiple daily email alerts that deal platforms typically send out. Without the fixed effects, one would expect consumers to always multi-home, as more deal options would always yield higher expected total utility. In reality, however, many consumers single-home, suggesting that there is a cost for consumers to consider multiple platforms.

Again, under the assumption that  $\varepsilon_{irmt}^{pf}$  is i.i.d. from a type I extreme value distribution, the market share for set  $r$  becomes

$$s_{rmt}^{pf} = \int s_{irmt}^{pf} dF(D_i, v_i) . \quad (7)$$

## 4.2 Supply model

Several considerations underline the key features of our supply-side model.

First, there are multiple pieces involved in the pricing decision between platforms and merchants. In this empirical setting, platforms charge a commission fee for facilitating the sales, which is a linear rate per sale. This essentially can be seen as the transaction cost from the merchant's perspective. The final price of the deal is what consumers pay to buy the voucher, which is equivalent to the discount level given the voucher's face value. This price, however, is determined through a negotiation between the merchant and the platform's salesforce. Groupon

and LivingSocial employ a large number of sales people to recruit merchants and negotiate terms with them on a deal-by-deal basis. Both merchants and platforms are incentivized to influence the price to favor their business objectives; but neither has full discretion on price setting in this market.

Second, a deal contributes to a platform's profits not only through generating revenues on the current deal sales but also through growing the platform's customer base and influencing sales of other deals on the platform. Hence, the platform recognizes the network effect of each deal and internalizes it in its pricing decision. In other words, a platform maximizes its total platform revenue rather than just the single deal revenue.

Third, a merchant works with a platform not only because it wants to generate some revenues from selling deals but also because it hopes to grow its business by retaining some of the customers acquired through the deal promotion. Therefore, a merchant internalizes both the current and future revenues in its pricing decision.

Given those considerations, we formally model the outcome of a price negotiation as the equilibrium of a bilateral Nash bargaining problem in the sense that neither the platform nor the merchant wants to deviate from the determined price. The Nash model, developed by Horn and Wolinsky (1988), has become the workhorse for empirical work on bilateral negotiations. In our application, the prices maximize the Nash product of the payoffs to the platform and to the merchant with an agreement relative to the payoffs without an agreement. That is, a deal price solves

$$\max_{p_{jk}} \left[ \pi_{kmt}(p_{mt}) - d_{jkmt} \right]^{b_k(j)} \left[ q_{jk}(p_{jk})(h_{jk} - c_j) + \sum_{t=1}^{\infty} \lambda^t q_{jk}(p_{jk})(FV_j - c_j) \right]^{b_j(k)} \quad \forall j \in J_{kmt}, \quad (8)$$

where inside each bracket is the payoff with and without an agreement for the platform and merchant, respective. For platform  $k$ , it receives the total market-level profits  $\pi_{kmt}$  if an agreement is reached between  $j$  and  $k$ , versus a disagreement payoff  $d_{jkmt}$  if an agreement is not reached. The platform's profits depend on the deal demand and margin<sup>11</sup> across all the deals it offers in the market:  $\pi_{kmt} = \sum_{j \in J_{kmt}} q_{jk}(p_{jk})(p_{jk} - h_{jk})$ , where  $h_{jk}$  is the payment made by the platform to the merchant. For the merchant offering deal  $j$ , its payoff consists of the revenue from selling the deal,  $q(p_{jk})(h_{jk} - c_j)$ , and a future flow of revenues from the acquired customers,  $\sum_{t=1}^{\infty} \lambda^t q(p_{jk})(FV_j - c_j) = (1 - \lambda)^{-1} q(p_{jk})(FV_j - c_j)$ , where  $c_j$  is the merchant's marginal cost of serving a customer. We assume that a constant fraction of the acquired customers would return repeatedly in the future, captured by  $\lambda$ , and that on average they consume goods or services worth of the face value of the deal. Note that this is a fairly restrictive assumption: we do not allow  $\lambda$  to vary by merchants hence miss the heterogeneity in their ability to retain customers; nor do we capture the real consumption value for the acquired customers. Without detailed transaction data from each merchant, these parsimonious assumptions are necessary. Despite the limitations, this simplification provides a means to capture the future payoff that the merchant internalizes when negotiating the price with the platform.

In the bargaining literature, there are multiple ways to capture the disagreement payoffs. Following Horn and Wolinsky (1988), we assume that other contracts—possibly including those

---

<sup>11</sup> Notice that we assume the platform have zero marginal cost in selling an additional deal; this is reasonable because deal sites operate online.



between the platform and other merchants—would not be renegotiated if platform  $k$  and merchant  $j$  do not reach an agreement. The platform’s disagreement payoff thus becomes  $d_{jkm} = \pi_{kmt}(\mathbf{p}_{mt}; J_{kmt} \setminus \{j\})$ ; that is, the profits for platform  $k$  in market  $m$  during time  $t$  given the prices of all remaining deals.

Parameter  $b_j(k) \geq 0$  is the price-bargaining power of merchant  $j$  when facing platform  $k$ , and  $b_k(j) \geq 0$  is  $k$ ’s bargaining power when facing  $j$ . Bargaining parameters are not separately identifiable, hence we normalize them by  $b_k(j) + b_j(k) = 1$ . If  $b_k(j) = 1$ , the platform sets the price and the merchant uses a take-it-or-leave-it strategy. Vice versa for  $b_j(k) = 1$ . Hence, this Nash bargaining model nests the Bertrand pricing model as a special case.

We solve the first-order condition (FOC) of this Nash bargaining problem and obtain the following pricing equation:

$$p = h + \frac{b_k}{b_k + b_j} \frac{q}{-q'} (1 - h') + \frac{b_j}{b_k + b_j} \left[ (p - h) - (\pi - d) \left( \frac{1}{q} + \frac{1}{q'} \frac{h'}{h - c + \frac{1}{1 - \lambda} (FV - c)} \right) \right]. \quad (9)$$

Equation (9) shows that the equilibrium price equals the payment transfer from the platform to the merchant plus a weighted average of two terms. The first is the platform’s Bertrand-Nash best-response markup, weighted by the platform’s relative bargaining power. Therefore, it is obvious that, if platforms have full control over price setting, the equilibrium price becomes the platform’s Bertrand-Nash best-response price. The second term, capturing the departure from the platform’s most preferred markup, contains (1) the merchant’s relative bargaining power, (2) its

“externality value” to the platform,  $\pi_{kmt} - d_{jkmt}$ , and (3) its future average markup,  $(1 - \lambda)^{-1}(FV_j - c_j)$ .

To better understand the equilibrium properties, we re-write the pricing equation as:

$$p - h = \frac{q}{-q'}(1 - h') + \frac{b_j}{b_k}(\pi - d) \left[ \frac{1}{-q'} \frac{h'}{h - c + \frac{1}{1 - \lambda}(FV - c)} - \frac{1}{q} \right]. \quad (10)$$

We see that the equilibrium price is pulled closer to the merchant’s preferred price and further away from the platform’s preferred price, when the merchant has higher relatively bargaining power,  $b_j / b_k$ , or a higher externality value,  $\pi_{kmt} - d_{jkmt}$ . It’s worth noting that platforms do not always prefer a lower price than the merchants. The relationship between platforms’ and merchants’ best-response prices depends on the marginal cost of the merchant (see Appendix). Merchants with higher marginal costs would prefer a higher deal price than the platforms, and merchants with lower marginal costs would benefit from a price lower than the platforms’ best-response price. Interestingly, if the proportion of future versus current payoff for the merchant is bigger, i.e., larger  $(1 - \lambda)^{-1}(FV_j - c_j)$ , then the merchant would always prefer a lower price as to gain a higher future traffic; in this case, the equilibrium price would necessarily be driven down as to generate the sales-boosting effect.

Note that  $h_{jk}$  is unobservable to researchers. Per industry practice, we model it as proportional to deal price:  $h_{jk} = \kappa_k p_{jk}$ , where  $1 - \kappa_k$  is the platform’s commission rate and will be estimated to be platform specific.

After regrouping the terms, we rewrite Equation (10) and further parameterize the price-bargaining ratio as a function of observable platform and merchant characteristics,  $\chi_{jkmt}$ , and the unobservable,  $\varsigma_{jkmt}$ :

$$g(p_{jk}; \kappa) = \frac{b_{jmt}(k)}{b_{kmt}(j)} = \chi_{jkmt} + \varsigma_{jkmt}, \quad (11)$$

where the left-hand side of the equation,  $g(p_{jk}; \kappa) = g_1 / g_2$ , where

$$g_1 = p_{jk} - h_{jk} + q_{jk} \cdot (q_{jk}')^{-1} (1 - h_{jk}') \quad \text{and} \quad g_2 = (\pi_k - d_{jk}) \left[ \frac{1}{-q_{jk}' h_{jk} - c_j + (1 - \lambda)^{-1} (FV_j - c_j)} - \frac{1}{q_{jk}} \right]$$

can be constructed as a function of data and the estimated demand structural parameters. We describe the choice of observable characteristics,  $\chi_{jkmt}$ , and other estimation details in Section 5.2.

## 5 Estimation, Identification, and Computation

In this section, we present our estimation strategy, discuss parameter identification, and provide details on the computation.

### 5.1 Estimation of the demand-side parameters

We adopt the BLP method to address price endogeneity and incorporate consumer heterogeneity in deal preference and price sensitivity. The parameters are estimated by minimizing an objective function based on a set of moment conditions as defined in the generalized method of moments (GMM) (Hansen 1982).

#### 5.1.1 Deal-demand estimation

We begin by describing the variables used in the deal-demand specification. The vector of observable deal characteristics,  $x_j$ , includes price, the voucher's face value, the month in which the deal was offered (to capture any seasonal effect), deal category, and the size of the market

(dummy variables indicating top-20 markets, markets ranked 21 to 40 markets, and otherwise). We use the logarithm of prices in the estimation to address the skewness in this variable. Deals may be substantially different even within the same category. For example, a ticket package to a premium children’s play, such as “How to Train Your Dragon” at the IZOD Center, is priced around \$80-90, while a fine play like “Sesame Street Live: Can’t Stop Singing” typically have a face value around \$30. We include the voucher’s face value to at least partially control for deal heterogeneity.

For individual characteristics,  $D_i$ , we include annual income, household size, and age. As in Equation (2), we allow the overall deal preference and price sensitivity to depend on those individual characteristics. We simulate the values for each variable based on its empirical distribution.

When estimating the price parameter,  $\alpha_c^p$ , we need to account for a nonzero correlation between  $p_j$  and  $\xi_j$ . Because a deal with higher demand shocks,  $\xi_j$ , may cost more but still end up with higher sales, failing to account for endogeneity would bias the price estimate towards zero. A valid price instrument should be correlated with  $p_j$  but exogenous to  $\xi_j$ . We choose as price instruments (a) the average price of all the deals from the same category in other markets during the same month on the focal platform and (b) the same average for the other platform. These instruments are similar to those used in (Hausman 1996; Nevo 2001). Because the instruments are averaged across deals of the same category around the same time, they should be correlated with  $p_j$ , due to common cost shifters at the category level. Because the averages are based on deals from other markets, it is reasonable to assume that the price instruments are uncorrelated with the demand shocks in the focal market. We set the restriction criteria as

$E(Z \cdot \xi_j) = 0$ . Note that these instruments would be invalid if they were only weakly correlated with the focal deal's price (causing weak-instrument problems) or if the unobservable demand shocks were correlated across markets (violating the exogenous requirement). We provide diagnostic statistics for the instruments in the results section.

### 5.1.2 Platform demand estimation

Equation (6) specifies the total utility that a consumer expects to derive from each platform set. We further use  $\Gamma_{ikmt}(\delta(\theta_1); \gamma, \theta_2) = \sum_{c \in C_k} \gamma_c EU_{ickmt}$  to denote the part of the utility directly related to deals being offered. Here,  $\gamma$  is the vector of category-specific taste preferences,  $\delta$  is the vector of deal mean utilities, and  $\theta_2$  is the vector of nonlinear utility parameters in the deal demand.

After plugging in  $\Gamma_{ikmt}(\delta(\theta_1); \gamma, \theta_2)$  and regrouping terms, we write the aggregated market shares for platform sets as

$$\begin{aligned}
s_{r,mt}^{pf} &= \int \frac{\exp(\Gamma_{irmt}(\delta(\theta_1); \gamma, \theta_2) + \omega_i + \eta_{rm} + \Delta\eta_{rmt})}{1 + \sum_{r=\{g,l,gl\}}} dF(D_i, v_i) \\
&= \int \frac{\exp(\Gamma_{irmt}(\delta(\theta_1); \gamma, \theta_2) + \delta_{rmt}^{pf})}{1 + \sum_{r=\{g,l,gl\}}} dF(D_i, v_i) \quad , \text{ for } r = \{g, l\} \text{ and} \\
s_{gl,mt}^{pf} &= \int \frac{\exp(\Gamma_{i,g,mt}(\delta(\theta_1); \gamma, \theta_2) + \Gamma_{i,l,mt}(\delta(\theta_1); \gamma, \theta_2) + \omega_i + \eta_{gl,m} + \Delta\eta_{gl,mt})}{1 + \sum_{r=\{g,l,gl\}}} dF(D_i, v_i) \\
&= \int \frac{\exp(\Gamma_{i,g,mt}(\delta(\theta_1); \gamma, \theta_2) + \Gamma_{i,l,mt}(\delta(\theta_1); \gamma, \theta_2) + \delta_{gl,mt}^{pf})}{1 + \sum_{r=\{g,l,gl\}}} dF(D_i, v_i)
\end{aligned}$$

where  $\sum_{r=\{g,l,gl\}} = \exp(\Gamma_{i,g,mt} + \delta_{g,mt}^{pf}) + \exp(\Gamma_{i,l,mt} + \delta_{l,mt}^{pf}) + \exp(\Gamma_{i,g,mt} + \Gamma_{i,l,mt} + \delta_{gl,mt}^{pf})$ .

The linear component of the aggregated platform shares becomes

$$\begin{aligned}
\delta_{g,mt}^{pf} &= \omega_t + \eta_{g,m} + \Delta\eta_{g,mt} \\
\delta_{l,mt}^{pf} &= \omega_t + \eta_{l,m} + \Delta\eta_{l,mt} \quad . \\
\delta_{gl,mt}^{pf} &= \omega_t + \eta_{gl,m} + \Delta\eta_{gl,mt}
\end{aligned} \tag{12}$$

Here, we are also concerned with potential endogeneity problems. A popular platform may offer more and better deals, introducing a nonzero correlation between  $\eta_{rm}$  and  $\Delta\eta_{rmt}$ . To address this concern, we use the within-group fixed-effect estimator and use the first-differences transformation to eliminate the fixed effects. After the transformation, Equation (14) becomes

$$\begin{aligned}
\delta_{rm,t}^{pf} - \delta_{rm,t-1}^{pf} &= (\omega_t - \omega_{t-1}) + (\Delta\eta_{rm,t} - \Delta\eta_{rm,t-1}) \\
D\delta_{rm,t}^{pf} &= D\omega_t + D\Delta\eta_{rm,t} \quad ,
\end{aligned} \tag{13}$$

where  $D\delta_{rm,t}^{pf} = \delta_{rm,t}^{pf} - \delta_{rm,t-1}^{pf}$ ,  $D\omega_t = \omega_t - \omega_{t-1}$  and  $D\Delta\eta_{rm,t} = \Delta\eta_{rm,t} - \Delta\eta_{rm,t-1}$ . We then form the identification restriction as  $E(D\omega_t \cdot D\Delta\eta_{rmt}) = 0$ .

### 5.1.3 BLP computation

Generally perceived as a nested fixed-point (NFXP) algorithm, the BLP method incorporates a contraction mapping step in which one inverts the demand system to recover a vector of mean utility,  $\delta$ , that equates the predicted market shares with the observed market shares. In the BLP scheme, this contraction mapping step is an inner loop nested within an outer loop to search for the nonlinear utility using GMM.

Berry et al. (1995) prove that the fixed-point iteration used in the BLP scheme is guaranteed to converge. While this global convergence property is appealing, the BLP contraction mapping can be time-consuming, especially when the sample size exceeds 5,000. In order to speed up convergence, a common technique is to (a) relax the inner-loop tolerance value ( $\varepsilon_{in}$ ) in regions where the minimization of the GMM objective function is far from the true solution and (b) tighten the tolerance criterion as the minimization gets closer to the truth. However, this

procedure may lead to incorrect estimates, as Dube, Fox, and Su (2012) show that the inner-loop tolerance must be set at  $10^{-14}$  with the outer-loop tolerance at  $10^{-6}$ .

To accelerate the convergence without being penalized for estimation bias, we adopt the squared polynomial extrapolation method (SQUAREM), a state-of-the-art algorithm that can operate directly on the fixed-point formulation of the BLP contraction mapping. Originally developed to accelerate the expectation-maximization (EM) algorithm, SQUAREM has been shown to be not only faster but also more robust (in terms of the success rate of convergence) than the original contraction mapping procedure used in BLP (Reynaerts et al. 2012; Varadhan and Roland 2008). The advantage of SQUAREM is even more substantial when the sample size is large (as in our case) and when the initial values of the parameters are far from the truth.<sup>12</sup>

It is noteworthy that estimating the deal-demand parameters separately from the platform-demand parameters may also yield inaccurate estimates due to selection bias. Consumers may self-select onto different platforms depending on their preferences and the platforms may tailor their offerings accordingly, introducing another source of endogeneity. We therefore jointly estimate Equations (1) and (6) by iteratively solving for the deal-demand and the platform-demand parameters during the optimization search.

The details of the estimation routine are as follows:

1. For each market, simulate NS=300 individuals
  - a. with observable characteristics from empirical marginal distributions,  $F(D_i)$

---

<sup>12</sup> For example, in our application, one search for the vector  $\delta_j$  took 26 iterations and 3.5 minutes using the SQUAREM accelerator and over 5,000 iterations and 3 hours using the BLP contraction mapping with the inner-loop tolerance set at  $10^{-14}$ .

- b. with unobservable idiosyncratic shocks,  $v_i$ , simulated from a multivariate standard normal distribution
2. Assign initial values for  $\alpha$ ,  $\alpha^p$ , and  $\beta$  and calculate the initial value for  $\delta_j$ :
- $$\delta_j^{(0)} = \alpha_c + \alpha_c^p p_j + \beta x_j$$
3. Given  $\theta_2$  and  $\delta_j$ , predict the share for each deal
- a. given  $\theta_2 = (\pi, \pi^p, \sigma, \sigma^p)$ :  $\mu_{ij} = (\pi D_i + \sigma v_i) + (\pi^p D_i + \sigma^p v_i^p) p_j$
- b. given  $\delta_j$ :  $\hat{\sigma}_j(\delta_j, \mathbf{x}_j; \theta_2) = \frac{1}{NS} \sum_i \frac{\exp(\delta_j + \mu_{ij})}{\exp(u_0) + \exp(\delta_j + \mu_{ij})}$
4. Conduct BLP contraction mapping with SQUAREM accelerator to search for  $\delta_j$  such that  $\hat{\sigma}_j(\hat{\delta}_j, \mathbf{x}_j; \theta_2) = s_j$  as long as  $\|\delta_j^{(h+1)} - \delta_j^{(h)}\| < \varepsilon_{in}$ , where  $\varepsilon_{in}$  is the inner-loop tolerance set as  $10^{-14}$
5. Given  $\delta_j = \alpha_c + \alpha_c^p p_j + \beta x_j + \xi_j$ , analytically solve for the deal-demand linear parameters,  $\theta_1 = (\alpha_c, \alpha_c^p, \beta)$
6. Given  $\delta_j$ ,  $\theta_1$ ,  $\theta_2$ , and  $\gamma$ , compute  $\Gamma_{irmt}(\delta(\theta_1); \gamma, \theta_2) = \sum_{k \in r} \left( \sum_{c \in C_k} \gamma_c EU_{ickmt} \right)$ ,
7. Given  $\Gamma_{irmt}$ , compute the predicted platform shares
- $$\hat{\sigma}_{rmt}^{pf} = \frac{1}{NS} \sum_i \frac{\exp(\Gamma_{irmt}(\delta(\theta_1); \gamma_c, \theta_2) + \delta_{rmt}^{pf})}{1 + \sum_q \exp(\Gamma_{iqmt}(\delta(\theta_1); \gamma_c, \theta_2) + \delta_{qmt}^{pf})}$$
8. As in step 4, perform BLP contraction mapping with the SQUAREM accelerator to search for  $\delta_{rmt}^{pf}$  so that  $\hat{\sigma}_{rmt}^{pf} = s_{rmt}^{pf}$  as long as  $\left\| \left( \delta_{rmt}^{pf} \right)^{(h+1)} - \left( \delta_{rmt}^{pf} \right)^{(h)} \right\| < \varepsilon_{in}$



9. Form the GMM moment conditions based on  $E(Z \cdot \xi_j) = 0$  and  $E(D\omega_t \cdot D\Delta\eta_{rmt}) = 0$ , and repeat from step 3 for each iteration of the optimization.

## 5.2 Supply model estimation

The function  $g(p_{jkmt}; \kappa)$  requires the implied deal demand,  $q_{jkmt}$ , the platform's profits with the agreement,  $\pi_{kmt}$ , its disagreement payoffs,  $d_{jkmt}$ , and the price elasticity term,  $q'_{jkmt}$ , all of which are constructed based on our demand parameter estimates. We further parameterize the marginal cost of a deal to be a fixed proportion of its face value,  $c_j = \psi_c FV_j$ , where the proportion is category-specific, i.e., being held constant across deals from the same category. Parameter  $\chi_{jkmt}$  captures how the relative bargaining power depends on the merchant observables. We parametrize the merchant difference based on three important merchant characteristics: the number of employees, whether the merchant belongs to a chain (1=chain; 0=independent), and the annual sales for the merchant. We take the logarithm transformations for the continuous variables, such that our estimates are less influenced by extreme values. Again, we use GMM to solve for the parameters and we set the moment condition as  $E[Z^s \zeta_{jk}] = 0$ . The supply-side instruments,  $Z^s$ , are set to be  $\chi_{jkmt}$ , under the assumption that the observables are exogenous to  $\zeta_{jk}$  after controlling for all the included variables.

## 5.3 Identification

The linear parameters,  $\theta_1$ , are straightforwardly identified via the cross-sectional variation across deals. The nonlinear parameters,  $\theta_2$ , are identified through the variation in deals that have similar observables but end up with different sales quantities in markets with varying consumer characteristics. Imagine that two identical deals are offered in markets A and B, where market A

has higher average incomes, a larger average household size, and an older population. The difference in sales of the deals between the markets would help identify the random intercepts associated with those individual characteristics. Now suppose two deals, identical except for price, are again offered in markets A and B. We can then use the demand differences of these two deals between the markets to identify the effect of the characteristics on the price coefficient.

The taste parameters for deal categories,  $\gamma$ , are identified through the within-market variation in deal offers and platform market shares. In a given market, if the change in a platform's market share is positively and substantially associated with a change in its offerings in a particular category (e.g., restaurant deals) the taste parameter for that category would be estimated to be large.

The supply parameters include the commission rate charged by the platforms,  $\kappa_k$ , the marginal cost to face value ratios per category,  $\psi_c$ , and the retention rate,  $\lambda$ . Note that  $\psi_c$  and  $\lambda$  are not separately identifiable: if a merchant accepts a low price, it could be either because he has a low marginal cost of serving customers or because he is expecting a high retention rate. Without better data to solve this problem, we fix  $\lambda$  to be 0.25, based on interviews with merchants. The rest of the estimation is straightforward. Parameters  $\kappa_k$  are identified through the data variation between platforms. Imagine two deals have identical category, bargaining power, face value, and marginal cost, but one is offered on Groupon and the other on LivingSocial. If they charge different prices for consumers, the price difference would help identify the commission rate for the platforms. Note that restricting the commission rate to be constant across deals from the same platform helps separately identify it from the bargaining parameter. The marginal cost parameters,  $\psi_c$ , are identified through the differences in price-to-face-value ratio and price elasticity across categories. In a nutshell, imagine two deals that are from different

categories but otherwise identical (i.e., platform, face value, and bargaining power). Given the category-specific price elasticity estimated from the demand side, the price difference between the deals helps identify marginal cost parameters for each category.

## **6 Results**

### **6.1 Demand parameter estimates**

We examine several specifications of the deal demand and present the linear parameter estimates in Table 1.5. The first specification is a homogeneous logit model without accounting for price endogeneity or heterogeneity across individuals. This is simply the ordinary least squares (OLS) estimate with the dependent variable being the logarithm of the deal share minus the logarithm of the outside share. Results from specification (1) are used as benchmark values.

In the second specification, we use the Hausman-type price instruments discussed in Section 5.1.1, though we still do not account for consumer heterogeneity. With IV, the main effect for price was estimated to be much stronger: -4.619 with IV versus -0.781 without it. The direction of the change is as expected when prices and the unobservable demand shocks are positively correlated: when popular deals are priced high and unpopular ones are priced low, the OLS estimate of the price coefficient would be attenuated towards zero, as in our case. To assess the validity of the instruments, we run the first-stage regression and find the F statistic to be 1460.9 ( $p < 0.01$ ). We also run the Stock and Yogo (Stata 2013) test for weak instruments: our F statistic is higher than the test-critical value of 19.9, rejecting the null hypothesis of weak instruments.

Table 1.5: Linear parameter estimates for deal demand

Variable	(1) Homogeneous logit without IV		(2) Homogeneous logit with IV		(3) Random-coefficient logit with IV	
	Est	SE	Est	SE	Est	SE
Price	-0.781***	0.012	-4.619***	0.247	-7.833***	0.307
Price X Beauty	0.188***	0.016	-2.608***	0.309	-2.370***	0.360
Price X Family	0.120***	0.019	4.522***	0.308	4.066***	0.352
Price X Fitness	-0.170***	0.026	0.192	0.239	-0.086	0.300
Price X Goods	-0.360***	0.018	2.933***	0.253	4.589***	0.329
Price X LifeSkill	-0.072***	0.023	-2.007***	0.345	-2.628***	0.415
Price X LiveEvents	0.129***	0.023	4.630***	0.301	5.769***	0.370
Price X Outdoor	-0.092***	0.019	1.885***	0.254	2.127***	0.319
Price X Personal	-0.363***	0.029	-0.337	0.361	-1.008**	0.450
Price X Restaurants	0.069***	0.017	1.770***	0.254	2.698***	0.316
Price X Sports	0.089***	0.030	2.412***	0.291	2.366***	0.377
Price X Travel	-0.706***	0.024	1.873***	0.240	2.545***	0.296
Face value	0.004**	0.002	0.644***	0.032	0.698***	0.037
Beauty	-0.714***	0.063	10.977***	1.230	10.266***	1.431
Family	-0.901***	0.074	-17.692***	1.161	-16.638***	1.336
Fitness	0.544***	0.097	-1.622*	0.914	-1.113	1.144
Goods	-0.170***	0.064	-12.713***	0.952	-18.272***	1.221
LifeSkill	0.131	0.092	7.483***	1.339	9.321***	1.617
Live events	-0.419***	0.079	-16.590***	1.075	-20.853***	1.328
Outdoor	0.745***	0.073	-6.828***	0.970	-7.908***	1.221
Personal	1.862***	0.114	1.737	1.404	4.127**	1.752
Restaurants	0.527***	0.058	-7.472***	0.940	-10.393***	1.172
Sports	0.086	0.113	-8.906***	1.093	-9.429***	1.412
Travel	2.345***	0.101	-7.474***	0.936	-10.311***	1.154
Controls	Included		Included		Included	
N	198,089		198,089		198,089	

\*\*\* p<0.01; \*\* p<0.05; \* p<0.10

Specification (3) is the random-coefficient aggregate logit model that uses the price instruments and also incorporates individual preferences as a function of income, household size, and age. As expected, the mean price coefficient is estimated to be negative and significant ( $\hat{\alpha}^p = -7.833, p<0.01$ ). The corresponding random coefficient estimates are reported in Table 1.6. We find significant variation in price elasticity across individuals: people who are older, have a

higher income, or come from a larger household are significantly less price-sensitive for daily deals. After controlling for those consumer characteristics, there is still significant heterogeneity in price elasticity ( $\hat{\sigma}^p = 0.928$ ,  $p < 0.01$ ). The overall deal preference varies by household size, but not by income level or age: our estimates indicate that consumers from a larger household tend to like deals less than otherwise. The remaining consumer heterogeneity in deal preference is small but significant ( $\hat{\sigma} = 0.415$ ,  $p < 0.05$ ).

Table 1.6: Nonlinear parameter estimates for deal demand

Coefficient	$\sigma$		Income		Household size		Age	
	Est	SE	Est	SE	Est	SE	Est	SE
Price	0.928***	0.029	0.746***	0.237	0.596***	0.174	0.666***	0.220
Intercept	0.415**	0.162	-1.128	0.884	-1.177***	0.328	1.036	0.744

\*\*\*  $p < 0.01$ ; \*\*  $p < 0.05$ ; \*  $p < 0.10$

Using the estimated demand parameters, we compute the mean price elasticity for each category (Table 1.7). Consistent with the parameter estimates, we see substantial variation in price elasticity across categories. Consumers are the most price-sensitive to deals on life skill class, the mean price elasticity being -10.45 with the interquartile range of (-10.58, -10.31). The next two price-sensitive categories are beauty deals (-10.19) and personal care deals (-8.83). To put these numbers in perspective, the average price elasticity for consumer packaged goods is around -2.50 (Tellis 1988). Soft drinks are typically considered elastic goods: Coca-Cola has an elasticity of -3.8 while Mountain Dew's is -4.4 (Ayers and Collinge 2003). Alcoholic beverages typically have elasticity between -1.0 and -1.5. Among the 12 categories, only three categories would be considered as relatively price insensitive: live events (-2.07), goods (-3.25), and family activities (-3.77). The mean price elasticity for all other categories was all greater than 5.

Table 1.7: Price elasticity by categories

Category	Mean	25% Percentile	75% Percentile
Beauty	-10.19	-10.32	-10.05
Family	-3.77	-3.91	-3.64
Fitness	-7.91	-8.04	-7.78
Goods	-3.25	-3.39	-3.11
Home and auto	-7.83	-7.95	-7.69
Life skill	-10.45	-10.58	-10.31
Live events	-2.07	-2.21	-1.94
Outdoor	-5.71	-5.84	-5.57
Personal care	-8.83	-8.96	-8.70
Restaurants	-5.14	-5.27	-5.00
Sports	-5.47	-5.60	-5.34
Travel	-5.29	-5.42	-5.16

Note: Price elasticities are calculated using deal-demand estimates for each simulated individual.

We report the means and quartiles across the elasticity distribution.

Next, we discuss consumers’ preferences for different deal categories as they choose platforms. The higher the estimate for  $\gamma_c$ , the more a category is able to attract consumers to a platform. Our results reveal substantial heterogeneity across categories in their capacity to grow a platform’s customer base. We find that beauty deals (such as haircuts, hair removals, and facials) have the highest appeal (0.149,  $p < 0.01$ ), followed by deals on home and automobile services (0.080,  $p < 0.01$ ). Seven other categories—life skill classes, live events, outdoor activities, personal care, restaurants, sports, and travel activities—are also effective in growing a platform’s customer base. The remaining three categories—family activities, fitness, and goods—exert minimal influence on a consumer’s choice of a platform. In general, these categories tend to have fewer deals, lower sales, or both, which partially explains why they are ineffective in attracting users to a platform.

Table 1.8: Parameter estimates for platform choices

Category	Est	SE
Beauty	0.149***	0.003
Family	0.003	0.002
Fitness	0.001	0.003
Goods	0.001	0.002
Home and auto	0.080***	0.003
Life skill	0.015***	0.002
Live events	0.015***	0.002
Outdoor	0.020***	0.002
Personal care	0.039***	0.003
Restaurants	0.024***	0.002
Sports	0.006***	0.002
Travel	0.026***	0.002
Controls	Included	

\*\*\* p<0.01; \*\* p<0.05; \* p<0.10

## 6.2 Supply parameter estimates

In Table 1.9, we present the estimates for the supply parameters<sup>13</sup> specified in Equation (11). First, we estimate that, on average, LivingSocial pays 51.4% of the price to the merchant, in other words, charging an average commission rate of 48.6%. The commission rate charged by Groupon is higher by 9.8%: Groupon keeps the 58.4% of the deal price as its commission, and pays the merchant 41.6% of the price. The incumbent status and the size of Groupon's customer base perhaps help explain why they are able to demand a larger slice of the pie than LivingSocial. This also rationalizes why Groupon, being the first and larger platform in most markets, was not able to capture the entire market: although Groupon can perhaps yield more sales for merchants

<sup>13</sup> The current estimates are based on a subset of data (N=17,470 deals). The smaller sample size is due to the time-consuming process of cleaning merchant names to match those in the OneSource database. As we have continued to increase the sample size, the reported patterns have so far been duplicated.

than LivingSocial does, some merchants still benefit from working with LivingSocial because they are offered a larger cut of the pie.

Next, we interpret how the relative bargaining power depends on the platform's and the merchants' characteristics. Note that the ratio of bargaining power,  $b_{jmt}(k)/b_{kmt}(j)$ , is specified as the merchant's bargaining power relative to the platform's; hence, higher parameter estimates correspond to higher relative bargaining power for merchants and 1 indicate an even split. We find strong evidence that merchants vary in their price bargaining power when dealing with platforms. For an independent merchant with an average number of employees and average annual sales, their bargaining power is lower than the platform's, (0.802,  $p < 0.01$ ), suggesting that the platform can dominate the price setting decision. However, an average-sized chain merchant actually would have higher price bargaining power for their deals (0.802+0.463), so that they can bargain for a price more favorable to their bottom line. Our estimates also indicate that larger merchants—in terms of the number of employees—tend to have higher price-bargaining power than smaller ones (0.084,  $p < 0.01$ ). The coefficient for merchant annual sales is estimated to be positive yet insignificant (0.024,  $p > 0.10$ ): after controlling for all other merchant characteristics, merchant's bargaining power is not related to their level of sales, which we originally expect to serve as another measure of merchant size. This result is perhaps caused by the fact that the level of annual sales reflects categories characteristics more than merchants' characteristics. Hence, caution is needed to interpret this insignificant effect.

Given the characteristics, merchants are found to have lower price bargaining power on Groupon than on LivingSocial (-0.077); the gap is yet significant at 0.10 level but still shows trend for being marginally significant ( $p = 0.11$ ). We plot the merchant-to-platform bargaining power ratio in Figure 1.3. It is clear to see that, first of all, most merchants do have some



bargaining influence in setting deal prices, although both Groupon and LivingSocial seem to dominate many merchants, which our estimates suggest tend to be small and independent ones. At the same time, some merchants are able to exert higher influence on pricing than platforms. Next, we see that the unconditional mean for merchant’s bargaining power is higher on LivingSocial (0.78) than on Groupon (0.73), consistent with the coefficient estimated in Table 1.9. With merchants roughly having similar bargaining leverage between platforms, it is the commission rate (i.e., the transaction cost for merchants’ to use platforms) that plays a critical role in merchant’s platform choices. The bargaining power, however, still matters in terms of determining the payoff level for merchants through affecting the total size of the pie, as we will show in the counterfactual analysis.

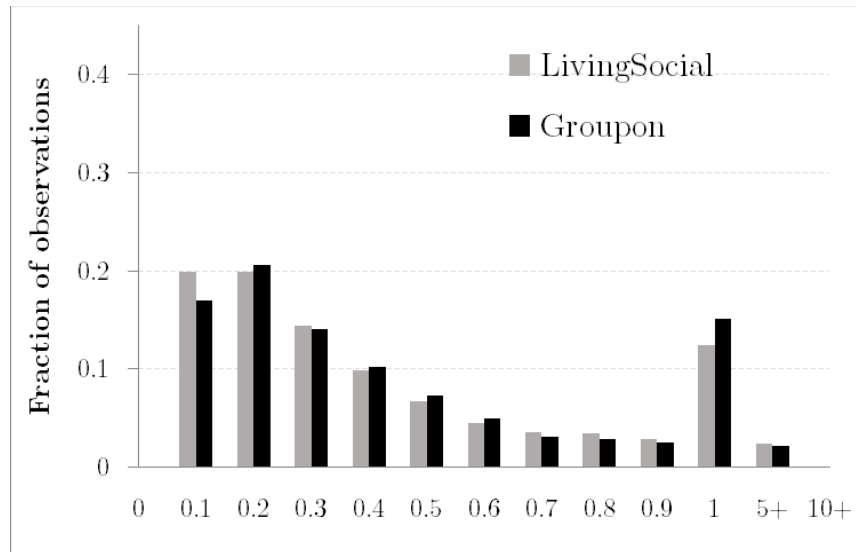


Figure 1.3: Histogram of merchant-to-platform bargaining power ratio

	N	Mean	SD	Min	Max
Groupon	10,105	0.73	2.07	0.09	48.6
LivingSocial	5,336	0.78	2.52	0.001	47.5

Note: The table summarizes the estimated merchant-to-platform bargaining power ratio for Groupon and LivingSocial. The difference in the mean ratio between Groupon and LivingSocial is not significant,  $t(15,439)=1.31$ ,  $p=0.187$ .

Our model also estimates the mean marginal cost, as a percentage of the deal's face value. The percentages vary from 40.3% (for goods category) to 52.5% (for beauty deals). According to our estimates, categories with relatively higher marginal costs are family activities (51.9%), restaurants (51.8%), personal care (51.6%), fitness services (50.6%), and home and automobile services (50.6%). The estimated costs are then used in the following counterfactual analyses.

Table 1.9: Supply parameter estimates

Parameter	Est	SE
$\kappa$		
Difference for Groupon	-0.098***	0.018
LivingSocial	0.514***	0.010
$b_j / b_k$		
Intercept	0.802***	0.039
Merchant size	0.084**	0.037
Chain	0.463***	0.088
Merchant annual sales	0.024	0.059
Groupon dummy	-0.077	0.047
$\psi_c$		
Beauty	0.525***	0.005
Family	0.519***	0.013
Fitness	0.506***	0.008
Goods	0.403***	0.018
Home and auto	0.506***	0.008
Life skill	0.431***	0.013
Live events	0.419***	0.014
Outdoor	0.440***	0.013
Personal care	0.516***	0.010
Restaurants	0.518***	0.006
Sports	0.451***	0.017
Travel	0.450***	0.032

\*\*\*  $p < 0.01$ ; \*\*  $p < 0.05$ ; \*  $p < 0.10$ .

Note: Estimates are from a subset of observations (N=17,470) for which we have an exact match of merchant names with the OneSource dataset. As calculating the disagreement payoffs is very computational intensive, using a sub-sample helps ease the computation burden. We use bootstrap method to compute the standard errors for platform commission rates and merchants' marginal costs.

### **6.3 Counterfactuals**

In this section, we conduct counterfactual analyses to disentangle the effect of platform size, commission rate, price bargaining power, and customer retention rate. Results of these analyses help us quantify the magnitude of each effect and shed light on how those factors jointly determine the pricing outcome and shape the platform choice in the daily deal market.

First, we set out to understand to what extent merchants' bargaining power matters on setting prices and determining profits. To conduct this analysis, we inflate merchants' price bargaining power by 10%, 20%, and 30%, respectively, and then compute the counterfactual price and demand given the new bargaining power. Results are presented in Table 1.10. When merchants are better at bargaining, the new prices would be closer to their profit-maximizing prices, which could be higher or lower than the current observed equilibrium prices. With higher bargaining power, we see that 89% of the merchants would be able to increase the final deal price and the remaining 11% push the price towards a lower point. With a 10% increase, merchants would either end up with a higher price (1.6%) but lower demand (-12.3%), or they would have a lower price (-1.8%) but higher sales (5.0%); the final profits would increase by 4.2% on average. When merchants' bargaining power increases by 20%, the merchants would increase the price by 2.8% or reduce their price by 4.2%, and the average profits increase by 8.9%.

Similarly, with a 30% increase in bargaining power, the price would increase by 3.8% or decrease by 5.3%, and the total profits increase by 11.6%. The general trend from this counterfactual analysis reveals how much merchants can improve profits simply through being better at bargaining. Therefore, the results provide us an upper bound on how much merchants are willing to invest to improve their bargaining ability. Note that in this analysis we take the merchant-platform pairs as given and do not explicitly account for the possibility that an existing pair may disagree on the new price. Nevertheless, the results still inform us on the importance of price-bargaining power in this empirical setting.

Table 1.10: Counterfactual results on percentage changes in price and sales

Increase in merchant's bargaining power		Price Increase (89%)		Price Decrease (11%)		Profits
		Price	Demand	Price	Demand	
10%	Mean	1.6%	-12.3%	-1.8%	5.0%	4.2%
	SD	0.8%	7.4%	2.7%	5.5%	4.4%
20%	Mean	2.8%	-21.7%	-4.2%	11.3%	8.9%
	SD	2.2%	13.1%	9.2%	19.0%	11.5%
30%	Mean	3.8%	-28.9%	-5.3%	14.7%	11.6%
	SD	3.1%	16.5%	10.4%	21.5%	13.5%

Note: The table reports the percentage changes in counterfactual price and demand with a certain percentage increase in merchant's bargaining power.

In the second counterfactual analysis, we examine to what extent the rate of returning customers influences the pricing decision for deals (see Table 1.11 for results). The intuition here is as follows. Merchants care about how much future profit can be generated through the deal promotion. If a higher percentage of current customers (introduced through deals) would return in the future, the merchant would put a relatively higher weight on the demand, i.e., reaching

more people through selling more deals; hence, they would want to accept a lower deal price even if it may mean lower current profit. Indeed, we find that when the rate of returning customers,  $\lambda$ , increases from 0.25 to 0.30, the average deal price would drop by 3.6% and the average profits would increase by 45.9%. In the same vein, if  $\lambda$  decreases, say to 0.20, the merchant would focus more on the current payoff and hence would be incentivized to increase the price by 3.1%, leading to an overall loss in profits by 32.2%. This calculation highlights the importance of retaining the deal customers as regular customers. One managerial implication here is that, by having the expected profit change for a given  $\lambda$ , the merchant knows how much he is willing to spend in order to obtain a higher customer retention rate, for example, through improving the decoration of the physical space (e.g., for restaurants or salons).

Table 1.11: Counterfactual results for different customer retention rate

$\lambda$	Price		Demand		Profits	
	Mean	SD	Mean	SD	Mean	SD
0.20	3.1%	1.2%	-23.6%	12.9%	-32.2%	12.9%
0.30	-3.6%	1.6%	28.5%	17.4%	45.9%	23.9%
0.35	-7.0%	3.1%	52.2%	27.5%	95.7%	44.8%

Note: The counterfactual results are percentage differences with respect to the baseline rate of returning customers (i.e., 0.25). A lower rate typically leads to a higher price and lower overall profits. Vice versa, a higher percentage of returning customers would incentivize the merchant to take a lower current price and still receive higher overall profits.

Our final counterfactual analysis directly speaks to the question of what types of merchants benefit from working with which platform (Table 1.12). To do this, we compute the counterfactual price, demand, and total discounted profits for two hypothetical merchants. The

first is an independent merchant with a small employee size (one standard deviation below the mean) and a voucher with a \$50 face value. The second one is a chain merchant with a large employee size (one standard deviation above the mean) and a voucher of a \$200 face value. Across both cases, the merchants pay a higher transaction cost and have lower bargaining power on Groupon than on LivingSocial. Our results suggest that, for each merchant, the deal would have a lower price and a higher demand on Groupon than on LivingSocial. Although a lower price on Groupon would mean an even lower per-deal payment received for selling deals, the higher number of sales there is able to make up for the difference, yielding a higher total profit (including both the current deal profit and the expected total future profit) for using Groupon. For the small and independent merchant, the total profit would be \$11,600 on Groupon and \$7,400 on LivingSocial; and for the large and chain merchant, the total profits would be \$189,500 versus \$118,100, respectively. This perhaps explains why the number of merchants on Groupon almost doubles that on LivingSocial, despite the fact that merchants have higher price bargaining power and lower transaction cost on LivingSocial. When we compare the relatively payoff between using the two platforms, we see that the large and chain merchant is able to benefit more on Groupon (60% more profits than on LivingSocial) than the small and independent merchant (57% more profits). Therefore, our results seem to suggest that the size of the platform dominates bargaining power or transaction cost and that large and chain merchants are more likely to use the larger platform, consistent with what we see in our model-free evidence.

Table 1.12: Counterfactual results for platform comparison

	Small and Independent Merchants		Large and Chain Merchants	
	Groupon	LivingSocial	Groupon	LivingSocial
$b_j / b_k$	0.60	0.68	1.27	1.35
Price (\$)	34.6	35.3	79.6	87.8
Demand	451.2	248.2	2,101.3	1,155.7
Payment per deal (\$)	14.4	18.2	33.2	45.2
Merchant Profits (000\$)	11.7	7.4	189.5	118.1

Note: We set the face value to be \$50 for the voucher from a small and independent merchant and \$250 for the voucher from a large and chain merchant, which roughly match the sample means respectively.

## 7 Conclusion

We study the pricing decision and platform choice a two-sided market. Using a unique and comprehensive dataset from U.S. daily deals, we specify a structural model that examines consumer behaviors and the strategic interactions between deal platforms and merchants. Our study contributes to the literature by incorporating merchant heterogeneity and allowing prices to be jointly determined by platforms and merchants, both of which are motivated by the real-world complexity but are challenging to model theoretically. When examining the pricing process, our supply-side model allows the platforms to internalize the external value of each deal and allows the merchants to incorporate the promotional effect of the deal. Our results confirm that merchants vary in price elasticity, their ability to grow the platform’s customer base, and their price-bargaining power. In addition, our estimates reveal how platform size, commission rate and bargaining power jointly determine how platforms are differentiated among different types of

merchants. We find that LivingSocial charges a lower commission rate and has lower bargaining power than Groupon, which perhaps explains why it can secure a decent market share despite its smaller platform size. However, merchants still benefit from working with Groupon if this larger platform can generate more sales to compensate its higher commission rate and bargaining power, which is more so for large and chain merchants than for small and independent ones.

To the best of our knowledge, this is one of the first empirical papers in marketing that studies the platform choice and price bargaining in a two-sided market. Due to data limitations, we leave a few interesting and important topics for future research. First, we assume in our paper that platforms are myopic and bargain with merchants to maximize the joint payoffs from the current transaction, regardless of how the outcome may influence its future returns. In fact, a platform may face a tradeoff between current and future payoffs. If it accepts a price more favorable for merchants, more merchants may be willing to join that platform rather than its competitors. Due to the network effect, this could increase the platform's customer base and boost profits in the long run. Therefore, if a platform behaves dynamically, it should negotiate a price that maximizes the product of the merchant's current profits and its own future discounted total profits, rather than merely its own current profits. Modeling such dynamic decision-making, however, requires a longer time horizon of observations on the platform's pricing decisions and its growth than we could manage in this study. Furthermore, by focusing on the current period's payoffs, we generate insights on how platforms and merchants internalize price-bargaining power in their strategic interactions; this lays a foundation for future research to study the forward-looking behavior of platforms.

Second, while the future payoff for merchants is important in this setting, we do not have merchant sales and revenue data outside the deal domain to study the details of customer



retention behavior. Instead, we have to assume a fixed retention rate for all merchants. Based on our sample of interviews, the rate we use represents what a merchant typically expects. In reality merchants may vary in terms of their ability to retain customers; however, we are not able to include this merchant heterogeneity due to data limitations. Nevertheless, our model and counterfactual results light on how retention rate influences the pricing decision and the merchant's bottom line. With better data, our model can be applied to generate more precise estimates, for example, on marginal cost for each merchant.

Third, we do not try to pin down the exact determinants for price-bargaining power. To an extent, bargaining ability may depend on the negotiation skills and the incentives behind individual negotiators. Unfortunately, the data in our study are not available to address the mechanisms of bargaining power. Yet, better understanding the determinants could lead to interesting insights concerning how managers can shape market outcomes by influencing their bargaining power.

Lastly, as is true for almost all empirical work, our results may be contingent on the specific characteristics of this study setting. In particular, we try to back out the pricing decision given the observed pairs of merchants and platforms. We do not formally model the formation of the networks among merchants and platforms, which is known to be challenging both theoretically and empirically. In order to do so, we not only need to know the cost of using each platform, but also need to form expectations on what other competitor merchants would do given the choice for the focal merchant. Although modeling the formation of the merchant-platform network is beyond the current paper, our research helps understand how prices are set and profits are split, which could be helpful for future studies on this topic.

Viral Videos:  
The Dynamics of Online Video Advertising Campaigns

LINGLING ZHANG

CLARENCE LEE

ANITA ELBERSE

Abstract

Firms increasingly promote their products using advertisements posted on online video-sharing sites such as YouTube. There, web users often redistribute these “advertiser-seeded” advertisements, either in their original form or as altered, derivative works. Such user-generated “viral” placements can significantly enhance the true number of an advertising campaign’s impressions – in fact, across our data for movie and video-game trailers, the number of views generated by viral placements is three times as many as the number of views for the original “advertiser-seeded” placements. In this study, in order to help advertisers understand how their online video advertisements spread, we investigate the dynamics of viral video campaigns, modeling the interactions between marketers’ advertising actions and consumers’ decision to view and spread advertisements. We find that several instruments under the control of advertisers – notably the intensity and timing of original video placements –influence the extent to which campaigns benefit from user-generated content. Our results underscore that, with the right

strategy, advertisers can substantially and inexpensively increase the number of impressions that their online video campaigns yield.

## 1 Introduction

Firms increasingly promote their brands via advertisements posted on online video-sharing sites such as YouTube, Dailymotion, and Metacafe. There, advertisers encounter a phenomenon that they rarely experienced with traditional media: videos they upload are commonly redistributed by viewers either in their original form (as simple copies), or as altered “derivative” works (commonly referred to as spoofs, remixes, and mash-ups). Such user-generated “viral” placements can significantly enhance the true spread of advertising campaigns – in fact, in many instances, the number of views generated by viral placements tends to be several times greater than the number of views for the original “advertiser-seeded” placements. Consider the Sony Pictures movie *Angels and Demons*, starring Tom Hanks: of the roughly 6 million views that its trailers amassed on YouTube alone, more than 70% came from trailer videos placed by users, not the studio itself (see Figure 2.1). Ignoring these user-generated views thus could lead to a substantial underestimation of a campaign’s number of impressions, and a lack of understanding of “how videos go viral” may mean advertisers miss opportunities to significantly expand their reach at a low cost.

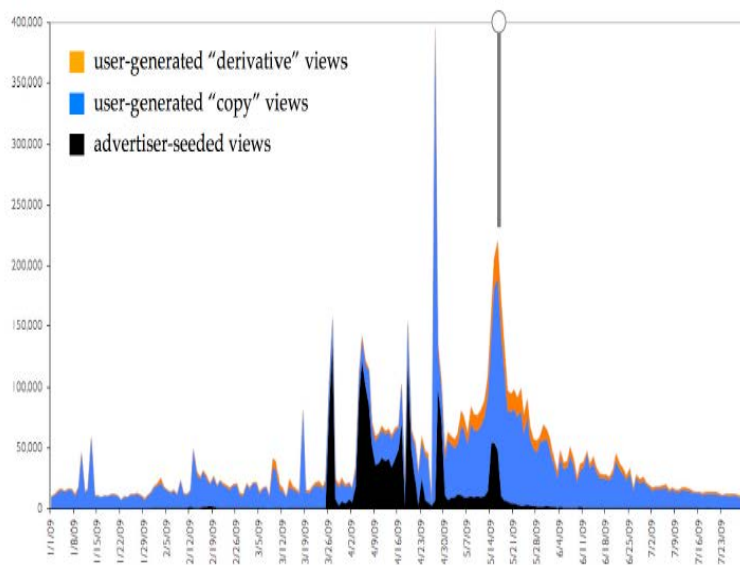


Figure 2.1: Views for the movie *Angels and Demons*. We plot the daily views the movie *Angels and Demons* received from its video advertising campaigns. A large proportion of views are from user-generated videos, either as an exact copy or with modifications (i.e., derivative videos), suggesting that user-generated videos are a force to be reckoned with when it comes to generating impressions. The vertical line corresponds to the date when the movie was released in theaters.

In this paper, we study the dynamics of viral video advertising, in particular, the interplay between user-generated content and marketer activities. Whereas the topic of user-generated content in social media has received rising attention recently, much of the extant research in this domain has focused on assessing the impact of user-generated content on sales (e.g., Berger et al. 2010; Chevalier and Mayzlin 2006; Chintagunta et al. 2010; Liu 2006). In contrast, our study seeks to answer what determines the extent to which users *view* online video advertisements and *generate* their own copies of those videos. We focus on two categories of entertainment goods, movies and video games, for several reasons. First, entertainment firms are among the biggest advertisers in the U.S.; and they are disproportionately active in online environments – especially on video-sharing sites. Second, films and video games tend to have short life cycles (they usually generate the lion’s share of their revenues in a matter of weeks or months), and the goal for their advertising campaigns is usually to generate impressions among a large audience (rather than, say, build a long-term brand image). Therefore, films and video games provide a good research setting to study the impact of marketer actions on generating online impressions. In addition, film studios and video-game publishers collectively release many products in a relatively short

time period, providing us with the cross-sectional observations needed to account for inherent differences among brands.

We compiled a comprehensive data set from multiple sources. Our sample includes 546 trailers (hereafter referred to as “videos”) from 141 titles (hereafter, “brands”) for which advertising campaigns were conducted on video-sharing sites from September 2005 to April 2011. For each brand, we track its videos available online—including the ones uploaded by advertisers and those generated by users—and collect daily placements and viewership counts for each video, which together yield a total of more than 63,300 placements and roughly 1.1 billion views. Our data also include daily spending for each brand on television advertising. Across our sample, we find that user-generated videos play a vital role in enhancing the overall reach of online video campaigns: 95% of the videos uploaded onto the video-sharing websites were generated by users, accounting for 78% of total advertising views (see Table 2.1). These numbers are not driven by a few brands: more than 75% of the videos in our sample had most of their impressions generated by user-uploaded copies, and for roughly half of the titles, the share of user-generated views was higher than 80%. Although per-video view counts tended to be higher for advertiser-uploaded videos, the large volume of user-generated content makes it a critical factor in determining the success of advertising campaigns.

To gain insights on the drivers of viral video impressions, we examine the inter-variable dynamics that capture the relations among advertiser-seeded and user-generated videos over time. We model five outcome variables—offline advertising, advertiser-seeded placements and views, and user-generated placements and views—in a system of dynamic linear regression models (DLM) and jointly estimate the parameters using the hierarchical Bayesian framework.

Our results provide insights into how managers can influence the dynamics of video advertising. For example, we find that advertisers' seeding intensity matters: when advertisers upload more of their own videos, users tend to follow the action and create more copies. Therefore, by manipulating their seeding intensity, advertisers can drive the total number of campaign impressions both directly and indirectly via the channel of user-uploaded videos. We also find *when* to seed advertisements matters: uploading a video closer to the release date can boost views within a shorter period of time. We further find *offline* advertising to exert a positive, albeit relatively small, spillover effect on online views, which suggests that, especially compared to (free) advertiser-seeded placements, it is more cost-effective for advertisers to "earn" impressions by manipulating seeding intensity and timing than to "buy" impressions through advertising in offline channels.

Furthermore, our results reveal how a user's decision to upload and view a video depends on how prior users interacted with it. When it comes to viewing, users exhibit a "jump-on-the-bandwagon" tendency: videos that have received a larger number of user ratings, more favorable user ratings, or just more views attract more future viewing. But when it comes to uploading copies, while the rating favorability still has a positive effect, the volume of user ratings exerts a *negative* effect. We think this may be explained by the underlying motivation to share (Wu and Huberman 2008): a user may be more motivated to share if the expected impact of adding one more video out-weighs the efforts involved. For videos that are perceived favorably by others (i.e., have a high average rating) but have yet to reach a large audience (i.e., have a low volume

of ratings), an additional copy is more likely to attract attention, and hence users are more inclined to upload copies of such videos<sup>14</sup>.

Our paper makes two main contributions. First, we extend the literature on online advertising by assessing the process of how advertising messages “go viral” in social media channels, specifically, on video-sharing sites. By disentangling the effect of advertisers’ and users’ activities, we show to what extent and in what way online users help advertisers spread their advertisements. These dynamics have hardly been documented in the advertising literature. The “one-to-many” web-based viral video campaigns likely diffuse quite differently from the “one-to-one” email and offline word-of-mouth campaigns that have been subjects of recent studies (e.g., Godes and Mayzlin 2009; Toubia et al. 2011). We are unaware of any studies on the impressions generated by viral advertisements placed on video-sharing sites, despite this media channel’s rapidly growing importance. Second, not content with merely describing the phenomenon, we show how advertisers can *influence* the extent to which users help propel their campaigns to higher numbers of impressions. This answers the call for more research on how firms should manage and organize social media efforts to optimize online communication with consumers (Aral, Dellarocas, & Godes 2013; Miller & Tucker 2013). We show that inexpensive instruments under the direct control of advertisers—the number and timing of advertising messages seeded—can substantially drive the total number of advertisement impressions. By quantifying the impact, we help advertisers understand how to effectively use the tools at their disposal to maximize advertising exposures for their brands.

---

<sup>14</sup> Among all user-uploaded videos in our sample, 74% were exact copies of the original content and the majority of the remaining videos were only with minor modifications.



## 2 Data

Our data cover four sets of variables: those related to video placements and views, brand-level television advertising expenditures, brand and video characteristics, and online user ratings of the videos. In this section, we describe each set and present summary statistics.

### 2.1 Video Placements and Views

We obtain video placements and views data from Visible Measures, a leading provider of measurements for online video campaigns. The company is behind the Top 10 Viral Video Ads chart published weekly in Advertising Age, and the Top 10 Online Film Trailers chart published weekly in the film trade magazine Variety. We compile the data using the Visible Measures Viral Reach Database, a repository of data on more than 100 million Internet videos across over 150 video-sharing destinations (such as YouTube and Metacafe). Movie and video-game trailers consistently rank among the most popular videos: in 2010 they made up over 15% of the so-called “100 Million Views Club” –videos that collectively accumulated more than 100 million views.

Our sample includes 546 trailers that were used to promote 141 movie and video-game titles. The title selection is representative of the spectrum of high- and low-budget movie and video-game brands released by major studios in North America. For each of the videos, we have daily number of advertiser-seeded placements (*seeded\_placements*) and user-generated placements (*ug\_placements*) as well as the corresponding daily views (*seeded\_views*, *ug\_views*). Both placements and views are specified as incremental measures, reflecting the number of new placements and views generated each day. YouTube is by far the most popular site, accounting for 77% of all placements and 85% of all views.

Table 2.1: Summary patterns of placements and views

Metrics	Movies	%	Video games	%	Total	%
Seeded Placements	1,988	5%	1,126	5%	3,114	5%
Copy Placements	31,915	75%	15,189	72%	47,104	74%
Derivative Placements	8,433	20%	4,652	22%	13,085	21%
Total Placements	42,336	100%	20,967	100%	63,303	100%
Seeded Views	156,687,170	19%	87,461,849	30%	244,149,019	22%
Copy Views	514,131,419	64%	108,687,506	38%	622,818,925	57%
Derivative Views	136,030,774	17%	91,160,109	32%	227,190,883	21%
Total Views	806,849,363	100%	287,309,464	100%	1,094,158,827	100%

Note: We present the total numbers of placements and views for both advertiser-seeded and user-generated videos. Although they make up only 5% of total placements, advertiser-seeded videos account for more than 20% of total views.

The terms “placements” and “views” deserve a careful definition. First, **placements** are the individual instances of videos associated with a given brand’s campaign, each with a unique URL. The more placements a campaign has, the more copies of campaign-related videos exist across the universe of video-sharing sites (simply linking to a video does not constitute a placement). For example, if Sony Pictures has uploaded a trailer for its film *Angels and Demons* onto YouTube and onto Metacafe, two placements of that video exist. Similarly, if Sony Pictures has uploaded the trailer onto YouTube under two different URLs, we also speak of two placements.<sup>15</sup>

Placements can be divided into advertiser-seeded placements and user-generated (or “viral”). **Advertiser-seeded placements** are videos uploaded by the content creator, be it a studio or an agency (hereafter referred to as “advertisers”). **User-generated placements** are all the videos that are connected with advertiser-seeded content but come from anyone other than the content

<sup>15</sup> This happens fairly regularly, for instance because it enables studios to track the effectiveness of marketing materials that make reference to URLs.

creator. There are two forms: copies and derivative works. Copy placements are duplicates of the original content unchanged in content. Derivative placements are versions that somehow deviate from the original content—such as mash-ups, parodies, spoofs, and remixes. For example, if a user were to copy the trailer uploaded by Sony Pictures and upload it again onto YouTube (but under his own name) that would constitute a user-generated copy placement. If that same user were to take the trailer but add a screen at the beginning of the video that reads “Brought to you by John Doe” and upload it onto YouTube, that would constitute a derivative placement. Very few derivative works are changed to such an extent that the original message is lost, most derivative placements being relatively straightforward alterations of original content that preserve the intended meaning.

Finally, **views** are the number of times a campaign’s placements are watched. Number of views is synonymous with number of impressions. Views are categorized according to placement type: those of videos originated by content creators are **advertiser-seeded views**, those of videos originated by users, **user-generated views**.

For each video, our data collection begins with the first placement and ends the day before the release of the brand. Table 2.2 presents summary statistics for the variables. On average, a video has 3.56 advertiser-seeded placements (which generate around 3,000 daily views) and 94.6 user-generated placements (which altogether attract more than 10,300 daily views). For the analysis, we focus our attention on up to five months before the release, for managerial and empirical considerations. For movies and video games, the lion’s share of marketing activities happens relatively close to the product release date, and managers are also most concerned with enhancing brand awareness during this period. The concentration of marketing efforts is

reflected in our data. Truncating the data therefore helps us avoid imprecise estimates for early periods.

Table 2.2: Summary statistics

Variable	Mean	Median	SD	Min	Max
By video, by day (N=82,462)					
seeded placements (cum)	3.56	2	4.99	0	68
user-generated placements (cum)	94.64	46	135.47	0	1,297
seeded views	2,961	6	29,092	0	3,990,643
user-generated views	10,308	1054	47,922	0	3,162,621
seeded ratings valence	331.5	461.0	219.1	0	500
seeded ratings volume	574	34	2,215	0	31,213
user-generated ratings valence	423.0	464.7	125.3	0	500
user-generated ratings volume	3,099	982	6,308	0	73,084
television advertising (\$)	74,039	0	310,331	0	12,952,472
By video (N=546)					
sentiment	0.98	1	0.14	-1	1
seeding days prior to release	112	85	105	8	1,057
By brand (N=141)					
production budget (\$million)	55.7	35	46.0	1.8	200
critics review score	70.1	73	19.5	18.0	98
	Frequency	%			
franchise	55	43			
rating: restricted/mature	56	43			
genre					
action	56	43			
adventure	14	11			
comedy (reference)	38	30			
drama	20	16			
romance	13	11			
thriller	11	9			
animation	7	5			
first-person shooter	7	5			
other	27	21			

Note: Variables are organized according to the measurement unit. We present daily metrics first, followed by video characteristics and brand characteristics.

## 2.2 Advertising Expenditures

We collect daily television advertising expenditures from TNS Media Intelligence, which tracks network, cable, spot, and syndication television advertising. Although most brands also advertise in other media (such as magazines, newspapers, radio, outdoor, and on the web), television advertising tends to represent the bulk of the spending for products of our context. Within television, cable television advertising (48%) is the most popular, closely followed by broadcast network television advertising (34%).

## 2.3 Brand and Video Characteristics

We turn to various sources for characteristics of the movie and video-game brands<sup>16</sup>. One challenge we face is the need to account for differences among brands due to their intrinsic quality. Product quality, a source of heterogeneity across brands, likely affects the level of advertising, video placements, and views. Luckily, for entertainment products, rich information sources exist that provide good proxies for product quality (e.g., Elberse and Eliashberg 2003). In this study, we include two such metrics: production budgets (*budget*) and critics' reviews (*critic\_review*), to at least partially account for the heterogeneity due to product quality.

Other brand-level characteristics include indicators for genres, an indicator for product rating (*rating\_R*, 1 being “restricted” for movies and “mature” for games), an indicator for property being part of an existing franchise, and an indicator for video games (as opposed to movies).

To account for heterogeneity in video content, we measure for each video the positive and negative emotion evoked among viewers. Visible Measures collects snapshots of the top 100 words that appear most frequently in the user comments for each video, which yields more than

---

<sup>16</sup> Brand characteristics were collected primarily from two online databases: the Internet Movie Database ([www.imdb.com](http://www.imdb.com)) and Box Office Mojo ([www.boxofficemojo.com](http://www.boxofficemojo.com)).

2,000 unique words across the videos in our sample. We asked two independent coders to categorize each word as having either a positive (e.g., “amazing”), negative (e.g., “boring”), or neutral meaning; a handful of words on which the coders disagreed were eliminated from the analysis. We gauge emotional response to a video by a sentiment measure (*sentiment*), calculated as  $(positive\_words^2 - negative\_words^2) / (positive\_words^2 + negative\_words^2)$ . This variable ranges between -1 and 1, higher values indicating greater positive appeals. The high average sentiment measure (mean = .98, SD = 0.14) confirms that online user reviews tend to be positive in general, which is consistent with findings from other studies (e.g., Anderson 1998; Chevalier and Mayzlin 2006; Liu 2006).

## 2.4 User Ratings

The Visible Measures’ Viral Reach Database tracks daily user ratings for both advertiser-seeded and user-generated videos. Ratings are measured as numeric scores ranging from 0 (the least favorable) to 500 (the most favorable). From these data we construct the volume (*rating\_volume*) and valence (i.e., the mean, *rating\_valence*) of the scores<sup>17</sup>. Mean ratings for the brands in our sample are predominately favorable at 331.5 (SD = 219.1) for advertiser-seeded, and 423.0 (SD = 125.3) for user-generated, videos.

## 3 Model

To answer our research questions, we examine the dynamics of five outcome variables—offline advertising, advertiser-seeded placements and views, and user-generated placements and views. These variables reflect decisions from two types of agents: advertisers and consumers. In the context of online video campaigns, advertisers need to decide how much to spend on offline

---

<sup>17</sup> Including both rating volume and valence does not introduce a multi-collinearity problem. The variance inflation factor is 2.08 for rating volume and valence on advertiser-seeded videos and 1.19 for user-generated videos.

advertising and how many videos to upload onto video-sharing websites, and consumers need to decide whether to watch a video and whether to upload their own copies. To understand the drivers for these decisions, we explicitly model not only the temporal dynamics that summarize the base growth pattern for each variable but also the inter-variable dynamics that capture the time-varying association among them. For each variable, we specify a dynamic linear model (DLM) that best captures the underlying data-generation process; and we jointly estimate the system of five simultaneous equations. Next, we describe the specification for the outcome variables, one at a time.

### 3.1 Advertisers' Decision 1: Spending on Television Advertising

In the meantime of promoting products using online video campaigns, marketers often also advertise on TV. We model the amount of advertising spending for brand  $i$  video  $j$  at day  $t$  as a function of two main components (equation 1): one related to brand heterogeneity and the other related to time variation. First, the amount of ad spending is a function of a brand-specific intercept,  $\alpha_i^1$ , and an idiosyncratic shock,  $\nu_i^1$ , where  $\alpha_i^1$  is further related to brand-level observed characteristics,  $X_i$ <sup>18</sup>.

$$\text{advertising}_{ijt} = \alpha_i^1 + \delta_t + \tau^1 X_i + \varepsilon_{ijt}^1 \quad (1)$$

$$\alpha_i^1 = \kappa^1 X_i + \nu_i^1 \quad (1.1)$$

Second, ad spending is also a function of day-specific characteristics,  $X_t$ , and a time-varying intercept,  $\delta_t$ . Examples of  $X_t$  include days of the week and whether it is a holiday. We specify  $\delta_t$  to be time-varying: it summarizes how offline advertising evolves as it approaches the product release date. In particular, we assume that the amount of ad spending at day  $t$ ,  $\delta_t$ , persists from

---

<sup>18</sup> The superscripts in our equations denote the order of the outcome variable. Notable is that advertising spending is constant for all videos of the same brand  $i$ , and hence is only related to brand characteristics.

the amount at the previous day,  $\delta_{t-1}$ , and is augmented by an exogenous shock,  $\omega_t^\delta$ . Thus, through  $\delta_t$  we capture the temporal dynamics of offline advertising while being agnostic to the underlying mechanism. Equation 1.2 summaries the parameter evolution process.

$$\delta_t = G\delta_{t-1} + \omega_t^\delta \quad (1.2)$$

Here, the stochastic term,  $\omega_t^\delta$ , is assumed to follow a normal distribution,  $N(0, V_\omega^\delta)$ . We set the transition matrix to identity (i.e.,  $G = I$ ) and, by doing so, assume  $\delta_t$  to evolve following a random walk process. This specification strikes a balance between inter-temporal dependence and random variance, and is popularly employed by DLM modelers in marketing applications (e.g., Ataman et al. 2010; Teixeira & Wedel 2010).

It is noteworthy that we deliberately exclude online video placements and views (being it advertiser-seeded or user-generated) from the equation for offline advertising. Industry practice dictates that television ad spots are often purchased months before the advertisements are actually aired, and are not timely adjusted in response to how video advertisements are received on video-sharing sites. Therefore, online video diffusion should not be included as factors to determine the level of offline advertising.

### 3.2 Advertisers' Decision 2: Uploading Advertiser-seeded Videos

In parallel to advertising on TV, advertisers manage their online video campaigns by deciding how many video seeds to upload onto video-sharing sites. The number of video placements for brand  $i$  video  $j$  at day  $t$  is specified as follows.

$$\begin{aligned} \text{seeded\_placements}_{ijt} = & \alpha_i^2 + \beta_j^2 + \tau^2 X_t \\ & + \theta_{0t} + \theta_{1t} \text{ad\_stock}_{ijt} + \theta_{2t} \text{seeded\_views}_{ijt-1} \\ & + \gamma_1^2 \text{rating\_valence}_{ijt-1} + \gamma_2^2 \text{rating\_volume}_{ijt-1} \\ & + \varepsilon_{ijt}^2 \end{aligned} \quad (2)$$



By this specification, advertiser-seeded placements are a function of four components. First, it is related to brand- and video-specific attributes,  $\alpha_i^2$  and  $\beta_j^2$ , respectively, which we further specify as being a function of observed brand-level characteristics,  $X_i$ , and video-level characteristics,  $X_j$ , plus unobserved idiosyncratic shocks,  $v_i^2$  and  $\xi_j^2$ . That is,

$$\alpha_i^2 = \kappa^2 X_i + v_i^2 \quad (2.1)$$

$$\beta_j^2 = \eta^2 X_j + \xi_j^2 \quad (2.2)$$

Second, the number of uploaded videos also varies over time; part of this temporal variation could be explained by the observed day characteristics,  $X_t$ , and part of it is captured by the day-specific intercept,  $\theta_{0t}$ . Similar to  $\delta_t$  in equation 1, this day-specific intercept summarizes the temporal dynamics of seeded placements, not captured by  $X_t$  and independent from the influence of other time-varying variables.

Third, when deciding how many new videos to place online, advertisers often consider how much advertising has been offered offline: coupling online campaigns with offline advertising could generate a positive synergy between marketing channels and hence enhance the overall campaign reach. Because advertising effect is likely to persist over time, we construct an ad stock measure, which is defined as the decayed ad stock augmented by the daily-added advertising (Nerlove & Arrow 1962)<sup>19</sup>:

$$ad\_stock_{ijt} = \phi ad\_stock_{ijt-1} + advertising_{ijt} \quad (2.3)$$

---

<sup>19</sup> To estimate the carryover factor, we conducted a grid search similar to the procedure used in Manchanda et al. (2008). We set the factor to range from 0 to 1 with an increment of 0.1, estimated our main model for each value, and calculated the log-likelihood statistics. The log-likelihood statistic was the highest when the carryover factor equaled 0.7.

Lastly, advertisers also take into account how their existing videos have been received by users. We capture this by three metrics: how many times the videos were watched the day before ( $seeded\_views_{ijt-1}$ ), how many user ratings they received ( $rating\_volume_{ijt-1}$ ), and how much they were liked ( $rating\_valence_{ijt-1}$ ). It is noteworthy that we use only user ratings of *advertiser-seeded* videos to construct the volume and valence measures when modeling the placements and views for *advertiser-seeded* videos, and vice versa for *user-generated* videos.

We specify the effect of ad stock ( $\theta_{1t}$ ) and existing seeded views ( $\theta_{2t}$ ) to be time-varying, so as to allow the effects to vary at different stages of the campaign cycle. Our exploratory analyses (available upon request) indicate that the impact of user ratings on advertiser-seeded placements is largely constant over time; hence, we specify  $\gamma_1$  and  $\gamma_2$  to be time-invariant. The evolution process for the dynamic parameters is defined similar to equation 1.2.

$$\begin{aligned}\theta_t &= [\theta_{0t} \quad \theta_{1t} \quad \theta_{2t}]' \\ \theta_t &= G\theta_{t-1} + \omega_t^\theta\end{aligned}\tag{2.4}$$

Our exclusion of user-generated placements and views from equation 2 is a deliberate choice, stemming partly from conversations with advertisers who indicated that they do not take into account user-generated content when making seeding decisions. Even were advertisers to want to monitor the spread of user-generated versions of their advertisements, they would find it difficult to do so, especially in a comprehensive and timely manner.

### 3.3 Consumers' Decision 1: Viewing Advertiser-seeded Videos

Once advertisers have uploaded videos, consumers decide whether to watch them. The individual decisions aggregate into the number of views received for advertiser-seeded placements, denoted as  $seeded\_views_{ijt}$ . We model this variable as a function of four components. The first is related to brand- and video-specific attributes (i.e.,  $\alpha_i^3$  and  $\beta_j^3$ ), which are further

specified as a function of observed characteristics,  $X_i$  and  $X_j$ , and idiosyncratic shocks,  $v_i^3$  and  $\xi_j^3$ . Second, seeded views are also a function of observed day characteristics,  $X_t$ , and an unobserved day-specific intercept,  $\lambda_{0t}$ .

Third, the extent to which consumers watch advertiser-seeded videos also depends on the advertisers' marketing actions that are used to promote the brand, namely, how many seeded copies advertisers have placed on video-sharing sites (*seeded\_placements\_cum<sub>ijt</sub>*) and how much they have spent on offline advertising (*ad\_stock<sub>ijt</sub>*). We specify the right-hand-side placement variable as the cumulative stock of placements, i.e., the number of copies available for video  $j$  of brand  $i$  on day  $t$  (and not just the number of *new* placements that day), because the cumulative number best reflects the assortment of videos available to viewers.

Lastly, we allow future consumers' viewing decisions to be influenced by existing users' viewing decisions and feedback. Prior work in this area (e.g., Salganik and Watts 2009) has shown that, with limited information available to guide product selection decisions, consumers tend to sample what others have chosen, leading to a "cumulative advantage" or "success-breeds-success" phenomenon. This reinforcement loop could occur in three ways in the context of online video campaigns. For one, a large volume of current views may attract more future users to view the video; we capture this effect by allowing  $\lambda_{0t}$  to persist over time (equation 3.3). The second reason for a potential "success-breeds-success" trend is because consumers who become aware of the brand through user-generated videos may further decide to watch the official advertise-seeded videos. We model this spillover effect by including user-generated views on the previous day, *ug\_views<sub>ijt-1</sub>*, in the equation specification. Third, the previous day ratings from current users, *rating\_valence<sub>ijt-1</sub>* and *rating\_volume<sub>ijt-1</sub>*, also enter the information set for future users when they decide what to view.

The above considerations lead to the following equation for advertiser-seeded views:

$$\begin{aligned}
seeded\_views_{ijt} = & \alpha_i^3 + \beta_j^3 + \tau^3 X_t \\
& + \lambda_{0t} + \lambda_{1t} ad\_stock_{ijt} \\
& + \lambda_{2t} seeded\_placements\_cum_{ijt} + \lambda_{3t} ug\_views_{ijt-1} \\
& + \gamma_1^3 rating\_valence_{ijt-1} + \gamma_2^3 rating\_volume_{ijt-1} \\
& + \varepsilon_{ijt}^3
\end{aligned} \tag{3}$$

$$\alpha_i^3 = \kappa^3 X_i + \nu_i^3 \tag{3.1}$$

$$\beta_j^3 = \eta^3 X_j + \xi_j^3 \tag{3.2}$$

As in equation 2, we allow the coefficients of offline advertising, seeded placements, and user-generated views to be time-varying, to account for the possibility that these effects vary in magnitude as the time approaches the product's release date. We again assume a random-walk process for the dynamic parameters:

$$\begin{aligned}
\lambda_t = & [\lambda_{0t} \quad \lambda_{1t} \quad \lambda_{2t} \quad \lambda_{3t}]' \\
\lambda_t = & G\lambda_{t-1} + \omega_t^\lambda
\end{aligned} \tag{3.3}$$

### 3.4 Consumers' Decision 2: Posting User-generated Videos

Identifying the drivers for user-initiated video uploading is one of our key objectives. Similarly, we specify the underlying data-generation process as containing four components. The first relates to observed brand- and video-specific characteristics ( $X_i$  and  $X_j$ ) and those unobserved ( $\nu_i^4$  and  $\xi_j^4$ ); and the second relates to observed day characteristics ( $X_t$ ) and unobserved time variation ( $\psi_{0t}$ ).

Third, consumers can also be influenced by advertisers' activities, in particular, the intensity of online seeding ( $seeded\_placements\_cum_{ijt}$ ) and offline advertising ( $ad\_stock_{ijt}$ ), which further determines how many new videos would be generated and uploaded by web users.

It is noteworthy that the right-hand-side placements are again the cumulative number of placements available to potential viewers online.

Lastly, user-generated placements could also be influenced by how much existing videos are being viewed ( $ug\_views_{ijt-1}$  and  $seeded\_views_{ijt}$ ) and liked ( $rating\_valence_{ijt-1}$  and  $rating\_volume_{ijt-1}$ ). Motivation-based theory (Wu & Huberman 2008) would suggest that consumers are more motivated to upload new videos they expect to have a greater impact, either because the not-yet-popular content has potential, or in the hope that a new version might help a popular, but not yet sufficiently “liked,” video.

Combining these four sets of factors yields the following specification for user-generated placements:

$$\begin{aligned}
ug\_placements_{ijt} = & \alpha_i^4 + \beta_j^4 + \tau^4 X_t \\
& + \psi_{0t} + \psi_{1t} ad\_stock_{ijt} + \psi_{2t} seeded\_placements\_cum_{ijt} \\
& + \psi_{3t} seeded\_views_{ijt} + \psi_{4t} ug\_views_{ijt-1} \\
& + \gamma_1^4 rating\_valence_{ijt-1} + \gamma_2^4 rating\_volume_{ijt-1} \\
& + \varepsilon_{ijt}^4
\end{aligned} \tag{4}$$

$$\alpha_i^4 = \kappa^4 X_i + \nu_i^4 \tag{4.1}$$

$$\beta_j^4 = \eta^4 X_j + \xi_j^4 \tag{4.2}$$

$$\psi_t = [\psi_{0t} \quad \psi_{1t} \quad \psi_{2t} \quad \psi_{3t} \quad \psi_{4t}]' \tag{4.3}$$

$$\psi_t = G\psi_{t-1} + \omega_t''$$

### 3.5 Consumers' Decision 3: Viewing User-generated Videos

We specify the equation for user-generated views similar to the model for advertiser-seeded views in equation 3. In addition to characteristics related to brand, video, and day, views are also a function of users' goodwill attributed to advertising ( $ad\_stock_{ijt}$ ), the spread of user-generated placements ( $ug\_placements\_cum_{ijt}$ ), the attention raised by advertisers-uploaded seeds

(*seeded\_views<sub>ijt</sub>*), as well as the liking of existing videos (*rating\_valence<sub>ijt-1</sub>* and *rating\_volume<sub>ijt-1</sub>*). The specification is summarized as follows.

$$\begin{aligned}
 ug\_views_{ijt} = & \alpha_i^5 + \beta_j^5 + \tau^5 X_t \\
 & + \rho_{0t} + \rho_{1t} ad\_stock_{ijt} \\
 & + \rho_{2t} seeded\_views_{ijt} + \rho_{3t} ug\_placements\_cum_{ijt} \\
 & + \gamma_1^5 rating\_valence_{ijt-1} + \gamma_2^5 rating\_volume_{ijt-1} \\
 & + \varepsilon_{ijt}^5
 \end{aligned} \tag{5}$$

$$\alpha_i^5 = \kappa^5 X_i + \nu_i^5 \tag{5.1}$$

$$\beta_j^5 = \eta^5 X_j + \xi_j^5 \tag{5.2}$$

$$\begin{aligned}
 \rho_t = & [\rho_{0t} \quad \rho_{1t} \quad \rho_{2t} \quad \rho_{3t}]' \\
 \rho_t = & G\rho_{t-1} + \omega_t^p
 \end{aligned} \tag{5.3}$$

By now, we have finished discussing our model specification. Due to the complexity of the model, we present a summary of the parameter descriptions in Table 2.3.

Table 2.3: Summary of model parameters

Equation No.	(1)	(2)	(3)	(4)	(5)
Dependent Variable	offline advertising	seeded placements	seeded views	ug placements	ug views
<b>Dynamic Parameters</b>					
Time-specific intercept	$\delta_t$	$\theta_{0t}$	$\lambda_{0t}$	$\psi_{0t}$	$\rho_{0t}$
Effect related to:					
offline advertising		$\theta_{1t}$	$\lambda_{1t}$	$\psi_{1t}$	$\rho_{1t}$
seeded placements			$\lambda_{2t}$	$\psi_{2t}$	
seeded views		$\theta_{2t}$		$\psi_{3t}$	$\rho_{2t}$
ug placements					$\rho_{3t}$
ug views			$\lambda_{3t}$	$\psi_{4t}$	
<b>Time-invariant Parameters</b>					
Effect related to:					
brand-specific intercept	$\alpha_i^1$	$\alpha_i^2$	$\alpha_i^3$	$\alpha_i^4$	$\alpha_i^5$
video-specific intercept		$\beta_j^2$	$\beta_j^3$	$\beta_j^4$	$\beta_j^5$
Effect related to:					
day characteristics	$\tau^1$	$\tau^2$	$\tau^3$	$\tau^4$	$\tau^5$

brand characteristics	$\kappa^1$	$\kappa^2$	$\kappa^3$	$\kappa^4$	$\kappa^5$
video characteristics		$\eta^2$	$\eta^3$	$\eta^4$	$\eta^5$
ratings valence		$\gamma_1^2$	$\gamma_1^3$	$\gamma_1^4$	$\gamma_1^5$
ratings volume		$\gamma_2^2$	$\gamma_2^3$	$\gamma_2^4$	$\gamma_2^5$

Note: Our model consists of a system of five equations, each containing dynamic and time-invariant parameters. We provide a summary of these parameters. The dynamic parameters are denoted by the subscript “ $t$ ”. The superscripts for the time-invariant parameters correspond to the equation number.

## 4 Estimation and Identification

### 4.1 Estimation

Our model consists of a system of five equations with the error terms assumed to follow a multivariate normal distribution,  $MVN(\mathbf{0}, \Omega)$ :

$$\boldsymbol{\varepsilon}_{ijt} = \begin{pmatrix} \varepsilon_{ijt}^1 \\ \varepsilon_{ijt}^2 \\ \varepsilon_{ijt}^3 \\ \varepsilon_{ijt}^4 \\ \varepsilon_{ijt}^5 \end{pmatrix} \sim MVN \left( \begin{pmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{pmatrix}, \Omega_{5 \times 5} = \begin{pmatrix} \Omega_{11} & \Omega_{21} & \Omega_{31} & \Omega_{41} & \Omega_{51} \\ \Omega_{21} & \Omega_{22} & \Omega_{32} & \Omega_{42} & \Omega_{52} \\ \Omega_{31} & \Omega_{32} & \Omega_{33} & \Omega_{43} & \Omega_{53} \\ \Omega_{41} & \Omega_{42} & \Omega_{43} & \Omega_{44} & \Omega_{54} \\ \Omega_{51} & \Omega_{52} & \Omega_{53} & \Omega_{54} & \Omega_{55} \end{pmatrix} \right)$$

We jointly estimate the parameters in a hierarchical Bayesian framework. We assume conjugate priors and run Markov Chain Monte Carlo (MCMC) method with Gibbs sampling for 20,000 iterations. The first 40% of the iterations are treated as burn-in, and every other iteration from the remaining 12,000 runs are used for inference. Convergence of Gibbs samples was checked through visual inspection of the key model parameters. We log-transform the continuous variables in the estimation. Detail of the priors and estimation steps is described in the appendix.

## 4.2 Identification

Listing the relationships among the endogenous variables reveals the right-hand side of each equation to include only lagged endogenous variables or those already defined in the foregoing equations. Therefore, our model is fully recursive and identified (e.g., Wooldridge 2002)<sup>20</sup>.

$$\begin{aligned} ad\_stock_{ijt} &= f_1(.) \\ seeded\_placements_{ijt} &= f_2(seeded\_views_{ijt-1}, ad\_stock_{ijt}) \\ seeded\_views_{ijt} &= f_3(seeded\_placements\_cum_{ijt}, ug\_views_{ijt-1}, ad\_stock_{ijt}) \\ ug\_placements_{ijt} &= f_4(ug\_views_{ijt-1}, seeded\_placements\_cum_{ijt}, seeded\_views_{ijt}, ad\_stock_{ijt}) \\ ug\_views_{ijt} &= f_5(ug\_placements\_cum_{ijt}, seeded\_views_{ijt}, ad\_stock_{ijt}) \end{aligned}$$

Our proposed model assumes a temporal ordering among the endogenous variables, namely, offline advertising, seeded placements, seeded views, user-generated placements, and then user-generated views. Guiding by our conversations with industry practitioners, this ordering is reasonable: television advertising is usually determined months prior to the product release date and in current practice not adjusted in real time according to online video placements and views<sup>21</sup>; placements precede views and are often adjusted in real time; and advertiser-seeded videos precede the user-generated.

---

<sup>20</sup> We conducted a simulation analysis to confirm the recoverability of the model's parameters. We used the initial values of the five endogenous variables (i.e., values at five months prior to product release), set the parameter values, and generated the values for the endogenous variables in later time periods. The results of running the hierarchical Bayes estimation with the simulated data set showed our parameters can be recovered. The simulation data and results are available upon request.

<sup>21</sup> Our assumption that advertising is not affected by video placements and views may seem overly strong, considering the rich marketing literature that emphasizes the importance of treating advertising endogenously. As a robustness check, we fit a vector auto-regressive (VAR) model on each video and performed the Granger causality



### 4.3 Endogeneity

Two endogeneity concerns—omitted variables and simultaneity—are potentially relevant to this paper. First, movies and video games differ in their intrinsic characteristics: popular brands are likely to receive more offline advertising and their online advertisements would also tend to be more popular. In other words, intrinsic brand quality could be correlated with the level of advertising, placements, and views, but is unobserved by researchers. To at least partially alleviate this omitted variable concern, we include two important variables—production budget and critics review score, both of which are regarded as highly correlated with product quality in the movie and video game industry<sup>22</sup>, and should not be influenced by the scale of advertising campaigns used to promote the brand. Furthermore, by including multiple videos from the same brand, we account for the heterogeneity across brands and exploit variation across videos for the same brand, which further mitigates the omitted variable concern.

Second, simultaneity as a source of endogeneity bias is more challenging to address given the nature of the observational data. Our data being on a daily basis, the shorter observation interval makes simultaneity less a concern than for more aggregated data (e.g., collected on a weekly or monthly basis). Nevertheless, we have to make temporal assumptions for model identification, as described earlier. Discussions with industry practitioners confirm the reasonability of those assumptions.

---

test. We find that seeded placements, seeded views, user-generated placements, and user-generated views Granger-cause advertising spending in only 4%, 8%, 15%, and 11% of the videos, respectively, supporting our assumption that advertising is not endogenously determined by online video diffusion.

<sup>22</sup> To check their validity, we correlated these metrics with brand sales: the Pearson correlation between budget and sales was 0.48 ( $p < .001$ ) and between critic review and sales was 0.37 ( $p < .001$ ). This strengthened our confidence that the addition of these two variables helped address the omitted variable bias.

#### 4.4 Model Fit

To assess model fit, we compare our model (Model 1) to two alternative models that include the same five equations but set various restrictions on the time-varying parameters (i.e.,  $\delta_t$ ,  $\lambda_t$ ,  $\rho_t$ ,  $\psi_t$ , and  $\theta_t$ ). Model 2 specifies the time intercept and the advertising coefficient in each equation to be time-dependent (i.e.,  $\delta_{0t}$ ,  $\theta_{0t}$ ,  $\theta_{1t}$ ,  $\lambda_{0t}$ ,  $\lambda_{1t}$ ,  $\psi_{0t}$ ,  $\psi_{1t}$ ,  $\rho_{0t}$ , and  $\rho_{1t}$ ), and restricts all other dynamic parameters to be time-invariant. Model 3 furthers the restriction and allows only the time intercept in each equation to be time-varying (i.e.,  $\delta_{0t}$ ,  $\theta_{0t}$ ,  $\lambda_{0t}$ ,  $\psi_{0t}$ , and  $\rho_{0t}$ ).

We compare the model fit using two statistics: the log marginal likelihood (LML) and the deviance information criterion (DIC) suggested by Gelman, Carlin, Stern, & Rubin (2003). Our final model has the largest LML statistic and the smallest DIC, hence is considered better than the two more restrictive models.

	Model 1*	Model 2	Model 3
LML	-2,923,176	-2,924,693	-2,925,642
DIC	5,802,714	5,891,058	5,814,315

\*Model with best fit.

### 5 Findings

In describing our findings, we focus first on the interplay between advertiser-seeded and user-generated videos, and then discuss other drivers that influence video placements and views. When summarize our insights on how advertisers' strategies affect the success of viral campaigns, we trace the effect of advertisers' activities through the five equations and calculate their return on views. Table 2.4 presents the summaries for the dynamic parameter estimates, and Table 2.5 lists the estimates for the time-invariant parameters.

#### 5.1 Dynamics between Advertiser-seeded and User-generated Videos

First, our results confirm that, intuitively, placements drive views. This effect is captured by  $\lambda_{2t}$  for advertiser-seeded videos and by  $\rho_{3t}$  for user-generated videos. The average of the posterior means is estimated to be 1.55 for  $\lambda_{2t}$  and 1.56 for  $\rho_{3t}$  (see Table 2.4), indicating that, for both advertiser-seeded and user-generated videos, every 10% increase in placements brings about roughly 16% instantaneous increase in views<sup>23</sup>. The highest posterior density (HPD) intervals for  $\lambda_{2t}$  and  $\rho_{3t}$  do not cross zero at all time points, suggesting that the effect of placements on views is consistent and strong.

While this “return for placements” is almost identical in percentages (i.e., 16%) for advertiser-seeded and user-generated videos, the sheer number of user-generated placements makes them a force to be reckoned with when it comes to generating impressions. In addition, this 16% increase does not include the indirect effect implied by our system of equations. We leave the calculation of the total effect to section 5.5.

Second, our results shed light on the interplay between advertiser-seeded and user-generated videos. We find a positive spillover effect of advertiser-seeded placements and views on user-generated placements and views. Every 10% increase in advertiser-seeded placements brings about a 1.2% (the average posterior mean for  $\psi_{2t} = 0.12$ ) increase in user-generated placements and every 10% increase in seeded views brings about 0.8% (the average posterior mean for  $\rho_{2t} = 0.08$ ) contemporaneous increase in user-generated views. The HPD intervals for  $\rho_{2t}$  do not cross zero for 126 out of 150 (84%) days, although not throughout the entire time period. We also observe a positive feedback effect from user-generated views to advertiser-seeded views: for every 10% increase in user-generated views the advertiser-seeded views would grow by roughly

---

<sup>23</sup> With an average number of 3.56 cumulative seeded placements available on video-sharing websites, we have to be careful in extrapolating the elasticities to values beyond those covered in our sample.

1.1% (the average effect for  $\lambda_{3t} = 0.12$ ); and the corresponding HPD intervals are positive and do not cross zero for all 150 days, ranging between 0.06 and 0.19.

Third, we show how advertisers can influence impressions through offline advertising. Our results indicate that advertisers' offline advertising can drive both advertiser-seeded and user-generated views, although the timing of the effect is noticeably different. For advertiser-seeded videos, the effect of offline advertising on views is only positive and substantial two weeks before the product release date: the average posterior means for  $\lambda_{1t}$  during the last two weeks is estimated to be 0.57, corresponding to 0.5% more daily seeded views with every 10% increase in offline advertising. In contrast, offline advertising drives views much earlier for user-generated videos, concentrating between two and three months before the release date, although the magnitude of the effect is somewhat comparable: for every 10% increase in offline advertising, the average instantaneous growth in user-generated views is estimated to be 0.4%. We think that this difference is perhaps due to the intrinsic differences between users who watch user-generated and those who watch advertiser-seeded videos. Hard-core fans are more likely to watch user-generated content and also tend to pay attention to advertising earlier than casual fans do. This is consistent with the finding in user review literature that hard-core fans participate in online content generation earlier than casual fans (e.g., Li & Hitt 2008).

We do not find offline advertising to have a spillover effect on user-generated placements; the HPD intervals for  $\psi_{1t}$  contain zero for every time point  $t$ . That is, investing in television advertising has no direct impact on motivating online users to generate placements.

Table 2.4: Estimates of the average effect of the dynamic parameters

Equation	Dependent Variable	Right-hand-side Variable	Parameter	Estimates	
				Mean	SD
(1)	offline advertising	intercept	$\delta_t$	3.45	0.25
(2)	seeded placements	intercept	$\theta_{0t}$	0.02	0.02
		advertising	$\theta_{1t}$	3.6E-04	3.2E-04
		seeded views (lag)	$\theta_{2t}$	4.8E-04	3.9E-04
(3)	seeded views	intercept	$\lambda_{0t}$	1.13	0.19
		advertising	$\lambda_{1t}$	-0.01	0.004
		seeded placements	$\lambda_{2t}$	1.55	0.03
		ug views (lag)	$\lambda_{3t}$	0.12	0.01
(4)	ug placements	intercept	$\psi_{0t}$	-0.31	0.04
		advertising	$\psi_{1t}$	0.001	0.001
		seeded placements	$\psi_{2t}$	0.12	0.01
		seeded views	$\psi_{3t}$	2.9E-05	0.001
		ug views (lag)	$\psi_{4t}$	0.05	0.002
(5)	ug views	intercept	$\rho_{0t}$	1.12	0.15
		advertising	$\rho_{1t}$	0.01	0.003
		seeded views	$\rho_{2t}$	0.08	0.003
		ug placements	$\rho_{3t}$	1.56	0.02

Note: For ease of exposition, we present the mean effect of each dynamic parameter averaged over 150 days, which can be interpreted as the overall effect of the corresponding term. Standard deviations reflect the variation of the dynamic effect over time.

## 5.2 User ratings

Moving to how user ratings influence the placements and views for video advertisements, first, our results show that more favorable ratings attract more future views. Holding all else constant, one standard deviation increase in the average rating would boost advertiser-seeded views by 36.4% ( $\widehat{\gamma}_1^3 = 0.31$ , see Table 2.5) and user-generated views by 27.4% ( $\widehat{\gamma}_1^5 = 0.24$ ). Furthermore, the volume of user ratings drives seeded views: a standard deviation increase in

rating volume would bring about 4.8 ( $\widehat{\gamma}_2^3 = 1.76$ ) times more seeded views, but, somewhat surprisingly, it does not affect user-generated views ( $\widehat{\gamma}_2^5$  insignificant).

Second, estimates of  $\gamma_1^4$  and  $\gamma_2^4$  offer insights into how user ratings help motivate content generation. Consistent with the “selection effect” identified by Moe and Schweidel (2012), we find more favorable ratings to encourage and more negative ratings to discourage ( $\widehat{\gamma}_1^4 = 0.02$ ) the generation of future placements. Everything else being equal, one standard deviation increase in the average user rating leads to an additional 1.5% increase in user-generated placements. This positive association between rating valence and future placements also agrees with our estimate for the video sentiment metric: videos that receive more positive comments tend to be associated with more placements ( $\widehat{\eta}^4 = 0.07$ ). Collectively, these evidences seem to suggest that online content creators exhibit “jump on the bandwagon” behavior, tending to contribute to videos already “approved,” and to steer away from those perceived negatively, by the general public. Interestingly, the volume of user ratings is negatively related to advertiser-seeded ( $\widehat{\gamma}_2^2 = -0.03$ ) and user-generated ( $\widehat{\gamma}_2^4 = -0.25$ ) placements: if a video has already attracted a large number of user ratings, *ceteris paribus*, fewer additional placements will be generated in the future. We find the opposite effect of rating valence and volume to be interesting. While the average rating underscores users’ sentiment towards a video, the volume of the ratings more reflects its level of popularity (Liu 2006). Putting things together, our results seem to suggest that a video that is liked by users but has yet become popular tend to receive more future user-generated placements. Such a video is more likely to draw incremental viewership than one that is less favorably perceived or has already been popular. If we assume that the objective for users to upload a video is to attract more views, it is reasonable to think that videos that are favorably received but have yet become popular tend to attract more uploads from users.

### **5.3 Brand, video, and day characteristics**

Estimates for  $(\kappa, \eta, \text{ and } \tau)$  shed light on other factors that influence video placements and views. We find differences among genres: drama-themed brands attract the most advertiser-seeded views and users tend to generate their own copies of videos more for romance-themed brands. Brands from an existing franchise or rated as “restricted” (for movies) or “mature” (for games) also tend to have more user-uploaded videos. Production budget and critical reviews, however, do not predict the intensity of user-generated placements. In addition, we observe more video placements on weekdays and more viewing on weekends. Wednesdays experience the lowest level of viewing traffic.

Table 2.5: Posterior means and the 95% highest posterior density (HPD) intervals for time-invariant parameters

Equation No.	(1)		(2)		(3)		(4)		(5)		
Dependent Variable	offline advertising		seeded placements		seeded views		ug placements		ug views		
Variable	Effect	95% HPD	Effect	95% HPD	Effect	95% HPD	Effect	95% HPD	Effect	95% HPD	
rating_valence	$\gamma_1$	--	--	-0.001	(-0.004, 0.003)	0.31*	(0.27, 0.35)	0.02*	(0.01, 0.02)	0.24*	(0.22, 0.26)
rating_volume	$\gamma_2$	--	--	-0.03*	(-0.03, -0.02)	1.76*	(1.70, 1.81)	-0.25*	(-0.26, -0.23)	-0.004	(-0.05, 0.05)
sentiment	$\eta$	--	--	0.00	(-0.004, 0.01)	-0.06	(-0.18, 0.06)	0.07*	(0.04, 0.09)	0.05	(-0.05, 0.14)
Monday	$\tau_1$	-0.20*	(-0.32, -0.08)	0.02*	(0.01, 0.03)	-0.02	(-0.08, 0.05)	0.04*	(0.03, 0.06)	0.05	(-0.01, 0.10)
Tuesday	$\tau_2$	-0.21*	(-0.33, -0.08)	0.02*	(0.01, 0.02)	-0.04	(-0.1, 0.03)	0.03*	(0.01, 0.05)	0.01	(-0.04, 0.06)
Wednesday	$\tau_3$	-0.32*	(-0.43, -0.22)	0.02*	(0.01, 0.03)	-0.07*	(-0.13, -0.01)	0.04*	(0.02, 0.06)	-0.09*	(-0.13, -0.05)
Thursday	$\tau_4$	-0.32*	(-0.43, -0.21)	0.03*	(0.02, 0.03)	0.00	(-0.05, 0.06)	0.05*	(0.04, 0.07)	-0.02	(-0.07, 0.02)
Friday	$\tau_5$	-0.53*	(-0.66, -0.41)	0.02*	(0.01, 0.02)	0.02	(-0.05, 0.09)	0.03*	(0.01, 0.05)	0.01	(-0.04, 0.06)
Saturday	$\tau_6$	-0.37*	(-0.48, -0.26)	-0.01*	(-0.01, -0.001)	0.01	(-0.05, 0.08)	0.00	(-0.02, 0.02)	0.01	(-0.04, 0.06)
holidays	$\tau_7$	0.49*	(0.30, 0.68)	-0.01	(-0.01, 0.001)	0.02	(-0.06, 0.11)	0.00	(-0.02, 0.03)	-0.04	(-0.10, 0.03)
budget	$\kappa_1$	0.44*	(0.25, 0.63)	0.01	(-0.02, 0.03)	-0.03	(-0.24, 0.18)	-0.01	(-0.07, 0.05)	-0.23*	(-0.38, -0.06)
critic_review	$\kappa_2$	0.02	(-0.17, 0.20)	0.00	(-0.02, 0.02)	-0.01	(-0.18, 0.15)	-0.05	(-0.11, 0.003)	-0.06	(-0.20, 0.06)
franchise	$\kappa_3$	-0.01	(-0.37, 0.34)	0.00	(-0.05, 0.04)	-0.30	(-0.68, 0.07)	0.12*	(0.002, 0.23)	0.15	(-0.09, 0.39)
rating_R	$\kappa_4$	-0.15	(-0.52, 0.23)	-0.01	(-0.05, 0.04)	0.37	(-0.003, 0.73)	0.13*	(0.02, 0.25)	0.07	(-0.19, 0.31)
genre_action	$\kappa_5$	-2.38*	(-2.99, -1.73)	-0.01	(-0.09, 0.07)	-0.29	(-0.80, 0.30)	-0.36*	(-0.54, -0.19)	-0.48*	(-0.91, -0.07)
genre_adventure	$\kappa_6$	0.74*	(0.18, 1.27)	0.02	(-0.04, 0.08)	0.13	(-0.37, 0.66)	0.04	(-0.10, 0.19)	0.14	(-0.34, 0.56)
genre_drama	$\kappa_7$	-0.71	(-1.43, 0.03)	-0.02	(-0.11, 0.08)	0.02	(-0.73, 0.70)	0.13	(-0.09, 0.36)	-0.38	(-0.86, 0.15)
genre_romance	$\kappa_8$	-0.48	(-1.11, 0.14)	-0.01	(-0.07, 0.06)	1.03*	(0.42, 1.65)	-0.08	(-0.24, 0.09)	-0.23	(-0.73, 0.19)
genre_thriller	$\kappa_9$	0.16	(-0.56, 0.86)	0.02	(-0.06, 0.10)	0.52	(-0.13, 1.31)	0.25*	(0.04, 0.45)	0.25	(-0.30, 0.83)
genre_animation	$\kappa_{10}$	-0.09	(-0.72, 0.57)	-0.01	(-0.09, 0.07)	0.43	(-0.20, 1.05)	0.04	(-0.15, 0.24)	-0.42	(-1.00, 0.16)
genre_shooter	$\kappa_{11}$	0.61	(-0.31, 1.53)	-0.01	(-0.12, 0.09)	0.49	(-0.21, 1.17)	0.08	(-0.15, 0.32)	0.41	(-0.20, 1.03)
genre_other	$\kappa_{12}$	-0.17	(-1.17, 0.85)	0.01	(-0.11, 0.13)	0.26	(-0.54, 1.05)	0.25	(-0.04, 0.53)	-0.07	(-0.67, 0.52)

Note: \* indicates that the 95% HPD interval does not contain 0.



## 5.4 Baseline Diffusion Pattern

Lastly, we turn to the baseline temporal patterns for video placements and views. As stated in the model section, we use the time-varying intercepts (i.e.,  $\delta_{0t}$ ,  $\theta_{0t}$ ,  $\lambda_{0t}$ ,  $\psi_{0t}$ , and  $\rho_{0t}$ ) to capture the evolution of the five endogenous variables after controlling for the inter-variable dynamics. The estimates are depicted in Figure 2.2. We find substantial differences between advertiser-seeded and user-generated videos. Overall, the rate of adding an advertiser-created placement is relatively constant during a campaign's life-cycle. The baseline trajectory for advertiser-seeded views is also relatively stable, with a slight boost towards the brand release date. In contrast, user-generated videos experience strong growth over time, indicating that users in general are more engaged in disseminating videos when a brand is approaching the release date.

## 5.5 The Return of Advertisers' Activities

Synthesizing our results on how advertisers can affect the success of viral video campaigns yields additional insights on the magnitude of the "return on investment." We calculate the total impact for two advertiser-controlled instruments, namely, intensity of advertiser-seeded video placements and offline advertising spending, by tracing their direct and indirect effects throughout the system of equations.

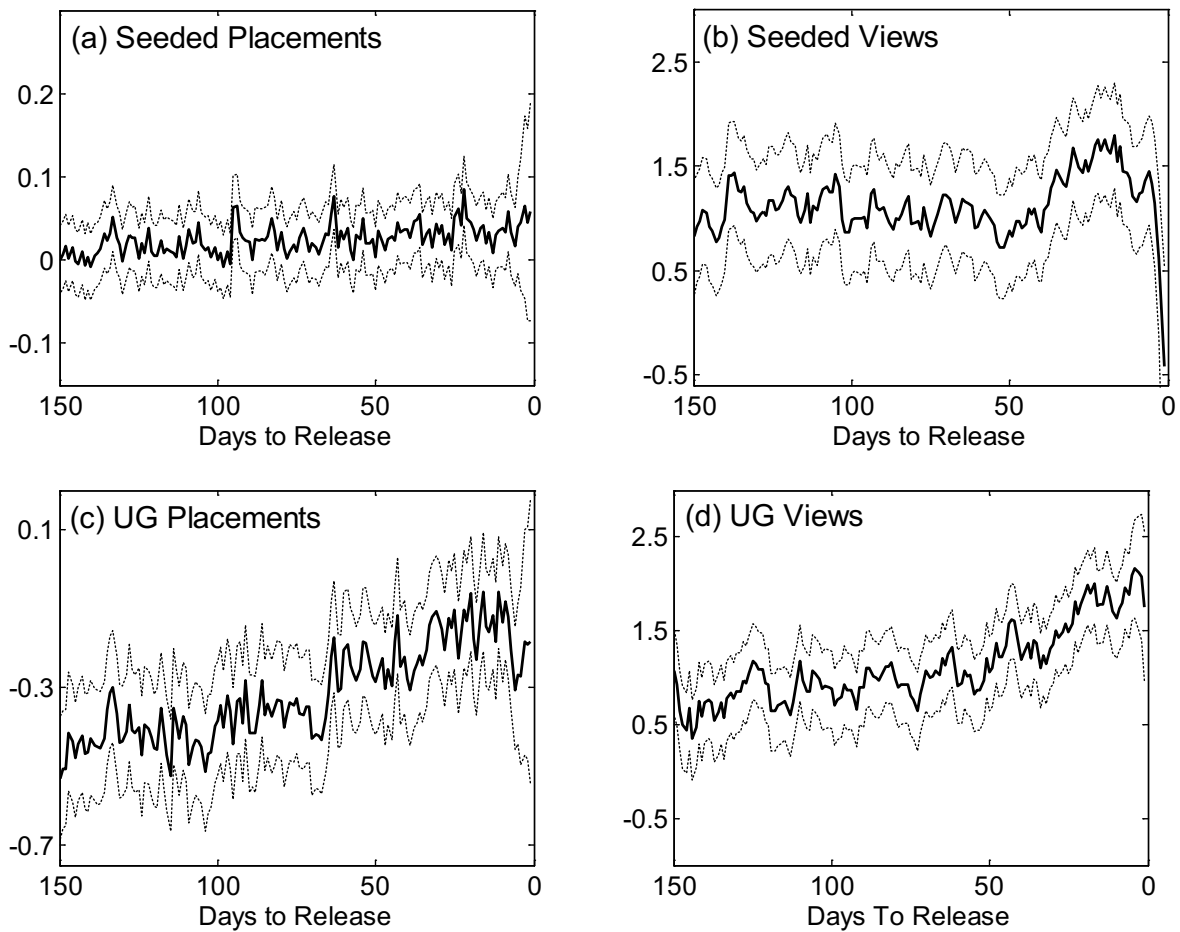


Figure 2.2: Baseline diffusion of video placements and views. Note: Figures 2.2a through 2.2d plot the posterior means and the 95% highest posterior density (HPD) intervals of the dynamic intercepts for advertiser-seeded placements, advertiser-seeded views, user-generated placements, and user-generated views, respectively. The solid line in each plot corresponds to the posterior means; the dashed lines represent the upper and lower bounds of the intervals. The estimates are plotted from 150 days to one day before the product release.

First, we constructed a time-series data matrix for a “typical” video by using the means of the five endogenous variables (i.e., offline ad spending, advertiser-seeded placements, advertiser-seeded views, user-generated placements, and user-generated views). Next, we computed the

return in views resulting from a shock to the advertiser-controlled instruments. The magnitude of the shock to seeded placements was set to 20%, which roughly corresponded to one additional placement. For a shock taking place on the 90<sup>th</sup> day before the product release, we increased the advertiser-seeded placements by 20% at the 90<sup>th</sup> day before release and held the values of the other variables constant. Then, we used the parameter estimates to predict the variables from the 89<sup>th</sup> day forward. To calculate the variable values in absence of the shock, we followed the same procedure but used the original data at the 90<sup>th</sup> day before release. The differences between the two simulated time series yield the effect for a 20% increase in seeded placements 90 days prior to the release. We repeated the procedure for each of the 6,000 draws of the parameter posterior distributions and summarized the results across the draws to derive the mean effect and the 90% HPD interval. The same steps were repeated for three other timings: 60 days, 30 days, and 10 days before product release. Following similar procedures we then computed the effect of offline advertising on online viewership generation.

Table 2.6 presents the changes in seeded and user-generated views resulting from a 20% shock in seeded placements and offline advertising, respectively. The first insight speaks to the relative effectiveness for advertiser-seeded placement intensity and offline advertising. Our calculation indicates that, compared with offline advertising, seeding intensity is a much more cost-effective instrument to drive online views. For example, a 20% increase in seeded placements 90 days prior to the product release would bring more than 40,000 views (approximately 440 daily views), among which 85% are from advertiser-seeded views and 15% from user-generated views. Ignoring the increase in user-generated views would lead to underestimation of advertisement reach. In contrast, the total effect of offline advertising support on views is positive and statistically significant, but small: a 20% increase in advertising (roughly

\$400,000 in our sample) over a 90-day window prior to the release yields about 12,000 total views (133 per day), a 70% reduction from the effect of increasing advertiser-seeded placements.

As far as the insights on the timing of the marketing investments is concerned, managers seeking to grow viewership for their videos will want to seed a placement closer to the brand release date. When moving a seed from 90 days before to 10 days before the release, the total effect on seeded views increases from 6.2% to 26.3% and that on user-generated views grows from 0.6% to 2.2%. Does this mean that managers should always wait until the last moment to upload their videos? Hardly. Although the boosting effect of advertiser-seeded placements is more pronounced closer to the release date, a longer duration provides more time for a video to grow an audience base. In terms of actual impression volume, it thus pays to seed earlier rather than later; a seed placed 90 days pre-release generates roughly 40,000 views while a seed 10 days pre-release does only about 23,000 total views. Taking these numbers together, when to place a seed depends on an advertiser's objective: if it is to generate buzz quickly to create a momentum, a couple of weeks or so pre-release would seem to be a good timing; if it is to increase overall reach, managers should consider putting some seeds available months before the release date.

Table 2.6: Returns for seeding intensity, timing, and offline advertising

20% Shock to Seeded Placements	Impression Percentage Change		Impression Volume Change		
	seeded views	ug views	seeded views	ug views	Total
90 days prior to release	6.2% (6.0%, 6.5%)	0.6% (0.5%, 0.6%)	33,959 (30,148, 37,824)	6,437 (4,225, 8,719)	40,396 (34,373, 46,543)
60 days prior to release	8.5% (8.2%, 8.9%)	0.9% (0.8%, 0.9%)	24,202 (21,544, 26,892)	7,137 (4,895, 9,461)	31,339 (26,439, 36,353)
30 days prior to release	13.5% (12.9%, 14.1%)	1.2% (1.1%, 1.3%)	28,965 (25,684, 32,268)	4,824 (3,174, 6,523)	33,789 (28,859, 38,791)
10 days prior to release	26.3% (24.8%, 27.8%)	2.2% (1.9%, 2.5%)	19,060 (16,924, 21,251)	3,680 (2,461, 4,932)	22,740 (19,385, 26,183)
20% Shock to Offline Advertising	Impression Percentage Change		Impression Volume Change		
	seeded views	ug views	seeded views	ug views	Total
90 days prior to release	0.73% (0.59%, 0.86%)	0.47% (0.38%, 0.56%)	5,774 (2,189, 9,422)	5,879 (1,436, 10,439)	11,653 (3,625, 19,862)
60 days prior to release	1.04% (.86%, 01.21%)	0.34% (0.27%, 0.41%)	5,528 (2,743, 8,369)	3,252 (1,183, 5,408)	8,780 (3,927, 13,777)
30 days prior to release	1.92% (1.58%, 2.25%)	0.33% (0.24%, 0.42%)	5,269 (2,723, 7,864)	1,651 (699, 2,669)	6,920 (3,421, 10,533)
10 days prior to release	3.79% (2.84%, 4.72%)	0.34% (0.25%, 0.44%)	3,129 (1,797, 4,491)	642 (326, 1,016)	3,771 (2,122, 5,506)

Note: we present the returns in campaign impressions (posterior means and the corresponding 90% HPD intervals) in response to a 20% shock in seeding intensity and offline advertising, respectively. Our calculations suggest that advertisers can influence the campaign reach through both levers; however, manipulating online seeding intensity turns out to be a much more cost-effective way to boost views than spending on television advertising.

The timing of offline advertising on seeded views works similarly to that of seeded placements: a shock in offline advertising closer to the release date can help create a momentum, but a longer duration allows more room for volume growth. However, delaying the offline advertising until product release does not help earn user-generated views, consistent with our finding that the significant direct effect of advertising on user-generated views happens two to three months prior to the product release. If managers were to use offline advertising to drive user-generated views, they want to place some pulses of TV advertising early in the advertising campaign.

## **6 Conclusions**

In this paper, we examine the diffusion of viral video campaigns, by modeling the spread of advertiser-seeded and user-generated videos as two interrelated processes. Although advertisers' decisions to place their advertisements on video-sharing sites may not always be aimed at triggering users to spread those messages, advertisers ignore that aspect of online channels at their peril.

Answering a call by Aral et al. (2013), we identify a number of instruments advertisers can employ to effectively influence how campaigns go viral. The extent to which users actively participate in, and thereby help to propel, online video campaigns is driven by the intensity and timing of advertiser-seeded placements. Campaigns adopting the right strategy on seeding intensity and timing are associated with higher numbers of impressions, both directly (through the relation between seeded placements and seeded views) and indirectly (through the spillover association between seeded and user-generated placements, and between seeded and user-generated views).

Our findings further yield insights into the dynamics between user ratings, views, and content generation; more views and better ratings help earn more future growth in views. There is also a tight interplay between advertiser-seeded and user-generated views: the popularity of one could spill over to the other. Interestingly, we uncover a pattern associated with user-generated video placements: videos that are rated high but have yet become popular are most likely to inspire user-generated copies.

### **Implications for Advertisers**

What are the implications for advertisers? It is a common practice for marketers in the entertainment space to approach videos-sharing sites with caution, and even to request sites like YouTube to remove user-generated copies of their advertisements. Film and video-game executives may have valid reasons for doing so. A desire to control the advertising environment and generate traffic to other company properties was often cited as reasons by practitioners with whom we spoke in the course of the present study. Our findings suggest, however, an accompanying strong negative effect in the severe curtailment of the spread of an online campaign. In fact, our study suggests that advertisers can increase advertising impressions through the low-cost practice of seeding videos with the express intent of letting users copy them and thereby expand campaign viewership.

The estimated limited role of offline advertising should be put in the proper context of our study. Buying more offline media seemingly does little to move the needle on the kinds of viral video marketing campaigns we examined. However, offline advertising may be very effective at generating impressions in traditional channels (as in fact research suggests, e.g., Elberse and Anand 2007), and it may well be that additional spending would greatly benefit the virality of smaller-scale entertainment products. More generally, our findings certainly do not indicate

advertisers should spend less on offline advertising—in fact, a certain threshold of awareness created by the large offline advertising budgets may be critical.

### **Future Research**

We believe three avenues for future research are particularly worthwhile. First, further research into the impact of online seeding advertising efforts on viral marketing campaigns for other product categories may be useful. Second, in our study, the overwhelming majority of user-generated advertisements consists of either exact copies or slightly altered derivative videos. In other contexts, users may often significantly alter the messages that advertisers seek to disseminate, possibly making those user-created versions far from helpful for advertisers. Research that investigates the balance between the positive and negative effects of viral campaigns in those contexts may be interesting. Third and finally, we think the marketing discipline would greatly benefit from studies that develop diagnostic tools for advertisers seeking to monitor and influence viral campaigns and their effects on market performance in real time. Controlled online experiments may be a relevant method in that respect.



The Air War versus the Ground Game:  
An Analysis of Multi-Channel Marketing in U.S. Presidential Elections

DOUG J. CHUNG

LINGLING ZHANG

Abstract

Firms increasingly use both mass-media advertising and targeted personal selling to promote products and brands. In this study, we jointly examine the effects of advertising and personal selling in the context of U.S. presidential elections, where the former is referred to as the “air war” and the latter as the “ground game.” Specifically, we look at how different types of advertising—the candidate's own advertising versus outside advertising—and personal selling—in the form of field office operations—affect voter preference. Furthermore, we ask how these campaign activities affect voting decisions through their diverse effects on various types of people. We compiled a unique and comprehensive dataset from multiple sources that record vote outcomes and campaign activities for the 2004-2012 U.S. presidential elections. Individuals' voting preference is modeled via a random-coefficient aggregate discrete-choice model, in which we incorporate individual heterogeneity and use instrumental variables to account for the endogeneity concern associated with campaign resource allocation. Among the many results, we find that personal selling has a stronger effect on partisan voters than on nonpartisans, while a

candidate's own advertising is better received by nonpartisans. We also find that outside ads behave very differently from candidate's own ads by mainly affecting partisan voters. Our findings may help candidates decide how to design effective campaigning by allocating resources both across multiple channels and within each channel, especially if the support from particular types of voters is weak.

## 1 Introduction

It is no secret that multi-channel marketing has increasingly been regarded as a competitive strategy critical to market success. Firms that understand the effect of and the dynamics behind their marketing channels are likely to reach customers more effectively and, hence, win over their customers. Among the channels, mass-media advertising and personal selling are usually the biggest arsenal at firms' disposal. Advertising has the advantage of reaching a large-scale audience via standard and well-scripted communication messages. Its importance goes without saying: global advertising spending was reportedly around \$128 billion in 2013.<sup>24</sup> Personal selling, on the other hand, happens at a micro level and takes the form of direct customer contacts, which may include regular and ad-hoc visits, distribution of fliers, and telemarketing, to name just a few. It often relies on a sales force to carry out the actual persuasion or mobilization, whether it is face-to-face or over the phone. Similar to advertising, personal selling is of great importance to many businesses. In the United States alone, the total spending on sales force has been reported to be more than four times the total spending on advertising (Zoltners et al. 2006), and approximately 11% of the nation's labor force is directly involved in sales or sales-related activities.<sup>25</sup> As advertising and personal selling are foremost in the minds of marketers, it is essential to understand the effect of the channels, in particular, the relative effectiveness of each on various types of consumers.

In this paper, we study mass-media advertising and personal selling in the context of U.S. presidential elections. Choosing the right product (the “president”) every four years is perhaps among the most critical decisions faced by many consumers of this country (the “voters”).

---

<sup>24</sup> *Nielsen*, 2013.

<sup>25</sup> *U.S. Bureau of Labor Statistics*, May 2013.

Presidential candidates carefully present themselves to people through strategic and expensive campaigns. The amount of marketing efforts behind each campaign is colossal: the 2012 election alone witnessed close to \$2 billion spending in campaigning across the Democratic and Republican candidates, making it one of the most expensive elections in the U.S. history and perhaps outnumbering any marketing campaigns that a consumer-packaged-goods company can possibly put together.<sup>26</sup>

Similar to the marketing of consumer-packaged-goods though, presidential campaigns have increasingly employed a multi-channel strategy. One notable phenomenon is the occurrence of large-scale personal selling efforts in the form of candidate's field operations during recent elections. President Barack H. Obama deployed an unprecedented field operation in 2008 such that many, including the *Denver Post*, attributed his election success to his on-the-ground efforts: "Obama's effective organization (of the field teams) could be a harbinger for how successful elections are won in battlegrounds in years to come."<sup>27</sup> Indeed, credit often goes to the winner's campaign for shaping the election results. For example, the day after President Obama was first elected, the *New York Times* claimed that "the story of Mr. Obama's journey to the pinnacle of American politics is the story of a campaign that was, even in the view of many rivals, almost flawless."<sup>28</sup> But, how much of this is true in reality? And if presidential campaigns are critical to voting outcomes, what can we marketers learn from them?

There are at least a couple of reasons why presidential elections provide a good setting for jointly studying the effect of advertising and personal selling. First, because campaign activities

---

<sup>26</sup> *The New York Times*, "The Money Race", 2012.

<sup>27</sup> Sherry Allison, "Ground Game Licked G.O.P." *The Denver Post*, November 5, 2008

<sup>28</sup> *The New York Times*, November 5, 2008.

vary substantially between contested and non-contested states as well as across counties within each contested state, presidential elections yield the much-needed geographical data variation. And because the competitive landscape changes from one election to another, the changes in campaign resource allocation also provide data variation along the time dimension. Second, political campaigns primarily serve a short-term goal to make “sales” happen (i.e., win votes), rather than to build a brand or maintain customer relationships. Therefore, the potential long-term effect of advertising and personal selling is less relevant in our context, simplifying the analysis and allowing us to focus on the causal influence of the campaign effect.

In this study we are interested in two questions. First, how much do mass-media advertising and personal selling matter on voter’s preference for presidential candidates? Second, how do those campaign activities affect voters through their diverse effects on various types of people? Answers to these questions not only help political campaign organizers but also marketers in general, as long as multi-channel marketing is engaged.

However, assessing the effects of advertising and personal selling turns out to be non-trivial, due to the challenge in obtaining comprehensive data and the difficulty of making causal inference. As far as advertising is concerned, much research has been conducted to understand its effect in consumer product marketing (Bruce 2008; de Kluyver and Brodie 1987; Dekimpe and Hanssens 1995; Givon and Horsky 1990; Lodish et al. 1995) as well as in political campaigns (Gordon and Hartmann 2013; Shachar 2009; Shaw 1999). A constant challenge is that advertising is often studied in isolation from other instruments of the marketing mix (Albers et al. 2010). The few exceptions in the context of political campaigning are Shachar (2009) and Shaw (1999). Shachar (2009) examined the relationship between the intensity of competition and two campaign activities—television advertising and grassroots campaigning—for the 1996-2004

presidential elections. His main finding is that close competition caused more campaigning, which further led to higher turnout rates. If anything, Shachar (2009) provides empirical evidence that campaign activities are endogenously determined according to the competition intensity. Built upon this finding, our paper addresses the endogeneity concern in estimating the causal effect of campaigning. More importantly, we also examine how the effect of campaigning varies by voter characteristics; hence, our results could provide more direct implication on allocating campaign resources. Shaw (1999) is another example studying those two campaign activities and found them to affect the statewide voting preferences. However, his results are generated via regression models on state-level observational data and may not adequately account for the endogeneity concern underlying campaign variables.

Compared to advertising, our knowledge on personal selling is even more limited, despite that it has long been regarded as an essential element of the marketing mix (Borden 1964; Weitz 1981). Many extant marketing papers only have aggregate-level measures on personal selling (e.g., Gatignon and Hanssens 1987; Narayanan et al. 2004) and hence cannot generate insights on where to allocate the sales force, an important implication for marketers aiming to yield the best possible outcomes. Researchers in political science also have difficulty collecting reliable data on personal selling. For example, Shachar (2009) used self-reported campaign contacts to measure grassroots efforts, which is restrictive in that one third of the states in his data had fewer than five respondents. Not surprisingly, a noisy measure like this prompted researchers to adopt an experimental design to study the effect of personal selling on voting. Gerber and Green (2000) conducted field experiments and found that face-to-face visits increase turnout rates. Alvarez et al. (2010), through another field experiment, concluded that delivering partisan messages in person can have an even bigger effect than previously reported. However, those results are not

exempt from common critiques for experiments: data usually come with a limited scale and the external validity of the results may be questionable. With the recent developments in data collection methods, better measures of personal selling have now become available for political scholars. Masket (2009) examined the placement of Democratic field offices in the 2008 presidential election and found them to significantly boost the vote shares. Darr and Levendusky (2014) tracked the deployment of field offices in several recent presidential elections and quantified the magnitude to be around 1% vote share increase per one additional field office. However, neither studies adequately addressed the endogeneity concern associated with allocating field offices; hence, their results are correlational rather than causal.

Finally, empirical papers on the campaign effect largely leave out individual heterogeneity, perhaps because individual-level characteristics are challenging to obtain on a large scale. However, understanding how various marketing activities may have a diverse effect on different segments of individuals is essential for designing targeted marketing and allocating resources. For example, Carroll et al. (1985) jointly estimated the effect of salesforce and advertising on Navy enlistment through a large scale field experiment. As they pointed out, one limitation of their study is the inability to examine heterogeneous marketing effects using aggregated campaign data. We have reasons to believe that incorporating heterogeneity is important in the context of political elections, because people with varying predisposition may likely respond differently to different marketing activities. One paper exploring this issue is Lovett and Peress (2015). They combined political advertising data with viewer profiles of television shows, and found that political advertising is primarily effective on the segment of swing voters. However, their paper only included the 2004 election and did not control for other campaign activities, both of which the authors acknowledged as a limitation.

To better understand the campaign effect in presidential elections, we set out to compile a unique and comprehensive dataset and carefully design our analysis to jointly examine advertising and personal selling while addressing endogeneity and consumer heterogeneity. Our data are integrated from multiple sources and include a total of 18,650 observations on vote outcomes and campaign activities. We collect detailed records of field operations for candidates from both parties, down to the county level. Our data on television advertising cover ad impressions at the designated-market-area (DMA) level and include not only the ads made by candidates but also those by outside political groups. The rapid growth of outside advertising in recent presidential elections, especially the 2012 one, has made it too important to be ignored. In addition, we control for the total candidate spending in digital campaigning—a relatively new channel that has started entering the candidates' marketing toolkit—as well as a large number of other control variables that signal the economic and political climate of the elections.

We model individuals' voting preference via a random-coefficient aggregate discrete-choice model, which allows the various campaign effects to differ by voter characteristics. Further, we use instrumental variables to account for the endogeneity concern associated with campaign activities. Our results show that field operations and advertising both have positive effects on voter preference. An addition of a field office would increase the vote share in a county by 1.143% for the Republican candidates and 3.305% for the Democrats, indicating a clear effect yet asymmetrical between the parties. We estimate the elasticity of candidate's own ads to be 0.059 for the Republicans and 0.081 for the Democrats, whereas the elasticity of outside ads is 0.032 and 0.045, respectively. We also find evidence that campaign effects depend on voters' baseline partisanship: field operations, often involving volunteers making face-to-face contacts with voters, are more effective among partisans than non-partisans, while candidate's



own advertising is only effective among non-partisans. Interestingly, we find that outside ads, which typically consist of negative and attacking messages, behave more like field operations than candidate's own ads, suggesting an interaction between the tone of ads and voters' partisan preferences.

To quantify the importance of campaign activities, we conduct counterfactual analyses using the parameter estimates. Overall, our estimates suggest that campaigns play an essential role in deciding the outcome of an election. Had field operations not been allowed in presidential elections, history would have been rewritten, with a different president being elected in 2008 and 2012. Interestingly, had the Democrats received more outside ads in 2004, the election would have ended up in a tie of 269 electoral votes on each side.

Our paper contributes to the literature in two ways. First, we jointly estimate the effect of mass-media advertising and personal selling, two of the most prominent marketing activities. Our data set includes almost all major campaign activities employed by recent presidential candidates and spans multiple election years, making it much more comprehensive than other extant data sets. One innovation of this paper is that we separate outside ads from candidate's own ads. To our best knowledge, this is one of the first attempts to systematically examine the effect of political ads sponsored by outside interest groups in presidential elections. Second, we carefully address the endogeneity concern for campaign activities and are able to make inference on the heterogeneous channel effect using only market-level aggregate data, which are more readily available than individual responses in many contexts. Therefore, our estimates can help allocate marketing resources both across different channels and within a channel across customer segments.

It is perhaps worthwhile to compare our paper with Gordon and Hartman (2013), which also studies the effect of advertising in U.S. presidential elections. This paper differs in at least four aspects: first, we have much more comprehensive data including almost all of the major marketing instruments utilized by campaigns—mass media advertising, personal selling, and digital campaigning; second, we distinguish between candidate own and outside advertising, and find them to have different effects; third, we have actual GRP data whereas Gordon and Hartmann (2013) estimate their ad exposure data; lastly and perhaps most importantly, we incorporate consumer heterogeneity and examine how different campaign instruments affect different types of voters.

The remainder of the paper is organized as follows. Section 2 describes the campaign activities and the data used for empirical analysis. Section 3 specifies the model and discusses the identification. Section 4 presents the parameter estimates and the counterfactual results. Section 5 concludes.

## **2 Data**

We compiled a unique dataset from multiple sources that includes actual voting outcomes and campaign activities for the 2004-2012 U.S. presidential elections. Our data are superior to those used in extant studies in at least four aspects. First, our data span a period of three presidential elections and, thus, the results are not confined to a particular combination of candidates. Second, our collection of multiple campaign activities encompasses a more comprehensive record of mass-media advertising and ground campaigning than ever seen before in previous studies. Knowing where and to what extent candidates choose to campaign enables us to assess the effects of various campaign activities after controlling for one another. Third, our unit of analysis is at the county level, which is as granular as it can be to reliably obtain the

voting outcomes. In addition, we also measure campaign activities at a granular level when possible. By having disaggregated data we are able to take a finer look at the campaign allocation and curtail the potential aggregation bias. Finally, we collect data on the registered party affiliation at the county level, which enables to examine how campaign effects differ according to voter partisanship.

## **2.1 Election Votes**

The dependent variable for this study is the number of votes cast for the presidential candidates in each county. We collected this variable from the CQ Press Voting and Elections Collection, a database that tracks major U.S. political elections.

We define each county as a “market” in which residents choose up to one “product” (candidate). In the subsequent analysis, we will use “market” and “county” interchangeably. We define a county's “market size” as the total number of resident citizens aged 18 and above, typically known as the Voting Age Population (VAP)<sup>29</sup>. We obtained the county-level age-specific population counts from the U.S. census database. The “market share” of each candidate is then the percentage of votes he or she receives out of the county VAP.

There are a total of 3,144 counties and county equivalents in the United States. We exclude Alaska from the analysis because its voting outcomes and population estimates are measured on different geographical units and hence are challenging to match. As a result, we end up with 3,110 counties from 49 states plus the District of Columbia, which is treated as a single market in the analysis. Table 3.1 lists the county-level summary statistics for vote outcomes. The

---

<sup>29</sup> A perhaps better measure for the market size of a county is the Voting Eligible Population (VEP), which equals the VAP minus ineligible felons. This metric, however, is available only at the state level. For a good introduction on how to estimate the Voting Eligible Population, see the United States Elections Project.

Republican candidate, George W. Bush, won more of the popular votes in the 2004 election, and the Democratic candidate, Obama, won more of the popular votes in the 2008 and 2012 elections. The average county-level vote share is always higher for the Republican candidates. The Republicans won many less-populated counties in 2008 and 2012, although they still lost to the Democrats in the total popular votes and electoral votes at the national level.

Table 3.1: Summary Statistics for Vote Outcomes by County

Year	Party	N	Mean	SD	Min	Max	Total
Votes							
2004	Democrat	3111	18,902	65,678	18	1,907,736	58,802,968
	Republican	3111	19,866	47,586	82	1,076,225	61,804,121
2008	Democrat	3106	22,289	77,146	8	2,295,853	69,230,895
	Republican	3106	19,221	44,883	94	956,425	59,701,115
2012	Democrat	3108	21,127	74,225	5	2,216,903	65,661,169
	Republican	3108	19,537	44,789	84	885,333	60,721,119
Vote Share							
2004	Democrat	3111	0.22	0.08	0.04	0.58	
	Republican	3111	0.35	0.09	0.05	0.76	
2008	Democrat	3106	0.25	0.10	0.04	0.64	
	Republican	3106	0.33	0.09	0.04	0.78	
2012	Democrat	3108	0.22	0.10	0.02	0.70	
	Republican	3108	0.33	0.10	0.02	0.73	
Combined Votes							
2004		3111	38,768	109,874	155	2,983,961	120,607,089
2008		3106	41,511	118,264	159	3,252,278	128,932,010
2012		3108	40,664	114,908	144	3,102,236	126,382,288
Turnout Rate							
2004		3111	0.57	0.09	0.19	0.98	
2008		3106	0.58	0.09	0.16	0.90	
2012		3108	0.55	0.09	0.15	0.99	

Note: We calculate the turnout rate as the sum of votes for the Democratic and the Republican candidates divided by the number of resident citizens aged 18 and above. The vote share for each candidate is calculated as the ratio of the focal candidate's received votes divided by the number of resident citizens aged 18 and above.

## 2.2 Ground Campaigning

To an average voter, presidential elections are perhaps most visible on the ground level through personal selling activities (henceforth, we will use ground campaigning and personal selling interchangeably). In the early stage of each election, presidential candidates establish field operations to organize the ground-level voter outreach; in particular, campaigns set up field offices from which staff and volunteers coordinate their door-to-door canvassing efforts, conduct telemarketing campaigning, and organize other outreach activities. We measure the scale of a candidate's field operations by the number of field offices deployed in each county. We collected the 2004 and 2008 field office data from the “Democracy in Action” project at George Washington University<sup>30</sup> and the 2012 data from Newsweek Daily Beast<sup>31</sup>, both of which scoured the Democratic and Republican campaign websites and gathered addresses for all the field offices. We then used the Geographic Information System (GIS) software to map the office addresses onto the corresponding county.

Table 3.2 displays the summary statistics for field operations. Across all elections, the Democratic candidates had an indisputable lead in establishing field operations: the ratio of the Democratic and Republican field offices was 3.51, 3.53, and 2.69 in 2004, 2008, and 2012, respectively. Furthermore, even between the Democratic candidates, field offices were more prominent in the Obama campaign than in the John Kerry campaign: while the latter had at least one field office in 237 (8%) counties, the former set up offices in 624 (20%) counties in 2008 and 439 (14%) in 2012.

---

<sup>30</sup> The URL for the project is: <http://www.gwu.edu/~action>. Accessed on 8/2/2013.

<sup>31</sup> The explanation of the data tracking method can be found at <http://newsbeastlabs.tumblr.com/post/34109019268/tracking-the-presidential-groundgame-as-the-two>. Accessed on 8/6/2013.

Table 3.2: Ground Campaigning by County

		N	Mean	SD	Min	Max	Total
Number of Field Offices							
2004	Democrat	3,111	0.10	0.45	0	12	313
	Republican	3,111	0.03	0.23	0	5	89
2008	Democrat	3,106	0.28	0.75	0	11	874
	Republican	3,106	0.08	0.46	0	17	247
2012	Democrat	3,108	0.24	0.93	0	21	750
	Republican	3,108	0.09	0.37	0	6	278
Presence of Field Offices							
2004	Democrat	3,111	0.08	0.27	0	1	237
	Republican	3,111	0.02	0.14	0	1	65
2008	Democrat	3,106	0.20	0.40	0	1	624
	Republican	3,106	0.06	0.24	0	1	192
2012	Democrat	3,108	0.14	0.35	0	1	439
	Republican	3,108	0.07	0.26	0	1	222

Note: The unit of observation is county. Field operations are measured by the number of field offices in each county. We also report the number of counties that had at least one field office.

It merits mentioning that we use the number of field offices in each county as a proxy of voters’ exposure to candidates’ field operations. This metric, becoming available only in recent elections, provides a more objective measure of field operations at a granular geographical unit than other alternative survey-based measures, which are prone to recall errors and non-response bias. Due to this advantage, the number of field offices has been used by several recent papers studying the effect of ground campaigning (e.g., Darr and Levendusky 2014; Masket 2009). However, it is not without limitations. For example, a field office may serve multiple purposes—coordinating voter contacts, organizing fund-raising events, or even laying groundwork to raise voter support for future party candidates—some of which may not be directly related to winning votes for the current election (Darr and Levendusky 2014). Yet, the primary goal for a field office during the general election should be around the target of “Race to 270”. Therefore, the

number of field offices should still indicate the degree to which a candidate uses ground campaigning to gain votes.

Nevertheless, a question may remain: how much does the number of field offices reflect the extent of voter exposures to ground campaigning. One way of assessing the validity of this metric is to correlate it with the number of voter contacts made by the ground campaign personnel, which we obtained from the American National Election Studies (ANES), a high-quality survey on voting and political participation. In the ANES 2004, 2008, and 2012 time series surveys, respondents were asked whether they had been contacted by a party about the campaign, and if yes, by which party. Based on the ANES responses we estimated the (weighted) number of respondents contacted by the Democratic and the Republican campaign teams, respectively. The correlation between voter contacts and the number of field offices was 0.76 for the Democrats and 0.73 for the Republicans, indicating that the number of field offices has a moderate to strong correlation with the self-reported individual exposure to ground campaigning. The ANES responses cannot be used in our analysis because they are only available at the census-region level. Therefore, we believe that the number of field office is the best proxy for field operations currently available to researchers. We acknowledge its limitations and think future research can benefit from improving the measurement for this variable.

### **2.3 Television Advertising**

There are three types of ad sponsors in the U.S. presidential elections: the candidates, their party committees—namely, the Democratic National Committee (DNC) and the Republican National Committee (RNC)—and some outside political groups. Because the candidates and party committees often coordinate advertising efforts, we combine the ads from these two types and label them as the candidate's own advertisements.

The third type of player—outside political groups, also known as the Political Action Committees (PACs)—buys television ad spots to support their preferred candidates and to attack their rivals. Although they have played a role in presidential elections for decades, PACs particularly took on a much greater prominence in recent elections, partly because, in 2002, a campaign finance reform law set stricter restrictions on fund-raising and spending, hence the PACs stepped in to fill the gap. Especially in the 2012 election, a relatively new kind of organization, the Super PAC, emerged as a major advertiser. Super PACs are made up of independent PACs that support a candidate with unlimited—and often anonymous—donations from unions, companies, or individuals. Due to the large number of PACs advertising in the presidential elections, it is challenging to track all of their ads. Fortunately, we are able to obtain the data for the top ad spenders, which, combined, are responsible for more than 90% of the total ad spending by the PACs.

We measure advertising using the gross rating points (GRPs), which quantify advertising impressions as a percentage of the target audience being reached. For example, if an ad aired in the Des Moines-Ames area reaches 25% of the target population, it receives a GRP value of 25; if the same ad is aired five times, the GRP value would be 125 ( $=5 \times 25$ ). GRPs are a better measure of ad exposures than dollar spending because the cost of advertising varies significantly across markets. For example, the same amount of ad dollars would yield far less exposure in Los Angeles than in Kansas City. Hence, GRPs provide a measure of audience reach, independent of the advertising cost.

We obtained television advertising data from Nielsen Media Research. Nielsen divides the U.S. media market into 210 designated market areas (DMA): residents from the same DMA receive largely the same television offerings, including advertising. Therefore, our advertising



metrics are measured at the DMA level. It is noteworthy that our outcome variable of interest is at the county level, with each county belonging to one and only one DMA. To link ad impressions to county-level votes, we assume that the percent of the audience reached in a county equals the percent of the audience reached in the DMA to which the county belongs. Take the Rochester-Manson City-Austin DMA, for example: This DMA consists of seven counties from Iowa (Cerro Gordo, Floyd, Hancock, Howard, Mitchell, Winnebago, and Worth) and five counties from Minnesota (Dodge, Fillmore, Freeborn, Mower, and Olmsted). During the week of October 21, 2012, Obama campaign ads reached 1048.8% of the DMA population. By assuming that advertising impressions are homogeneous within a DMA, we assign the Democratic candidate's own GRP value to be 1048.8 for each of the twelve counties during that week.

Because voting preference is revealed on Election Day, we calculate the cumulative GRPs that each DMA has received since September 1 of that year and use this cumulative measure in the subsequent analysis. Table 3.3 presents the summary statistics for candidate's own advertising and PAC advertising, respectively. For candidate's own advertising, the Democratic candidates outnumbered the Republicans by 20%, 50%, and 40% in the three elections, respectively. Interestingly, the PACs, which had less advertising than the candidates in 2004 and 2008, played a much bigger role in the 2012 election. In particular, the PAC ads supporting Mitt Romney were responsible for roughly 46% of the total advertising for Romney and outnumbered the PAC ads supporting Obama by almost seven times. Even though the Obama campaign had more advertising than the Romney campaign, the PACs filled the gap; in the end, 25% more pro-Romney ads were aired than pro-Obama ads in the 2012 election.

Table 3.3: Television Advertising by DMA

	N	Mean	SD	Min	Max	Total
Candidate Advertising (GRPs)						
2004 Democrat	206	1,420.4	2,374.1	0	8,933	292,611
Republican	206	1,885.9	2,764.0	0	8,440	388,501
2008 Democrat	206	3,809.9	3,797.8	255	13,838	784,848
Republican	206	2,075.3	2,422.7	77	8,452	427,517
2012 Democrat	206	2,232.2	4,143.7	0	15,779	459,827
Republican	206	1,320.7	2,576.5	0	9,535	272,055
Party Advertising (GRPs)						
2004 Democrat	206	1,942.0	2,211.0	0	7,561	400,054
Republican	206	868.0	1,271.1	0	5,858	178,814
2008 Democrat	206	1,766.5	2,324.9	0	12,277	363,905
Republican	206	1,553.5	1,890.5	0	11,035	320,013
2012 Democrat	206	1,063.8	1,604.2	0	7,270	219,144
Republican	206	1,069.4	1,774.5	0	11,044	220,291
Candidate and Party Advertising (GRPs)						
2004 Democrat	206	3,362.5	4,268.0	0	16,120	692,665
Republican	206	2,754.0	3,204.2	0	11,579	567,316
						1,148,75
2008 Democrat	206	5,576.5	4,760.6	255	18,418	4
Republican	206	3,628.8	3,713.1	77	17,965	747,530
2012 Democrat	206	3,296.0	4,941.1	0	19,849	678,971
Republican	206	2,390.0	3,611.8	0	19,597	492,346
PAC Advertising (GRPs)						
2004 Democrat	206	255.9	505.0	0	2,248	52,726
Republican	206	394.4	866.2	0	4,023	81,250
2008 Democrat	206	159.4	407.3	0	2,513	32,830
Republican	206	217.2	435.7	0	2,188	44,736
2012 Democrat	206	254.3	694.6	0	3,840	52,378
Republican	206	2,030.9	2,714.9	67	12,137	418,356
Total Advertising (GRPs)						
2004 Democrat	206	3,618.4	4,627.9	0	16,726	745,390
Republican	206	3,148.4	3,773.8	0	12,413	648,566
						1,181,58
2008 Democrat	206	5,735.8	4,951.8	255	19,592	3
Republican	206	3,846.0	3,965.4	77	19,704	792,267
2012 Democrat	206	3,550.2	5,533.5	0	22,943	731,349
Republican	206	4,420.9	6,091.8	67	29,295	910,702

Note: We measure television advertising using gross rating points (GRPs), which correspond to the percentage target population reached in each DMA. For PAC advertising, we obtained data

for the top spenders, which were responsible for more than 90% of the total PAC ad spending for each election. The total number of DMAs excludes those in Alaska.

## **2.4 Digital Campaigning**

In addition to field operations and television advertising, we also collected online campaigning data for the three elections. Online digital campaigning in political elections started to attract the mainstream's attention during the 2004 election, when the Democratic candidate, Howard Dean, adopted the then-innovative web-based campaign initiatives to raise a remarkable level of support in the early stages of the election. Since then, online campaigning has appeared on the radar screens and the candidates have been experimenting to incorporate the Internet into their standards of campaign activities. Understandably, the 2004 race largely regarded the web as a tool for fund raising or for insider communication rather than for advertising; a small amount of resources was dedicated to online campaign activities. On the Democrats side, the Kerry camp reportedly bought a \$1.3 million<sup>32</sup> worth of online ads and the DNC \$257,000; and on the Republicans side, the Bush campaign spent roughly \$419,000 and the RNC \$487,000<sup>33</sup>. In the next election cycle, a substantive increase in online campaigning was witnessed for candidates of both parties. In 2008, the online ad spending for the Obama campaign outnumbered that for the McCain campaign by 4:1, with roughly \$16 million for the former and \$4 million for the latter<sup>34</sup>. Digital campaigning grew more than three fold in the 2012 race: the Obama campaign spent \$52

---

<sup>32</sup> The 2004 and 2008 spending was inflated to the 2012 dollars.

<sup>33</sup> Pew Internet & American Life Project, 2004.

<sup>34</sup> Borrell Associates, 2014

million on online ads, followed by \$26.2 million by the Romney campaign.<sup>35</sup> We control for the total online campaign spending in our analysis.

## **2.5 Additional Variables**

We include a rich set of control variables that reflect the economic and political climate and may influence voter preference (see Table 3.4 for summary statistics). First, the presidential incumbency status captures the advantage for the incumbent candidates, as inertia alone has been shown to be able to generate votes (Campbell 1992). We assign 1 to the incumbent presidential candidates and 0 otherwise.

Second, we control for three state-level variables: (1) the home state advantage for presidential candidates, (2) the home state advantage for vice-presidential candidates, and (3) the governor advantage of the state. The home state variables take a value of 1 if the candidates are from the focal state and 0 otherwise. The governor advantage variable is also an indicator: for each campaign-state-party combination, the observation receives 1 if the governor of the state is from the same party that year, and 0 otherwise.

Lastly, we also include three sets of county-level contextual factors. The first is the percentage of African American residents to capture the racial composition of a county. The second group of variables, indicating the socio-economic conditions of the county, includes the median household income, the unemployment rate, the Gini index, the median house value, the percentage of residents dropping out of high school, and the percentage of residents living in poverty. Those variables are obtained from the U.S. Census Bureau databases. The third variable is the percentage of registered partisan voters, which we acquired from a proprietary database

---

<sup>35</sup> <http://www.businessinsider.com/infographic-obama-romney-final-ad-spend-2012-11>. Accessed on December 18, 2014.

tracking election data<sup>36</sup>. Based on this variable, we simulate the individual voter-level partisanship; we describe the simulation in more detail in Section 3.1.

Table 3.4: Summary Statistics of Additional Variables

	N	Mean	SD	Min	Max
Incumbent status					
2004	6,222	0.50	0.50	0	1
2008	6,212	0.00	0.00	0	0
2012	6,216	0.50	0.50	0	1
Home state advantage for presidential candidates					
2004	6,222	0.04	0.20	0	1
2008	6,212	0.02	0.14	0	1
2012	6,216	0.02	0.14	0	1
Home state advantage for vice-presidential candidates					
2004	6,222	0.02	0.14	0	1
2008	6,212	0.00	0.02	0	1
2012	6,216	0.01	0.11	0	1
Governor advantage					
2004	6,222	0.50	0.50	0	1
2008	6,212	0.50	0.50	0	1
2012	6,216	0.50	0.50	0	1
Percentage of African American residents					
2004	6,222	0.09	0.14	0	0.87
2008	6,212	0.09	0.15	0	0.86
2012	6,216	0.09	0.15	0	0.86
Median household income (\$)					
2004	6,222	46,458.1	12,257.3	20,193	121,266
2008	6,212	46,528.0	12,332.4	19,744	122,822
2012	6,216	44,901.4	11,549.5	19,624	122,844
Unemployment rate					
2004	6,222	0.07	0.03	0	0.36
2008	6,212	0.08	0.03	0	0.28
2012	6,216	0.09	0.04	0	0.27

<sup>36</sup> Because not all states require voters to declare party affiliation during registration, we have partisan information for 27 states in 2004 and 2008 and 28 states in 2012. Data come from a repository tracking U.S. elections (<http://uselectionatlas.org/>), where partisan numbers are extracted from various official websites such as the state's Secretary of State and the Office of Elections.

Gini index						
2004		6,222	0.43	0.04	0	0.62
2008		6,212	0.43	0.04	0	0.67
2012		6,216	0.44	0.04	0	0.60
Median house value (\$)						
2004		6,222	139,215.1	100,886.0	31,463	1,070,185
2008		6,212	141,710.5	95,805.9	17,148	1,014,468
2012		6,216	129,529.0	77,297.9	19,400	944,100
Dropout rate						
2004		6,222	0.07	0.06	0	0.58
2008		6,212	0.07	0.05	0	0.60
2012		6,216	0.06	0.05	0	0.63
Poverty rate						
2004		6,222	0.15	0.07	0	0.52
2008		6,212	0.16	0.07	0	0.53
2012		6,216	0.17	0.07	0	0.49
Percentage of registered partisans						
2004	Democrat	1,318	0.35	0.17	0.06	0.98
	Republican	1,318	0.33	0.16	0.03	0.90
2008	Democrat	1,319	0.35	0.16	0.06	1.00
	Republican	1,319	0.33	0.15	0.03	0.93
2012	Democrat	1,349	0.31	0.16	0.02	0.97
	Republican	1,349	0.33	0.15	0.03	0.87

Note: County-level control variables are obtained from the American Community Survey database. Data on registered voters by party are compiled from various official government sources. Some states do not require voters to declare party affiliation, hence, we have a smaller sample size for this variable.

## 2.6 Model-free Evidence

### 2.6.1 Campaign Effects

In this section, we present some model-free evidence. We first examine how ground campaigning and television advertising are related to vote shares. To account for the large cross-sectional variation across counties, we calculate the changes in vote shares and campaign activities from one election to the next and then examine the relationship between the changes.

Figure 3.1 depicts the association between vote shares and ground campaigning. The vertical axis of the figure corresponds to the change in vote shares—i.e.,  $s_{cj,t+1} - s_{cj,t}$ , where the vote share in county  $c$  for party  $j$  during election  $t$  is calculated as the vote counts for that party divided by county  $c$ 's VAP. The horizontal axis is the difference in the number of field offices—i.e.,  $G_{cj,t+1} - G_{cj,t}$ —and each dot corresponds to a county-party combination. We present the scatter plot and the best-fitting non-parametric polynomial curve with its 95% confidence interval. Figure 3.1 exhibits a positive relation: a candidate's vote share in a county increases with more field offices. The positive trend tails off and turns downward at the far right end; the decline is largely driven by a few outlier counties where the competition was intense and the candidates added five or more field offices. For example, in Broward County, Florida, the Obama campaign increased field offices from four in 2008 to ten in 2012; however, his vote share dropped from 36.6% to 35.7%, due to the intensity of the competition.

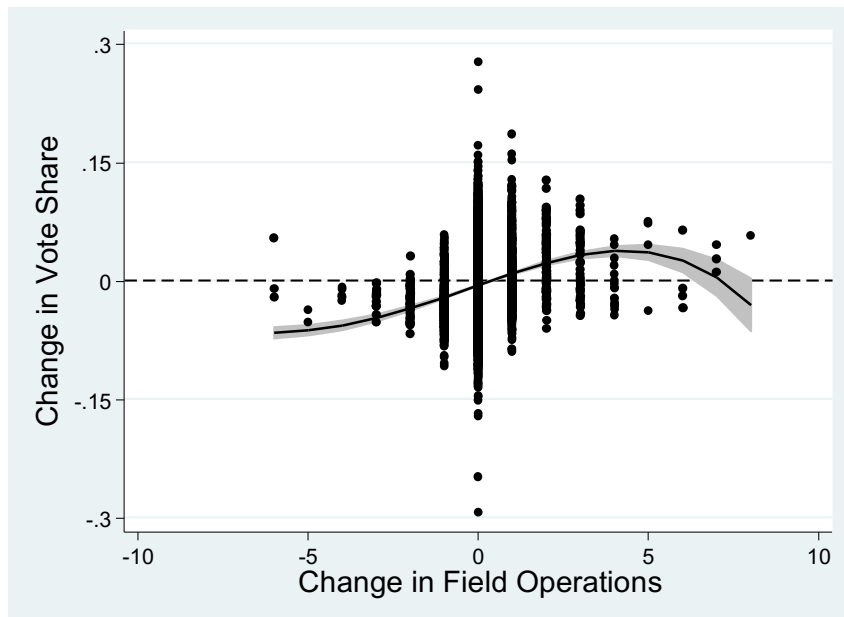
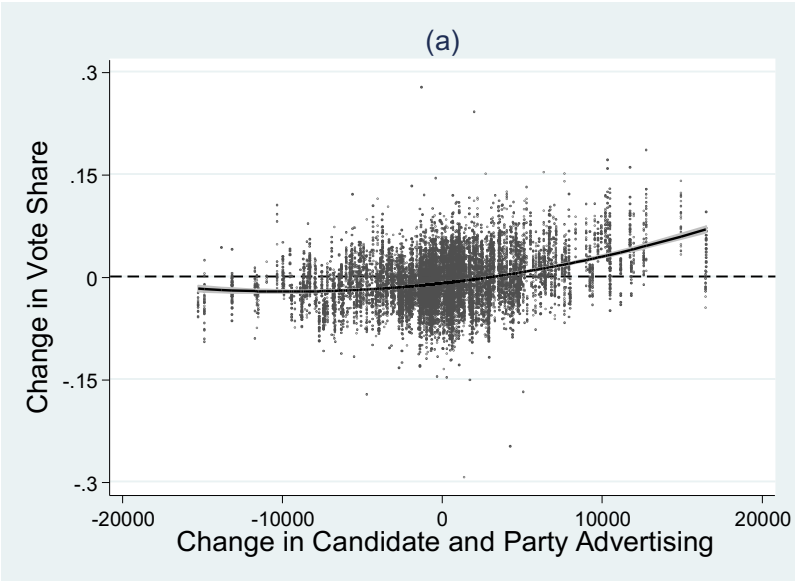


Figure 3.1: Vote Shares versus Ground Campaigning. Note: Each dot corresponds to a county-party combination. The line is the best-fitting non-parametric polynomial curve with its 95% confidence interval.

Similarly, Figure 3.2 depicts the changes in vote shares against the changes in advertising. We plot in Figure 3.2a, the ads sponsored by the candidates and their national committees, and in Figure 3.2b, those by the PACs. The horizontal axis now corresponds to the changes in advertising GRPs in each county-party combination. Once again, we observe a positive trend: a candidate's vote share goes up with an increase in advertising; this holds true for both the candidate's own advertising and the PAC advertising.

Figures 3.1 and 3.2 also show that ground campaigning and television advertising vary across elections, indicating that we have a sufficient amount of variation in our data for identification.








Figure 3.2: Vote Shares versus Television Advertising. Each dot corresponds to a county-party combination. The line is the best-fitting non-parametric polynomial curve with its 95% confidence interval.

### 2.6.2 Voter Heterogeneity

Is there heterogeneity in campaign effect? Would voters from different segments respond differently to campaigns? To gain an initial answer to these questions, we turn our attention to voter partisanship, a characteristic essential for signaling voters' political predisposition (Campbell 1992). For each county, we calculate the percentage of resident citizens who are registered as either a Democrat or a Republican. We then categorize a county as a high (low)-Democratic county if the percentage of registered Democrats there is above (below) the mean, and vice versa for the high-Republican and low-Republican counties. Figure 3.3 depicts the relation between vote shares and ground campaigning, separated into counties with a low or high percentage of partisan support, respectively. Again, for illustration, we show a scatter plot and the best-fitting non-parametric polynomial with its 95% confidence interval. The solid and dashed lines represent counties with high and low partisanship, respectively. We see that while

both lines exhibit a positive trend, the solid line has a much steeper slope, suggesting that ground campaigning seems to have a stronger effect in counties with a higher percentage of partisan voters. As Figure 3.3 only provides some initial suggestive evidence, we will in the next section specify how the effect of various campaign activities may depend on a voter's partisanship, after we control for other potential predictors of voter preference.




Figure 3.3: Effect of Ground Campaigning for Low-Partisan and High-Partisan Counties. Each dot corresponds to a county-party combination. The line is the best-fitting non-parametric polynomial curve with its 95% confidence interval.

### 3. Model of Voter Preference

We posit that individual  $i$  from county  $c$  has latent voting utility that she associates with the candidate from party  $j$  during election  $t$ , denoted as  $u_{icjt}$ . An individual faces three voting options—the Democratic candidate, the Republican candidate, and the outside option, which corresponds to voting for an independent candidate or choosing not to vote. Individual  $i$  chooses

the option that yields the highest utility, and the market shares for the three options are revealed from aggregating over individual choices. The conditional indirect utility is specified as

$$u_{icjt} = \Gamma_i(G_{cjt}, A_{cjt}) + \alpha_i + \eta X_{ct} + \xi_{mj} + \Delta \xi_{cjt} + \phi_t + \varepsilon_{icjt}. \quad (1)$$

The first component,  $\Gamma_i(G_{cjt}, A_{cjt})$ , captures how individual  $i$ 's goodwill towards candidate  $j$  is affected by how much she is exposed to the candidate's ground campaigning,  $G_{cjt}$ , and mass-media advertising,  $A_{cjt}$ . Because individuals may have diverse tastes for campaigns, we allow the effect to be heterogeneous in tastes and denote it with a subscript  $i$ . We will explain the specification for the campaign effect in Section 3.1.

The second component,  $\alpha_i$ , captures the remaining individual-specific heterogeneity in voting preference. It can be understood as the mean voting utility for  $i$  that is not explained by her exposures to campaigns. This term is further decomposed into three parts: (1) the grand mean across individuals,  $\alpha_1$ ; (2) the deviation from the mean that is attributed to observable individual characteristics,  $\alpha_2 D_{ijt}$ ; and (3) the individual departure from the mean related to all other unobservable individual characteristics,  $\sigma^\alpha v_i^\alpha$ , where we assume that  $v_i^\alpha$  is from a standard normal distribution. The unobserved characteristics include, for example, whether the individual gets a salary increase or loses her health insurance, which probably would shape her taste towards presidential campaigns but are usually missing from the data collection. We allow the three terms to enter utility linearly such that  $\alpha_i = \alpha_1 + \alpha_2 D_{ijt} + \sigma^\alpha v_i^\alpha$ .

The first and second components in Equation (1) capture the voter heterogeneity that could be attributed to observable or unobservable individual characteristics. The next four components describe the utility specific to the candidates, markets, and elections, but common to all individuals.

The term  $\eta X_{ct}$  captures how the voting utility is affected by observable county-election specific characteristics. Examples of such variables include the county's racial composition and socio-economic conditions such as the median household income and the unemployment rate, all of which may influence voter preference towards a candidate.

Next,  $u_{icjt}$  is also a function of unobservable characteristics related to a specific county-party-election combination. This could be further decomposed into three parts:  $\xi_{mj}$ ,  $\Delta\xi_{cjt}$ , and  $\phi_t$ .  $\xi_{mj}$  refers to the mean utility toward the candidate from party  $j$  across all the residents in the same media market  $m$ . People from the same media market likely exhibit similar political preferences due to exposures to the same media content (including news coverages), as well as to similar contextual conditions such as economic well-being. It is challenging to control for all the potential factors; thus, we use the fixed effect,  $\xi_{mj}$ , to absorb the cross-sectional variation among media markets and candidates.

The fifth component,  $\Delta\xi_{cjt}$ , is the county-party-election specific deviation from the mean utility,  $\xi_{mj}$ , which quantifies the hard-to-measure utility shifts over time. For example, when Hurricane Sandy hit the Northeastern part of the United States right before the Election Day in 2012, President Obama promptly committed to the relief operations and was praised for his crisis leadership, causing a positive boost in his support. Such unobserved factors would not be reflected in  $\xi_{mj}$  but would be captured by  $\Delta\xi_{cjt}$ . It is noteworthy that this county-party-election specific deviation is unobservable to the econometrician but is assumed to be observed by voters and candidates. This causes an endogeneity problem for estimating the parameters in  $\Gamma_i(G_{cjt}, A_{cjt})$ . We will discuss our solution to this problem in Section 3.3.

The sixth component of the utility is  $\phi_t$ , which captures the election-specific shocks to voting utility common to all county-party combinations. Finally,  $\varepsilon_{icjt}$  is the idiosyncratic utility shock that is assumed to be independently and identically distributed (i.i.d.) Type I extreme value across individuals, counties, candidates, and elections.

### 3.1 Specification of Campaign Effect

We postulate that the campaign effect,  $\Gamma_i(G_{cjt}, A_{cjt})$ , is a function of candidate's ground campaigning and mass-media advertising. As previously discussed, ground campaigning takes the form of field operations,  $G_{cjt}$ , and advertising has two primary types: own ads made by the candidates and their parties,  $A_{cjt}^o$ , and outside ads sponsored by the PACs,  $A_{cjt}^p$ . Both types of ads enter the model in log form to capture the diminishing return for advertising<sup>37</sup>. We allow those campaign activities to have a heterogeneous effect across individuals. To sum up, we specify the campaign effect in the following linear form:

$$\Gamma_i(G_{cjt}, A_{cjt}) = \beta_i G_{cjt} + \gamma_i A_{cjt}^o + \pi_i A_{cjt}^p . \quad (2)$$

The parameter,  $\beta_i$ , captures the voter  $i$ 's taste towards field operations and consists of three components: (1) the mean taste across individuals,  $\beta_1$ ; (2) the deviation from the mean that could be attributed to observable individual characteristics,  $\beta_2 D_{ijt}$ ; and (3) the individual departure from the mean related to all unobservable individual characteristics,  $\sigma^\beta \nu_i^\beta$ . Similarly, we decompose  $\gamma_i$  and  $\pi_i$  into three components such that

---

<sup>37</sup> We tested field operations in the log form and the quadratic form to examine a potential diminishing return for having more field offices. The linear form has the highest exploratory power to explain vote shares. This may be partially because the variable does not have enough variation to detect a non-linear effect: among counties with at least one field office, less than 5% had more than 4 offices.

$$\begin{aligned}
\beta_i &= \beta_1 + \beta_2 D_{ijt} + \sigma^\beta v_i^\beta \\
\gamma_i &= \gamma_1 + \gamma_2 D_{ijt} + \sigma^\gamma v_i^\gamma, \quad (3) \\
\pi_i &= \pi_1 + \pi_2 D_{ijt} + \sigma^\pi v_i^\pi
\end{aligned}$$

where each unobserved characteristic  $v_i$  is assumed to come from a standard normal distribution.

The individual characteristic  $D_{ijt}$  that we examine here is voters' party affiliation, which is believed to be an important factor affecting political preference towards candidates. A voter may be affiliated with either the Democrats or the Republicans, or neither. Because we observe the aggregate data of party affiliation on the county level, we assume the partisan variable to follow a multinomial distribution of three categories (i.e., Democrats, Republicans, neither), where the empirical means of the categories correspond to the observed percentages of registered partisan voters for each county. For example, if a county had 30% registered Democrats and 35% Republicans, our simulated individual partisanship would have roughly 30% being labeled as the Democrats, 35% as the Republicans, and the remaining 35% as neither.

### 3.2 Distributional Assumptions and Implied Market Shares

From Equations (1), (2), and (3), the utility function can be rewritten as

$$\begin{aligned}
u_{icjt} &= \alpha_1 + \alpha_2 D_{ijt} + \sigma^\alpha v_i^\alpha + (\beta_1 + \beta_2 D_{ijt} + \sigma^\beta v_i^\beta) G_{cjt} \\
&+ (\gamma_1 + \gamma_2 D_{ijt} + \sigma^\gamma v_i^\gamma) A_{cjt}^o + (\pi_1 + \pi_2 D_{ijt} + \sigma^\pi v_i^\pi) A_{cjt}^p \quad . \quad (4) \\
&+ \eta X_{ct} + \xi_{mj} + \Delta \xi_{cjt} + \phi_t + \varepsilon_{icjt}
\end{aligned}$$

We then rewrite Equation (4) as

$$\begin{aligned}
u_{icjt} &= \delta(G_{cjt}, A_{cjt}^o, A_{cjt}^p, X_{ct}; \theta_1) \\
&+ \mu(G_{cjt}, A_{cjt}^o, A_{cjt}^p, D_{ijt}, v_i; \theta_2) + \varepsilon_{icjt}, \quad (5)
\end{aligned}$$

where  $\theta_1 = (\alpha_1, \beta_1, \gamma_1, \pi_1, \eta, \xi_{mj}, \Delta \xi_{cjt}, \phi_t)$  and  $\theta_2 = (\alpha_2, \beta_2, \gamma_2, \pi_2, \sigma^\alpha, \sigma^\beta, \sigma^\gamma, \sigma^\pi)$ .

Hence, the utility is expressed in two parts: the mean utility across individuals,

$$\delta_{cjt} = \alpha_1 + \beta_1 G_{cjt} + \gamma_1 A_{cjt}^o + \pi_1 A_{cjt}^p + \eta X_{ct} + \xi_{mj} + \Delta \xi_{cjt} + \phi_t, \text{ and the individual departure from the mean, } \mu_{icjt} = (\alpha_2 D_{ijt} + \sigma^\alpha v_i^\alpha) + (\beta_2 D_{ijt} + \sigma^\beta v_i^\beta) G_{cjt} + (\gamma_2 D_{ijt} + \sigma^\gamma v_i^\gamma) A_{cjt}^o + (\pi_2 D_{ijt} + \sigma^\pi v_i^\pi) A_{cjt}^p.$$

We assume that  $\varepsilon_{icjt}$  follows an i.i.d. type I extreme value distribution, and normalize the utility for the outside option to  $u_{ic0t} = 0 + \varepsilon_{ic0t}$ . Based on the distributional assumption of the idiosyncratic shocks and the utility specification stated above, we define the probability of voter  $i$  in county  $c$  voting for the candidate from party  $j$  during election  $t$  as

$$s_{icjt} = \frac{\exp\left(\delta\left(G_{cjt}, A_{cjt}^o, A_{cjt}^p, X_{ct}; \theta_1\right) + \mu\left(G_{cjt}, A_{cjt}^o, A_{cjt}^p, D_{ijt}, v_i; \theta_2\right)\right)}{1 + \sum_{k=1}^2 \exp\left(\delta\left(G_{ckt}, A_{ckt}^o, A_{ckt}^p, X_{ct}; \theta_1\right) + \mu\left(G_{ckt}, A_{ckt}^o, A_{ckt}^p, D_{ikt}, v_i; \theta_2\right)\right)}. \quad (6)$$

We can obtain the county-level vote share by integrating over individuals such that

$s_{cjt} = \int s_{icjt} dP(D)dP(v)$ , where  $P(D)$  and  $P(v)$  are the distributions for the individual observable,  $D_{ijt}$ , and the idiosyncratic disturbances,  $v_i$ , respectively. Again,  $D_{ijt}$  is the partisan indicator, which we simulated, county by county, from an empirical multinomial distribution  $\hat{P}(D)$ , with the category means being the observed percentages of registered partisans for each party in that county.

### 3.3 Identification and Estimation

Per our model specification, we examine voter's choice of presidential candidates and allow individual heterogeneity in campaign effects. The challenge here is that the choices are observed at the aggregated county level. To address this, we employ the estimation approach developed by Berry, Levinsohn, and Pakes (1995), typically referred to as "BLP", which has been used in various marketing applications (e.g., Chung 2013; Gordon and Hartmann 2013; Sudhir 2001). The parameters are estimated via the method of moments (GMM) to minimize the

GMM objective function such that:  $\hat{\theta} = \text{argmin } g(\theta)' \cdot W \cdot g(\theta)$ , where  $g(\theta) = Z' \cdot \Delta \xi_{cjt}$  is the moment condition,  $Z$  is the vector of instruments assumed orthogonal to  $\Delta \xi_{cjt}$ , and  $W$  is the weight matrix (Hansen 1982).

Typically, the vector of the right-hand-side observables in Equation (1) can be used to form  $Z$ . However, we are concerned with an endogeneity problem. The county-party-election specific deviation from the mean utility,  $\Delta \xi_{cjt}$ , is observable to the candidates and PACs, and hence likely plays a role in determining the level of each campaign activity,  $G_{cjt}$ ,  $A_{cjt}^o$ , and  $A_{cjt}^p$ , causing a correlation between the error term and the campaign variables. For example, negative shocks (such as negative word-of-mouth, slow economic growth, and certain demographic shifts) of  $\Delta \xi_{cjt}$  may decrease voter preference towards a candidate, who is rightfully incentivized to increase the campaign intensity. Vice versa, in relatively safe counties where he or she sees sufficient voter support, a candidate may want to retain just the minimum level of campaigning and allocate the precious resources to where the competition is more intense. Shachar (2009) provides empirical evidence that candidates do more campaigning when the competition is more intense. Without accounting for this endogenous behavior, we may underestimate the true campaign effect.

A common approach to address endogeneity is to choose instruments that are correlated with the campaign activities but exogenous to  $\Delta \xi_{cjt}$ . The instruments we choose for advertising are the third-quarter DMA-level ad prices in the year before each election. The argument for the validity of those instruments is that price changes affect advertising cost and hence shift the amount of advertising, but the cause of the price fluctuation is assumed to be outside the system, i.e., independent of  $\Delta \xi_{cjt}$ . We use ad prices from the previous year instead of from the election



year to reduce the possibility that price changes are due to the changing demand of advertising in an ad-filled election year.

Our ad-price data come from the Kantar Media SRDS TV and Cable Source, and we collected prices for three dayparts: prime access, prime, and late news. Although invariant across candidate own ads and PAC ads, the ad costs are able to instrument both types of ads through the difference in airtime for each type. Using data from the University of Wisconsin Advertising Project (Goldstein and Rivlin 2008) we found that candidate's own ads were aired more during the prime access and the prime dayparts than PAC ads, while the latter more frequently appeared during the late news daypart. Therefore, the costs for different dayparts have varying effects on the two types of ads, providing the variation needed for identification. In particular, the unique exogenous variation in the ad price for prime access time and prime time helps identify candidate's own ads and the unique variation in the costs for late news daypart helps identify the effect for PAC ads. However, the instruments are constant across parties and would not provide any between-party variation, so that the first-stage fitted values for advertising (candidate's own ads and PAC ads) conditional on the instruments and the other covariates would be the same for the Democrats and Republicans. Hence, we include the interactions between the Democratic indicator and each of the ad cost instruments, which add between-party variation in the first-stage fitted values for the endogenous ad variables to help identification.

We use the real estate rental price in each county the year before the election to instrument field operations. The interaction with the Democratic indicator is also included to provide the between-party variation. The identification argument is similar to that of using lagged ad prices to instrument advertising. Specifically, lagged rental prices affect the demand for office rental and, hence, should be correlated with the number of field offices, but not directly correlated with

the unobservable utility shocks. Rental prices may not be valid instruments if some unobservable local economic conditions, say, an expected business boost, caused both an increase in the previous year's rent and a change in residents' candidate preference. By including a rich set of socio-economic variables for each county, we believe we have reasonably offset this potential bias because  $\Delta \xi_{cjt}$  now captures utility shocks not explained by the socio-economic shifts.

In addition to ad prices and rental costs, we included the interactions between rental price and each of the three ad prices. The rationale is to increase the first-stage predictive power hence increase the estimation precision in the final model (Angrist and Pischke 2008). Our final vector of the GMM instruments contains the lagged real estate price, the lagged ad prices, the interactions between the Democrat indicator and the cost shifters, the interactions among the cost shifters, and all of the exogenous variables in Equation (1) including the fixed effects. Because partisan information is available for slightly more than half of the states, we form separate moments for states with and without this variable, so that only the states with the partisan information contribute to the estimation of the random coefficients. Heuristically, the variation in vote shares for counties with different partisan density but the same campaign activities helps identify the mean of the random campaign effect distribution. For example, if two counties both have one more Democratic field office from one election to another, and if the one with a higher percentage of Democrats also sees a bigger change in vote share for the Democratic candidate, the partisan variable would be identified to positively moderate the effect for field operations. The same logic applies to how partisans and non-partisans respond to the ad effect.

## 4 Results

### 4.1 Parameter estimates

We estimate four specifications and present the results in Table 3.5. The first two specifications estimate the effect of ground campaigning and advertising in an ordinary least squares (OLS) regression with and without the DMA-party fixed effects, respectively. The third specification incorporates the instruments, and the fourth allows heterogeneous campaign effects across individual voters, which is our full model.

We begin with a brief discussion of the OLS estimates (see columns 1 and 2). First, adding fixed effects increased the model R-squared from 0.39 to 0.66; therefore, in the subsequent analyses, we always include the fixed effects. The OLS estimates in column 2 provide benchmark values of the campaign effect: without accounting for endogeneity, field operations and advertising—both the candidate and party ads and PAC ads—are positively correlated with vote shares. We also observe a positive and significant association between digital campaigning and vote outcomes.

Table 3.5: Parameter Estimates

	(1)	(2)	(3)	(4)		
	Est (SE)	Est (SE)	Est (SE)	Est (SE)	Partisan	Sigma
Field operations	0.074*** (0.011)	0.038*** (0.006)	0.368*** (0.053)	0.361*** (0.077)	0.119*** (0.042)	0.106 (0.232)
Candidate own ads	0.015*** (0.001)	0.006*** (0.001)	0.063*** (0.009)	0.121*** (0.011)	-0.173*** (0.005)	0.108*** (0.008)
PAC ads	0.022*** (0.001)	0.008*** (0.001)	0.012 (0.021)	0.013 (0.028)	0.167*** (0.004)	0.000 (0.1806)
Digital campaigning	0.064*** (0.018)	0.147*** (0.014)	0.107* (0.063)	0.032 (0.083)		
Year 2008	-0.111*** (0.035)	-0.276*** (0.028)	-0.388** (0.164)	-0.299 (0.212)		
Year 2012	-0.324*** (0.060)	-0.584*** (0.045)	-0.470*** (0.167)	-0.296 (0.222)		

Incumbent status	0.009 (0.010)	-0.044*** (0.008)	-0.067 (0.068)	-0.002 (0.087)		
Home state advantage for Presidential candidates	0.025 (0.019)	0.057*** (0.019)	0.061** (0.024)	-0.032 (0.031)		
Home state advantage for Vice-Presidential candidates	-0.113*** (0.029)	-0.078*** (0.026)	-0.079** (0.033)	-0.114*** (0.040)		
Governor advantage	0.091*** (0.007)	0.011* (0.006)	-0.001 (0.010)	-0.014 (0.012)		
Percentage AAs	-0.181*** (0.043)	-0.982*** (0.060)	-1.049*** (0.055)	-1.008*** (0.079)		
Percentage AAs X Democrat	1.242*** (0.063)	2.655*** (0.082)	2.401*** (0.087)	2.451*** (0.129)		
Median household income	-0.050 (0.040)	-0.182*** (0.051)	-0.313*** (0.049)	-0.339*** (0.072)		
Median household income X Democrat	-0.671*** (0.025)	0.108 (0.073)	-0.018 (0.065)	0.001 (0.098)		
Unemployment rate	-0.020*** (0.002)	-0.014*** (0.002)	-0.014*** (0.002)	-0.013*** (0.003)		
Unemployment rate X Democrat	0.030*** (0.003)	0.025*** (0.003)	0.026*** (0.003)	0.025*** (0.004)		
Gini index	0.028 (0.173)	0.319** (0.158)	-0.111 (0.165)	-0.015 (0.226)		
Gini index X Democrat	0.352 (0.252)	0.931 (0.222)	0.273 (0.237)	0.242 (0.332)		
Median house value	-0.190*** (0.017)	-0.084 (0.024)	-0.077*** (0.021)	-0.036 (0.031)		
Median house value X Democrat	0.542*** (0.023)	0.374 (0.035)	0.343*** (0.031)	0.326*** (0.047)		
Dropout rate	-0.357*** (0.089)	-0.586 (0.082)	-0.530*** (0.079)	-0.625*** (0.110)		
Dropout rate X Democrat	-0.778*** (0.144)	-0.027 (0.125)	0.008 (0.112)	-0.040 (0.158)		
Poverty rate	-0.031*** (0.002)	-0.035 (0.002)	-0.036*** (0.001)	-0.040*** (0.002)		
Poverty rate X Democrat	-0.002 (0.002)	0.024 (0.002)	0.022*** (0.002)	0.022*** (0.003)		
Intercept	2.112*** (0.438)	1.164** (0.478)	3.206*** (0.826)	3.978*** (1.186)	0.558*** (0.031)	0.575*** (0.089)
DMA-Party Fixed Effects	No	Yes	Yes	Yes		

Instruments	No	No	Yes	Yes
N	18,650	18,650	18,650	18,650
R <sup>2</sup>	0.39	0.66		

\*\*\* p<0.01; \*\* p<0.05; \* p<0.10

Note: We report results from four specifications. Column (1) estimates the marginal effects of ground campaigning and television advertising in OLS without fixed effects and column (2) with fixed effects. Column (3) estimates the marginal effects with instruments. Column (4) further incorporates voter heterogeneity in campaign effects.

Before discussing IV estimates in column 3, we first summarize some diagnostic statistics for the instruments. We estimated the first-stage regression equations and the reduced-form regression equation as outlined in Angrist and Pischke (2008) and present the results in Table 3.6. The first-stage regression results indicate a clear effect of the instruments on the three endogenous campaign variables—field operations, candidate’s own ads, and PAC ads. The partial F statistics are 26.15, 36.88, and 11.58, respectively. The instruments also have sufficient power to explain the vote shares after controlling for all the covariates in the reduced-form regression model. Those results provide initial evidence for the robustness of our instruments.

Table 3.6: Diagnostic Results for Instruments

Endogenous variable:	First-stage regression			Reduced-form regression
	Field operations	Candidate own ads	PAC ads	Vote shares (log)
Rent	-0.219* (0.104)	-1.934*** (0.426)	-0.367 (0.384)	-0.382*** (0.067)
Ad price I	0.569 (0.728)	9.101** (2.990)	-6.677* (2.692)	1.067* (0.472)
Ad price II	-2.672*** (0.544)	-23.325*** (2.234)	-5.881** (2.011)	-1.929*** (0.353)
Ad price III	1.844** (0.633)	14.006*** (2.597)	12.750*** (2.338)	0.788 (0.410)

Rent X Democrat	0.324*** (0.057)	0.214 (0.235)	0.617** (0.212)	0.477*** (0.037)
Democrat X Ad price I	-0.008 (0.083)	-1.791*** (0.341)	-0.047 (0.307)	-0.022 (0.054)
Democrat X Ad price II	0.101 (0.074)	0.143 (0.304)	-0.286 (0.274)	0.039 (0.048)
Democrat X Ad price III	-0.087 (0.074)	1.167*** (0.304)	0.697* (0.274)	0.002 (0.048)
Rent X Ad price I	-0.098 (0.107)	-0.911* (0.438)	0.861* (0.394)	-0.135 (0.069)
Rent X Ad price II	0.411*** (0.079)	3.182*** (0.324)	1.026*** (0.291)	0.276*** (0.051)
Rent X Ad price III	-0.264** (0.093)	-2.301*** (0.381)	-1.924*** (0.343)	-0.128* (0.060)
Control variables	Yes	Yes	Yes	Yes
R <sup>2</sup>	0.28	0.62	0.70	0.67
Partial F	26.15	36.88	11.58	22.35
Partial R <sup>2</sup>	0.01	0.02	0.01	0.01

Note: Control variables include all the exogenous variables as reported in Table 3.5.

After applying the instruments, we find the effect for field operations to be bigger than the corresponding OLS estimate; the direction of the change is expected with the presence of endogeneity. When candidates deploy more field offices in more intense competition, the OLS estimate would be attenuated towards zero, as is the case here. The IV estimate for candidate's own ads is also larger than the OLS estimate and is significantly positive. The IV estimate for PAC ads, although positive and larger than the OLS estimate, is no longer significant at 0.05 level.

Our final model (column 4 in Table 3.5) incorporates voter heterogeneity; in particular, we examine how the effect of various campaign activities depends on voter partisanship. The first column under specification (4) lists the parameter estimates for non-partisan voters; the second column is the estimated interaction effect with voter partisanship; and the third column corresponds to  $\hat{\sigma}$ , the estimated unobserved heterogeneity in each campaign effect.

We discover some interesting patterns regarding the effect of campaign activities in different voter segments. First, we estimate the effect of field operations to be 0.361, positive and significant for non-partisans and the effect is even stronger for partisan voters. In contrast, candidate's own advertising is found to have a positive and significant effect for non-partisans (0.121,  $p < 0.01$ ) but the effect reduces to being indistinguishable from zero for partisan voters (-0.052,  $p > 0.10$ ). That is, candidate's own ads are only effective for voters on the margin, i.e., those who have not yet developed a partisan affiliation with either party. Interestingly, PAC ads are found to be effective only among partisans ( $0.180 = 0.013 + 0.167$ ,  $p < 0.01$ ) and the effect is null among non-partisans (0.013,  $p > 0.10$ ). After controlling for partisanship, we find that the remaining variation in the effect of field operations and PAC ads is no longer significant across individuals. Candidate's own ads, on the other hand, still have heterogeneous effect among voters, suggesting that there may be additional voter segments along other dimensions of individual characteristics. Finally, voters' partisanship is found to have a strong main effect (0.558,  $p < 0.01$ ): not surprisingly, independent from the campaign effects, partisans tend to favor the candidates from their party. Nevertheless, our estimates indicate that campaign activities still influence voting outcomes beyond voters' baseline preference.

Among the control variables, we find some initial evidence that digital campaigning is positively associated with the candidate's vote shares; the IV estimate is 0.107 ( $p < 0.10$ ), and the final estimate is positive although insignificant at 0.10 level after incorporating individual heterogeneity. We also find that counties with fewer African American residents, lower median household income, lower unemployment rate, lower high school dropout rate, and lower poverty rate tend to have higher vote shares for the Republican candidates; in contrast, counties with

more African Americans, higher high school dropout rate, higher median household values, and higher poverty rate tend to favor the Democratic candidates.

## 4.2 Elasticity Estimates

In this section, we present the elasticities of various campaign activities. The field-operation elasticity, derived from the utility specification, is

$$\zeta_{jk,ct}^{\beta} = \frac{\partial s_{j,ct} G_{k,ct}}{\partial G_{k,ct} s_{j,ct}} = \begin{cases} \frac{G_{j,ct}}{s_{j,ct}} \int \beta_i s_{ij,ct} (1 - s_{ij,ct}) dP(D) dP(v) & \text{if } j = k \\ -\frac{G_{k,ct}}{s_{j,ct}} \int \beta_i s_{ij,ct} s_{ik,ct} dP(D) dP(v) & \text{if } j \neq k \end{cases},$$

which depends on the individual-specific taste parameter for field operations,  $\beta_i$ , integrated over individual voters. Elasticities for advertising are defined similarly.

Table 3.7 presents the elasticity estimates based on the estimates of specification (4) in Table 3.5. The numbers in the diagonal refer to the percentage change in vote share in response to a 1% increase in the party's own campaign efforts; and those in the off-diagonal correspond to the change in a party's vote share resulting from a 1% increase in the rival's campaign.

We begin with the elasticity estimates for advertising, as they are more straightforward to interpret. The elasticity for candidate's own ads is estimated to be 0.059 and 0.081 for the Republicans and Democrats, respectively: a 1% increase in the candidate's own advertising would result in a roughly 0.059% increase in vote shares for the Republicans and 0.081% for the Democrats. It is nontrivial to compare our estimates to others, as not many studies have carefully addressed both campaign endogeneity and voter heterogeneity as we do here. The paper closest to ours is Gordon and Hartmann (2013), which also uses ad costs to instrument advertising. Our elasticities roughly double their estimates of 0.033% for the Republicans and 0.036% for the Democrats. The difference may be because Gordon and Hartmann (2013) categorize ads by the



target candidate regardless of the sponsors and hence their ad effect is a composite effect of candidate's own ads, PAC ads, and other "hybrid/coordinated" ads. Our estimated elasticity for candidate's own ads should be greater than theirs when candidate ads have a stronger effect than other types of ads, as is what we find here. Furthermore, the fact that Gordon and Hartmann (2013) estimated ad GRPs while we directly observed the ad variables may also explain some discrepancy between our estimates. Another recent empirical paper on presidential advertising, Lovett and Peress (2015), estimated the ad effect size to be a 3.0% increase in vote share if the party can increase individual ad exposures by one standard deviation. It is hard to directly compare their numbers to ours, because of the use of different ad metrics. Nevertheless, their finding that ads are more effective among swing voters is qualitatively consistent with what we find.

The cross-ad elasticity estimates for candidate's own ads are noticeably smaller than the own-ad elasticities. If the Democrats raise their campaign ads by 1%, the Republican's vote share would decrease by 0.033%; the decrease is estimated to be 0.051% for the Democrats if the Republicans increase their campaign ads by 1%.

The average effect for PAC ads is found to be smaller than that for the candidate's own ads. The own-elasticity estimates are 0.032% for Republicans and 0.045% for the Democrats, and the cross-elasticity estimates are -0.011% and -0.020%, respectively.

For field operations, we calculated the percentage change in vote shares in response to one additional field office<sup>38</sup>. We find that the elasticity is much higher for the Democrats than for the

---

<sup>38</sup> The addition of one field office can be understood as a proxy for the average amount of voter contacts associated with a typical field office. The fact that the number of field offices is highly correlated with the self-reported voter contacts suggests that there may be a somewhat narrow distribution for the amount of voter contacts behind each field office. We thank an anonymous reviewer for raising this question.

Republicans: with one additional field office, the vote share is estimated to rise by 3.305% for the Democrats (versus 1.143% for the Republicans), suggesting that the Democratic field offices are more effective in driving vote shares than the Republican's. Similarly, the cross elasticity estimates reveal that the Democratic field offices are also more effective in converting Republican-leaning voters to the Democratic candidates, rather than the other way around (-1.889% versus -0.529%).

It is challenging to benchmark our field-operation elasticity estimates as empirical studies on this topic are scarce. In one exception, Darr and Levendusky (2014) identified a 1.04% boost in county-level vote shares with the presence of a Democratic field office, which corresponds to one third of our elasticity estimate for the Democrats. It is worth noting that Darr and Levendusky (2014) used an OLS model without adjusting for the potential correlation between field office deployment and the unobserved voter preference. As we have discussed, when field operations are condensed in competitive counties, ignoring this endogeneity concern may lead to underestimating the true effect for field offices. Our estimate is directionally consistent with what one would expect when treating field operations as endogenous. Indeed, the field-office elasticity is estimated to be 0.95% based on our OLS estimates, very similar to the estimate in Darr and Levendusky (2014).

Table 3.7: Elasticity Estimates for Ground Campaigning and Advertising

		1% increase from	
	Focal party	Republican	Democrat
Candidate own ads	Republican	0.059	-0.033
	Democrat	-0.051	0.081
PAC ads	Republican	0.032	-0.011
	Democrat	-0.020	0.045
		One additional office from	
	Focal party	Republican	Democrat
Field operations	Republican	1.143	-1.889

Democrat                      -0.529                      3.305

---

Note: The elasticities are computed based on estimates from our full model. The diagonal estimates are the own elasticities and the off-diagonal elements are the cross elasticities.

### 4.3 Counterfactual Analysis

With the structural parameter estimates we are now ready to answer the “what if” questions: what the election results would have been had the candidates campaigned differently. These counterfactual questions are crucial for understanding the true causal effect of campaign activities as causal effect is defined as the difference between factual and counterfactual inferences. For example, to pin down the extent to which each campaign activity matters to an election, we could eliminate that particular campaign activity while keeping others intact, predict the winner for each state, and then compare the results to the true state winners. We report the various counterfactual results in Table 3.8.

First of all, our results highlight the importance of field operations for the Democrats. Had neither party set up any field offices, the Republicans would have won the 2008 and 2012 elections. In other words, the Democratic field operations were responsible for a large portion of their total popular votes in 2008 and 2012. Without field operations, the Democrats would have lost seventeen states (Colorado, Florida, Indiana, Iowa, Maine, Michigan, Minnesota, Nevada, New Hampshire, New Mexico, North Carolina, Ohio, Oregon, Pennsylvania, Virginia, Washington, and Wisconsin) in 2008 and fifteen states (California, Colorado, Florida, Illinois, Iowa, Michigan, Minnesota, Nevada, New Hampshire, New Mexico, Ohio, Pennsylvania, Virginia, Washington, and Wisconsin) in 2012. After all, there is truth to the popular claim that Obama owed the victories to his unprecedented field operations.

As far as the candidate's own ads are concerned, zero advertising would have changed the results for some states; for example, the Republicans would have won Indiana in 2008 and four states (Connecticut, New Hampshire, Virginia, and Washington) in 2012. However, the national results would have remained the same for all three elections. This suggests that television advertising may not be a deterministic factor for driving the election results. The finding is somewhat expected, considering that the ad elasticity estimates are with a similar magnitude for the Democrats and Republicans and that the two parties had somewhat comparable levels of television ads. These counterfactual results also seem to suggest that the effect of field operations is more substantial than the ad effect, consistent with the finding in Carroll et al. (1985). Although in a different setting of Navy enlistment, they found that the elasticity of field salesforce was large and significant (i.e., 0.44%) while advertising was not significant.

How about the ads sponsored by PACs? Not surprisingly, we find that eliminating the outside ads barely moved the needle on the election results in 2004 and 2008, perhaps explained by the modest amount of PAC ads in those elections. However, if the 2012 election had allowed zero outside ads (without changing the actual amount of candidate advertising), the Democrats would have won with a much larger margin. The finding that the Democrats benefited more from banning outside ads could provide interesting insights into the consequences of the “People's Pledge,” pioneered in the 2012 Massachusetts Senate race. According to the pledge, the Republican candidate, Scott Brown, and the Democratic candidate, Elizabeth Warren, agreed not to accept any outside ads, aiming to curb the influence of third parties. Warren defeated Brown; thus, there has been a lot of speculation regarding whether the pledge had helped the Democrats more than the Republicans. Our counterfactual analysis suggests that banning PAC ads in presidential elections is more beneficial to the Democrats. To the extent that our finding can be

extended to a Senate race, one may conjecture that part of Warren's success is attributed to eliminating the outside ads.

Currently, PACs are prohibited from directly coordinating their advertising efforts with candidates. We also conducted a counterfactual analysis to understand the effect of this policy (row 5 in Table 3.8). Had PACs been allowed to donate their ad spots to the candidates— in other words, the GRPs of the candidates' own ads would have become the sum of the GRPs from the candidate campaign and from the leaning PACs— the Republican candidates would have won significantly more states, changing the election results for 2008 and 2012. This is primarily because of the large amount of PAC ads that the Republicans received in the two recent elections. One caveat of this counterfactual analysis is that PAC ads typically are broadcasted during less popular dayparts. Even if PACs gave all their ad spots to the candidates, in reality, their ads may not be as effective as the candidate's own ads, which more frequently aired during better dayparts. Hence, the consequences of eliminating PAC ads perhaps would be bounded by the two counterfactual scenarios that we conducted: simply removing the PAC ads and transferring all the PAC GRPs to the candidates.

It is noteworthy that our counterfactuals are not based on full equilibrium outcomes, in the sense that, when one activity is removed from the campaign mix, we retain the level of the other activities. Those partial equilibrium results are under the assumption that candidates do not adjust the amount of other campaign activities with the absence of the focal activity. A full equilibrium counterfactual analysis would require us to have a supply-side model that solves the new equilibrium level for all the remaining activities given a fixed campaign budget, which is

beyond the scope of the current study. Nevertheless, the partial equilibrium analysis still sheds light on the respective effect of each campaign activity while controlling for others.<sup>39</sup>

Table 3.8: Predicted Total Electoral Votes for Counterfactual Analyses

	2004		2008		2012	
	Democratic	Republican	Democratic	Republican	Democratic	Republican
actual	252	<b>286</b>	<b>365</b>	173	<b>332</b>	206
zero field operations	193	<b>345</b>	168	<b>370</b>	99	<b>439</b>
zero candidate own ads	243	<b>295</b>	<b>353</b>	185	<b>296</b>	242
zero PAC ads	249	<b>289</b>	<b>351</b>	187	<b>430</b>	108
PAC ads rolled over to candidate own ads	210	<b>328</b>	245	<b>293</b>	120	<b>418</b>

Note: For zero field operations, we assigned zero field office to each party candidate without changing other campaigning activities. Similar steps were taken for the other counterfactual conditions except for the last one, where we assumed that PAC GRPs became the candidate's. The reported numbers are the total final electoral votes won by each party candidate. The predicted winner of each election is in bold. For Alaska we used the actual results for each election.

If there were a chance to relive the history, what would it take for the losing parties to change their fate? In particular, could the Republicans have won the 2008 and 2012 elections if they had enhanced their field operations in swing states, as the public seemed to suggest?<sup>40</sup>

<sup>39</sup> We thank an anonymous reviewer for pointing this out.

<sup>40</sup> For example, days after President Obama was first elected in 2008, the *Washington Times* published an article claiming that “one of the keys to Mr. Obama’s success was building an unprecedented ground game.” Four years later, the *New York times* (April 17) ran an article saying that the extent to which “Mr. Romney can match the Obama’s footprint [in the ground game] in the swing states may prove critical in deciding the election.”

To answer this question, we conducted a counterfactual analysis to calculate the fewest *additional* field offices needed for the Republicans to win the national election. The computation took two steps. First, for each swing state that the Republicans lost, we used the model estimates to solve for the fewest additional field offices needed for the Republicans to win more popular votes in the state, holding constant the number and locations of the Democratic field offices and the other campaign activities.<sup>41</sup> The optimal numbers of field offices are presented in Table 3.9. In the second step, assuming that the cost of setting up a field office is constant across states, we selected the optimal combination of swing states that required the fewest additional field offices to reach the 270 electoral votes. The optimal combination takes into account the number of additional field offices per state and the electoral votes that each state carries and, hence, represents the most cost-effective way to allocate field operations in order to win the election.

To reach the 270 goal that year, the McCain campaign would have had to set up at least fourteen additional offices: two in Florida, one in Indiana, one in Nevada, one in North Carolina, three in Ohio, and six in Pennsylvania, conditional that the candidates locate the offices where they are expected to be the most effective, i.e., among partisans. Winning these states would have brought in a total of 99 electoral votes. Four years later, it would have taken fewer additional field offices for the Republicans to win the election, given that the Romney campaign

---

<sup>41</sup> The procedure, say for Colorado in 2008, goes as follows. First, we retain the value of the Democratic field operations in Colorado. Second, we solve for the number of Republican field offices per county, which minimizes the total number of field offices given that the Republicans would win the majority of votes. For the 53 counties in which the Republicans did not have an office in 2008, we restrict the value to fall between 0 and 10, as less than 0.5% of the counties in our sample ever had more than ten field offices. For the 11 counties with at least one office that year, we bound the variable between the current value and 10.

had already invested more in field operations: Romney could have won by adding merely six more field offices: one in Florida, one in Nevada, two in Ohio, and two in Virginia.

Table 3.9: Predicted Optional Field Offices

2008			2012		
State	Existing	Optimal	State	Existing	Optimal
Colorado (9)	11	15	Colorado (9)	14	17
Florida* (27)	0	2	Florida* (29)	48	49
Indiana* (11)	0	1	Michigan (16)	23	26
Iowa (7)	16	20	Minnesota (10)	0	2
Michigan (17)	14	19	Nevada* (6)	12	13
Minnesota (10)	13	16	New Hampshire (4)	9	10
Nevada* (5)	12	13	Ohio* (18)	38	40
New Hampshire (4)	4	6	Oregon (7)	0	4
New Jersey (15)	1	9	Pennsylvania (20)	24	27
New Mexico (5)	10	13	Virginia* (13)	28	30
North Carolina* (15)	18	19	Wisconsin (10)	24	27
Ohio* (20)	9	12			
Oregon (7)	0	5			
Pennsylvania* (21)	17	23			
Virginia (13)	18	21			
Wisconsin (10)	9	14			

Note: We report the optimal field operations that could have helped the Republicans win each battleground state. States with an asterisk make up the optimal state combination that requires the fewest field offices to win 270 electoral votes. We list the electoral votes in parentheses.

We then examine whether advertising could have helped the losing party in each election, and if so, by how much. Our analysis shows that advertising could have played a critical role in deciding the election, but only in a close competition such as the 2004 one. If his campaign had increased the ad coverage by 50% in New Mexico (worth of \$1.0 million spending) and 4.3



times in Virginia (worth of \$0.6 million spending), Kerry could have won additional 18 ECs, enough to help him reach the 270 goal and claim the victory! The Democrats-leaning PACs could have also helped Kerry win three more states by spending an additional \$0.03 million (the equivalent of 50% more GRPs) in New Mexico, \$0.11 million (50% more GRPs) in Iowa, and \$0.37 million (1.7 times more GRPs) in Nevada. Had that happened, the 2004 election would have resulted in a 269 to 269 tie. As directed by the 12th Amendment, members of the House of Representatives would have had to choose the president that year. To break the tie, the Democrats-leaning PACs would need to spend another \$1.1 million to win Colorado, which seems feasible given that year's total PAC ads budget of \$7.0 million in favor of the Democrats.

However, when the winner has a big competitive advantage, it is unlikely for a losing party to change the results solely through increasing advertising, at least not with a reasonable advertising budget. For example, for the 2008 election, the Republicans could have won Indiana with an extra budget of \$1.1 million, North Carolina with \$3.8 million, Iowa with \$7.3 million, Florida with \$11.8 million, Virginia with \$11.9 million, and Ohio with \$16.4 million. The extra spending adds up to \$52.3 million, roughly half of McCain's total campaign ad spending during the general election period; however, this would still make him short of 270 by 4 ECs. Similarly, if the PACs supporting McCain had spent \$6.8 million more, they could have won 48 additional electoral votes (i.e., North Carolina, Ohio, and Virginia); but this still could not have made up for the additional 97 electoral votes that he needed to win the election. In 2012, the Republicans could have increased their own ads to win New Hampshire (\$1.2 million more spending, the equivalent of 1.8 times more GRPs), Virginia (\$8.4 million more spending, doubling the existing GRPs), and Ohio (\$10.6 million more spending, the equivalent of 80% more GRPs). Or, the PACs could have helped Romney win Florida (\$7.4 million more spending), Minnesota (\$5.6

million), New Hampshire (\$2.7 million), and Pennsylvania (\$7.6 million). That additional spending would have exceeded one third of the total Republican-leaning PAC ad spending that year, but still could not have reached the 270 goal. And it would be prohibitively expensive for Romney to win more states simply by increasing advertising, being it the candidate's own ads or outside ads.

## **5 Conclusion**

We study the effect of mass-media advertising and personal selling—in the form of field operations—in the context of U.S. presidential elections. By linking various campaign activities to county-level vote results, we offer a comprehensive identification of the causal effect for various types of campaign activities. Different from most extant studies, we separate candidate campaign ads from those sponsored by outside political groups and examine how the ad effect varies by the types. Our results generate insights into the effectiveness of each campaign activity for different voter segments: field operations and outside ads are more effective for partisan voters, while candidate's own ads are only effective among non-partisans.

With our parameter estimates, we predict counterfactual election results under several hypothetical scenarios. Overall, we find that political campaigns play an essential role in shaping the election results, contrary to the “minimum effect of campaigning” view, which claims that most voters already have their minds made up and, hence, campaigns barely move the needle in terms of voting outcomes. We show that ground campaigning was critical to the Democrats: if neither party had implemented any ground operations, the Republicans would have won the 2008 and 2012 elections. We also find that advertising can play a critical role in a close election but not so when one party has a big advantage: with a modest amount of additional ads, the

Democrats would have won the 2004 election but the results would have been largely the same for the next two elections.

Some of the results merit further discussions. First, the finding that PAC ads behave similarly to field operations rather than candidate's own ads is surprising at a first glance. While candidate's own ads are found to be more effective for non-partisan voters, the opposite is true for PAC ads. We think that this is perhaps due to the difference in ad content: PAC ads are predominately negative and tend to attack rivals rather than promote the preferred candidates. Such a strong negative tone may work better to reinforce a partisan's beliefs than to persuade an undecided voter. This is complementary to Finkel and Geer (1998), which also found that voter partisan predisposition moderates the effect of negative ads.

Second, the finding that the Democratic field operations are more effective than the Republican's is also worth a closer examination. The field-operation own elasticity for the Democrats is estimated to be 2.9 times as large as that for the Republicans. We believe this is perhaps due to the quality of voter outreach activities resulting from the data available on voters and the techniques used to target and persuade them<sup>42</sup>. For example, personal voter interactions like door-to-door visits could be more powerful than indirect contacts such as telemarketing and door hangers. This is best echoed by a quote from an Obama field director in 2012: “Many field campaigns have historically favored quantity over quality. We do not. These are not phone calls made from a call center. They are done at the local level by our neighborhood team leaders, members and volunteers, who are talking to people in their communities.”<sup>43</sup> Despite the

---

<sup>42</sup> In the same vein, there might be reasons to suspect that field offices even from the same party may have a heterogeneous effect due to operational differences. We thank an anonymous reviewer for pointing this out. Due to data limitations we again exact away from this potential effect heterogeneity.

<sup>43</sup> *CNN*, “Analysis: Obama won with a better ground game.” November 7, 2012.

importance, detailed data on how field teams operate are challenging to obtain, especially at the county level. We acknowledge this data limitation and believe that future research could benefit from improving measuring the operation of ground campaigning.

In the same vein, our advertising measurement also has its limitation. In particular, we assume that the individuals from the same DMA (hence county) face homogeneous ad exposures, while in reality people may endogenously decide how much advertising to watch. We do not think this assumption would explain our main finding that candidate's own ads are only effective on non-partisans. The null effect on partisans is not due to the lack of ad exposures, because partisans tend to be more attentive and mindful to political ads than non-partisans (Finkel and Geer 1998). However, the difference in individual ad impressions could help explain the remaining variation in the ad effect, which could be interesting to explore further.

We would also like to point out that in this study we use the total spending of each party to measure the level of their digital campaigning. Note that the term, digital campaigning, is an umbrella concept encompassing various forms of campaign activities on digital platforms. Google search words, text-based banner display ads, online video ads, and social media ads are just several common examples that have entered the toolkit of presidential campaigns. In this study we are agnostic to the mechanism and the effectiveness of those various types, which could be a full-fledged paper on its own. Digital campaigning, despite its rapid growth, remained a relatively small venue in the recent elections. For this reason, we consider that controlling for the total online spending is sufficient for estimating the causal effect of advertising and ground campaigning. But we do see that examining the role of digital campaigning in presidential elections is an important and fruitful direction for future work.

## REFERENCES

- Albers, S., M. K. Mantrala, S. Sridhar. 2010. Personal selling elasticities : A meta-analysis. *Journal of Marketing Research* **47**(5) 840–853.
- Alvarez, R. M., A. Hopkins, B. Sinclair. 2010. Mobilizing Pasadena Democrats: measuring the effects of partisan campaign contacts. *The Journal of Politics* **72**(1) 31–44.
- Anderson, E. W. 1998. Customer satisfaction and word of mouth. *Journal of Service Research* **1**(1) 5–17.
- Angrist, Joshua D., J. Pischke. 2008. *Mostly harmless econometrics: An empiricist's companion*. Princeton, NJ: Princeton University Press.
- Aral, S., C. Dellarocas, D. Godes. 2013. Social media and business transformation: A framework for research. *Information Systems Research* **24**(1) 3–13.
- Armstrong, M. 2006. Competition in two-sided markets. *The RAND Journal of Economics* **37**(3) 668–691.
- Ataman, M. B., H. J. van Heerde, C. F. Mela. 2010. The long-term effect of marketing strategy on brand sales. *Journal of Marketing Research* **47**(5) 866–882.
- Ayers, R. M., R. A. Collinge. 2003. *Microeconomics: Explore and Apply*. Prentice Hall, London.
- Berger, J., A. T. Sorensen, S. J. Rasmussen. 2010. Positive effects of negative publicity: When negative reviews increase sales. *Marketing Science* **29**(5) 815–827.
- Berry, S., J. Levinsohn, A. Pakes. 1995. Automobile prices in market equilibrium. *Econometrica* **63**(4) 841–890.
- Berry, S. T. 1994. Estimating discrete-choice models of product differentiation. *The RAND Journal of Economics* **25**(2) 242–262.
- Borden, N.H. 1964. The concept of the marketing mix. *Journal of Advertising Research* **4**(2) 2–7.

- Bruce, N. I. 2008. Pooling and dynamic Forgetting effects in multitheme advertising: Tracking the advertising sales relationship with particle filters. *Marketing Science* **27**(4) 659–673.
- Caillaud, B., B. Jullien. 2003. Chicken & egg: competition among intermediation service providers. *The RAND Journal of Economics* **34**(2) 309–328.
- Campbell, J. E. 1992. Forecasting the presidential vote in the states. *American Journal of Political Science* **36**(2) 386–407.
- Carroll, V. P., A. G. Rao, H. L. Lee, A. Shapiro, B. L. Bayus. 1985. The Navy enlistment marketing experiment. *Marketing Science* **4**(4) 352–374.
- Chevalier, J. A., D. Mayzlin. 2006. The effect of word of mouth on sales: Online book reviews. *Journal of Marketing Research* **43**(3) 345–354.
- Chintagunta, P. K., S. Gopinath, S. Venkataraman. 2010. The effects of online user reviews on movie box office performance: Accounting for sequential rollout and aggregation across local markets. *Marketing Science* **29**(5) 944–957.
- Chung, D. J. 2013. The dynamic advertising effect of collegiate athletics. *Marketing Science* **32**(5) 679–698.
- Crawford, G. S., A. Yurukoglu. 2012. The welfare effects of bundling in multichannel television markets. *American Economic Review* **102**(2) 643–685.
- Darr, J. P., M. S. Levendusky. 2014. Relying on the ground game: The placement and effect of campaign field offices. *American Politics Research* **42**(3) 529–548.
- De Kluyver, C. A., R. J. Brodie. 1987. Advertising-versus-marketing mix carryover effects: An empirical evaluation. *Journal of Business Research* **15**(3) 269–287.
- Dekimpe, M. G., D. M. Hanssens. 1995. The persistence of marketing effects on sales. *Marketing Science* **14**(1) 1–21.

- Draganska, M., D. Klapper, S. B. Villas-Boas. 2010. A larger slice or a larger pie? An empirical investigation of bargaining power in the distribution channel. *Marketing Science* **29**(1) 57–74.
- Dube, J. P., J. T. Fox, C. L. Su. 2012. Improving the numerical performance of static and dynamic aggregate discrete choice random coefficients demand estimation. *Econometrica* **80**(5) 2231–2267.
- Dubé, J. -P., G. J. Hitsch, P. K. Chintagunta. 2010. Tipping and concentration in markets with indirect network effects. *Marketing Science* **29**(2) 216–249.
- Elberse, A., B. Anand. 2007. The effectiveness of pre-release advertising for motion pictures: An empirical investigation using a simulated market. *Information Economics and Policy* **19**(3) 319–343.
- Elberse, A., J. Eliashberg. 2003. Demand and supply dynamics for sequentially released products in international markets: The case of motion pictures. *Marketing Science* **22**(3) 329–354.
- Finkel, S. E., J. G. Geer. 1998. A spot check : Casting doubt on the demobilizing effect of attack advertising. *American Journal of Political Science* **42**(2) 573–595.
- Gatignon, H., D. M. Hanssens. 1987. Modeling marketing interactions with application to salesforce effectiveness. *Journal of Marketing Research* **24**(3) 247–257.
- Gelman, A., J. B. Carlin, H. S. Stern, D. B. Rubin. 2003. *Bayesian Data Analysis* (2nd Edition). CRC Press, 182–183.
- Gerber, A. S., D. P. Green. 2000. The effects of canvassing, telephone calls, and direct mail on voter turnout : A field experiment. *The American Political Science Review* **94**(3) 653–663.
- Givon, M., D. Horsky. 1990. Untangling the effects of purchase reinforcement and advertising carryover. *Marketing Science* **9**(2) 171–187.

- Godes, D., D. Mayzlin. 2009. Firm-created word-of-mouth communication: Evidence from a field test. *Marketing Science* **28**(4) 721–739.
- Goldstein, Kenneth, J. Rivlin. 2008. *Presidential advertising Combined File*. Madison, WI: The University of Wisconsin Advertising Project, The Department of Political Science at The University of Wisconsin-Madison.
- Gordon, B. R., W. R. Hartmann. 2013. Advertising effects in presidential elections. *Marketing Science* **32**(1) 19–35.
- Gowrisankaran, G., A. Nevo, R. Town. 2015. Mergers when prices are negotiated: Evidence from the hospital industry. *American Economic Review* **105**(1) 172–203.
- Grennan, M. 2013. Price discrimination and bargaining: Empirical evidence from medical devices. *American Economic Review* **103**(1) 145–177.
- Hagiu, A., R. S. Lee. 2011. Exclusivity and control. *Journal of Economics and Management Strategy* **20**(3) 679–708.
- Hansen, L. P. 1982. Large sample properties of generalized method of moments estimators. *Econometrica* **50**(4) 1029–1054.
- Hausman, J. A. 1996. Valuation of new goods under perfect and imperfect competition. In *The Economics of New Goods*. Chicago: University of Chicago Press, 207–248.
- Ho, K. 2006. The welfare effects of restricted hospital choice in the US medical care market. *Journal of Applied Econometrics* **21**(7) 1039–1079.
- Horn, H., A. Wolinsky. 1988. Bilateral monopolies and incentives for merger. *The Rand Journal of Economics* **19**(3) 408–419.
- Jullien, B. 2005. Two-sided Markets and Electronic Intermediaries. *CESifo Economic Studies*, **51**(2-3) 233–260.



- Katz, M. L., C. Shapiro. 1994. Systems competition and network effects. *Journal of Economic Perspectives* **8**(2) 93–115.
- Lee, R. S. 2013. Vertical integration and exclusivity in platform and two-sided markets. *American Economic Review* **103**(7) 2960–3000.
- Li, X., L. M. Hitt. 2008. Self-selection and information role of online product reviews. *Information Systems Research* **19**(4) 456–474.
- Liu, Y. 2006. Word-of-mouth for movies: Its dynamics and impact on box office revenue. *Journal of Marketing* **70**(3) 74–89.
- Lodish, L. M., M. Abraham, S. Kalmenson, J. Livelsberger, B. Richardson, M. E. Stevens, B. Lubetkin. 1995. How TV advertising works : A meta-analysis of 398 real world split cable TV advertising experiments. *Journal of Marketing Research* **32**(2), 125–139.
- Lovett, M., M. Peress. 2015. Targeting political advertising on television. *Quarterly Journal of Political Science* **10**(3) 391–432.
- Manchanda, P., Y. Xie, N. Youn. 2008. The role of targeted communication and contagion in product adoption. *Marketing Science* **27**(6) 961–976.
- Masket, S. E. 2009. Did Obama’s ground game matter? The influence of local field offices during the 2008 presidential election. *Public Opinion Quarterly* **73**(5) 1023–1039.
- Miller, A. R., C. Tucker. 2013. Active social media management: The case of health care. *Information Systems Research* **24**(1) 52–70.
- Misra, S., S. K. Mohanty. 2008. Estimating bargaining games in distribution channels. *Working paper*.
- Moe, W. W., D. A. Schweidel. 2012. Online product opinions: Incidence, evaluation, and evolution. *Marketing Science* **31**(3) 372–386.

- Narayanan, S., R. Desiraju, P. K. Chintagunta. 2004. Investment implications for pharmaceutical role expenditures : Interactions marketing-mix. *Journal of Marketing* **68**(4) 90–105.
- Nerlove, M., K. J. Arrow. 1962. Optimal advertising policy under dynamic conditions. *Economica* **29**(114) 129–142.
- Nevo, A. 2001. Measuring market power in the ready-to-eat cereal industry. *Econometrica* **69**(2) 307–342.
- Reynaerts, J., R. Varadhan, J. C. Nash. 2012. *Enhancing the convergence properties of the BLP (1995) contraction mapping*. No. 35 working paper Katholieke Universiteit Leuven, Faculteit Economie en Bedrijfswetenschappen, Vives.
- Rochet, J. C., J. Tirole. 2003. Platform competition in two-sided markets. *Journal of the European Economic* **1**(4) 990–1029.
- Rochet, J.C., J. Tirole. 2006. Two-sided markets : A progress report. *The RAND Journal of Economics* **37**(3) 645–667.
- Rossi, P., G. Allenby, R. McCulloch. 2005. *Bayesian Statistics and Marketing*. John Wiley & Sons, 66.
- Salganik, M. J., D. J. Watts. 2009. Web-based experiments for the study of collective social dynamics in cultural markets. *Topics in Cognitive Science* **1**(3) 439–468.
- Shachar, R. 2009. The political participation puzzle and marketing. *Journal of Marketing Research* **XLVI**, 798–815.
- Shankar V., B. L. Bayus. 2003. Network effects and competition: An empirical analysis of the home video game industry. *Strategic Management Journal* **24** 375-384.
- Shao, L. 2015. *Determinants of exclusivity , platforms' contracting decisions, and social welfare*, working paper.

- Shaw, D. R. 1999. The effect of TV ads and candidate appearances on statewide presidential votes, 1988-96. *The American Political Science Review* **93**(2) 345–361.
- Stata, 2013. Stata base reference manual release 13. , 950.
- Subramanian, U. R. C. Rao. 2016. Leveraging experienced consumers to attract new consumers: An equilibrium analysis of displaying deal sales by daily deal websites. *Management Science*.
- Sudhir, K. 2001. Competitive pricing behavior in the auto market : A structural analysis. *Marketing Science* **20**(1) 42–60.
- Teixeira, T., M. Wedel. 2010. Moment-to-moment optimal branding in TV commercials: Preventing avoidance by pulsing. *Marketing Science* **29**(5) 783–804.
- Tellis, G. J. 1988. The price elasticity of selective demand: A meta-analysis of econometric models of sales. *Journal of Marketing Research* **25**(4) 331–341.
- Toubia, O., A. T. Stephen, A. Freud. 2011. Viral marketing: A large-scale field experiment. *Economics, Management, and Financial Markets* **3** 43–65.
- Train, K.E. 2009. *Discrete Choice Methods with Simulations* 2nd ed., Cambridge University Press.
- Varadhan, R., C. Roland. 2008. Simple and globally convergent methods for accelerating the convergence of any EM algorithm. *Scandinavian Journal of Statistics* **35**(2) 335–353.
- Weitz, B. A. 1981. Effectiveness in sales interactions: a contingency framework. *Journal of Marketing* **45**(1) 85–103.
- Wilbur, K.C. 2008. A two-sided, empirical model of television advertising and viewing markets. *Marketing Science* **27**(3) 356–378.

Wooldridge, J. M. 2002. *Econometric Analysis of Cross Section and Panel Data*. Cambridge, Mass.: MIT Press, 752.

Wu, F., B. A. Huberman. 2008. How public opinion forms. In *Internet and Network Economics*. Springer Berlin Heidelberg, 334–341.

Zoltners, A. A., P. Sinha, S. E. Lorimer. 2006. *The Complete Guide to Sales Force Incentive Compensation: How to Design and Implement Plans That Work*. New York: AMACOM.