Essays on Consumer Switching and Search Behavior

Submitted in partial fulfillment of the requirements for

the degree of

Doctor of Philosophy

in

Engineering and Public Policy

Qiwei Han

B.S., Computer and Software Engineering, Beihang University

M.S., Information Networking, Carnegie Mellon University

M.S., Engineering and Public Policy, Carnegie Mellon University

Carnegie Mellon University
Pittsburgh, PA

May, 2017

For my family

Abstract

As recommender systems have increasingly become prevalent to guide consumers to find their desired products in many industries, understanding the impact of recommender systems on consumer choices is critical to the business performance and raises important policy implications. In this thesis, we examine the role of different recommendation schemes, spanning from interpersonal recommendations in social environment given by peers to product display recommendations in physical shopping environment given by sellers on consumers' switching and search behavior in two distinct case studies. In the first study, we look at the effect of peer recommendations on subscriber churn in a large mobile network. We find that consumers' propensity to churn increases with the number of friends that churn and in particular with the number of strong friends that churn. In the second study, we implement an in-vivo randomized field experiment to measure the effect of product display recommendations as book placement on shopper behavior in a physical bookstore. We leverage video tracking technologies to monitor how shoppers respond to random book placement, which induces random search costs. We find that books recommended at the edge of the table are more likely to be picked and taken than those placed at the center of the table. More interestingly, we also find that conditional on being picked, shoppers are equally likely to take books placed at the edge and at the center of the table, suggesting that display recommendations positively affect consumer choice mainly through its effect on the search process and not through its effect on the consideration process. Therefore, we empirically show that provision of recommendations, although in different schemes, may generally help to reduce consumers' search costs in product or service discovery process, relative to what they would do without such an intervention. Moreover, we perform a comparative analysis between offline and online applications of recommender systems to systematically investigate the current practices, future prospects and policy perspectives when applying recommender systems in physical retailing. All these issues highlight opportunities for physical retailers to design, implement and evaluate their recommender systems that offer convenience benefits and appropriate protection to consumers.

Acknowledgments

This thesis would not be possible to be completed without the support from so many people that I have great appreciations in mind.

First of all, I would like to acknowledge the financial support by the Fundação para a Ciência e a Tecnologia (Portuguese Foundation for Science and Technology) under Grant [UID/EEA/50009/2013] and through the Carnegie Mellon Portugal Program under Grant [SFRH/BD/51153/2010]. I would also like to acknowledge the two industrial partners that grant permission for me to have access to the dataset and perform field experiment on site.

I am deeply indebted to my advisor Prof. Pedro Ferreira, for his longstanding support throughout my Ph.D. studies. I would never forget the day that I met with Prof. Ferreira in a beautiful May afternoon in the EPP office many years ago, when he encouraged me to pursue Ph.D. degree in EPP and accepted me as his student. I benefit so much from his valuable guidance on my research.

I am also indebted to my co-advisor Prof. João Paulo Costeira for his faithful mentorship all along my work. Not only does he provide valuable inputs to my research, but also he takes good care of my personal needs that are very important in the late stage of my Ph.D. studies.

I am furthermore appreciative to my committee members, Prof. Nicolas Christin, Prof. Nuno Nunes and Prof. Leid Zejnilović for their interest and useful feedback to my thesis.

I am grateful to excellent colleagues from CMU-Portugal program, EPP department and Signal and Image Processing Group at ISR, whom I had great pleasure to work with. I give my special thanks to João Carvalho, who helped so much in my research and as a great companion in the lab.

I also owe my sincere thanks to all the administrative staff at EPP and IST whose help makes my life much easier. Especially, I am very grateful to Ana Mateus, who cared about me like a family and I wish her a great future.

Finally, I would like to express my gratitude to my family, who provide me with continued support throughout my life. In particular, I am extremely thankful to Dr. Chen Wang, my best friend and life partner for her understanding and encouragement that always lead me to achieve our goals.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

> Human beings, viewed as behaving systems, are quite simple. The apparent complexity of our behavior over time is largely a reflection of the complexity of the environment in which we find ourselves.
>
> *Herbert A. Simon, The Sciences of the Artificial*

Consumers, in most cases, need to search for the product information prior to make their choices. Typically, the decision making process is bounded by the cognitive limitations to gather information up to a point beyond which the benefits of acquiring additional information is no longer worth the costs of search (Stigler, 1961). Technological changes have significantly affected the information structure of the market, such as the types and costs of information available to consumers, which in turn affect individual choices and market outcomes. For example, information systems now serve as intermediaries that allow consumers to be better informed about the product offerings and share quality information to others, which reduce consumers' search costs and uncertainty that are believe to be the main deterrents in product purchase decisions (Bakos, 1997). In particular, as it is often necessary for consumers to seek product recommendations from other people, either by word of mouth or product ratings and reviews, recommender systems have increasingly become prevalent to automate this social process in large scale (Resnick and Varian, 1997). As such, understanding the impact of recommender systems on consumer behavior is critical to business performance in many industries and raises important policy implications.

In this thesis, we aim to study the role of recommendations in influencing consumer choices in different contexts by addressing several key challenges. First, collecting consumer data, such as interpersonal communications for peer recommendations, or interaction between consumers and product recommendations is not so obvious without appropriate support from digital technologies. Second, measuring the causal effect of recommendations on consumer choices requires robust identification strategies and careful experimental designs to avoid the misattribution of endogenous provision of recommendations. Third, although recommender systems are now ubiquitous in the online channel, systematic investigation on the uses in offline channel is still very limited.

Essentially, our studies cover a variety of recommendation schemes, spanning from interpersonal recommendations in social environment given by peers to product display recommendations in physical shopping environment given by sellers. We postulate that recommendations, although in different schemes, may generally help to reduce consumers' search costs in product or service discovery process. In the two distinct case studies, we empirically show that consumers' switching and search behavior are influenced by the provision of recommendations, relative to what they would do without such an intervention. Moreover, we perform a comparative analysis between offline and online applications of recommender systems and focus specifically on the policy perspectives that arise from the potential uses in offline world.

In the first study (Chapter 2), we examine the effect of peer recommendations on consumer switching behavior (also known as churn) in a large mobile network. We try to disentangle it from other effects that drive simultaneous churn across friends but that do not relate to peer influence. We analyze a random sample of roughly 10 thousand subscribers from large dataset from a major wireless carrier over a period of 10 months. We apply survival models and generalized propensity score to identify the role of peer influence. We show that the propensity to churn increases when friends do and that it increases more when many strong friends churn (Han and Ferreira, 2014). We also check the robustness of our results using alternative functional forms of the model. Therefore, our results suggest that churn managers should consider strategies aimed at preventing group churn.

In the second study (Chapter 3), we implement an in-vivo randomized field experiment to measure and analyze shopper behavior at the point of purchase in a physical bookstore. We leverage video tracking technologies to monitor how shoppers respond to random book placement, which induces random search costs. More specifically, we randomize the position of newly released books on the top of a large table with several rows and columns such that each book's search cost becomes independent of the book's characteristics. We use advanced 3D cameras and vision-understanding algorithms that can track human motions in real-time to overcome the large costs associated to large-scale video data. We find that on an average day books placed at the edge of the table are both picked and taken more often by consumers than books placed in the center of the table. However, the likelihood of taking a book that was picked is on average similar for books placed at the edge and at the center of the table, that is, books at the edge of the table sell more only because they are, on average, picked more often. Armed with this knowledge, the bookstore manager may maximize profit by placing books with higher margins at the edge of the table.

In the third study (Chapter 4), we perform an interdisciplinary approach to systematically investigate the current practices, future prospects and policy implications when applying recommender systems in physical retailing. We extend the conceptual framework for personalization and discuss about the existing solutions and potential opportunities in designing and implementing recommender systems in each stage of recommendation process. More specifically, we find that physical recommender system applications are mostly isolated prototypes that focus on solving technology challenges but fail to reflect the impact on consumer behavior. However, as physical retailers may have great potential in building ubiquitous shopping environment to capture consumer preferences through many different sensors in retail stores, they can leverage the rich consumer data to incorporate contextual information to improve the recommender system design. We discuss about policy implications for the use of physical recommender system from three perspectives: 1) welfare implications when physical retailers seek to increase profits through manipulating the design of recommender systems; 2) privacy concerns when consumer data is collected, stored and processed using emerging technologies; and 3) importance of human interpretability in algorithmic recommendation techniques. All these issues highlight opportunities for retailers to design, implement and evaluate their recommender systems that offer convenience benefits and appropriate protection to consumers.

Chapter 2

Role of Peer Influence in Subscriber Churn in Wireless Networks

2.1 Introduction

Over the past two decades, wireless industry experienced rapid change and strong technological progress from 2G to 4G. According to ITU (2015), over 7 billion mobile subscriptions worldwide increasingly saturate the wireless market. In parallel, deregulation opened up sectors to multiple entrants supporting both competition and technological innovation. Consequently, carriers need to invest heavily in acquiring spectrums and upgrading their networks to provide quality communications and novel services as well as to ensure that they healthily profit from existing subscribers. However, subscribers have many providers to choose from and can ever more easily transfer from one provider to another, as more information about products and services abounds and switching costs reduce (Cho et al., 2012).

Churn rates measure the proportion of subscribers discontinuing service during a certain period of time. As reported by wireless carriers across the world, average monthly churn rates vary between 1.5% and 5% (e.g. Wei and Chiu, 2002; Ahn et al., 2006; FCC, 2011). In other words, wireless carriers can lose over 20% of their subscriber base every year, which poses significant challenges for profitability and growth. Subscriber churn may represent significant economic loss. This loss can be estimated by multiplying the average cost to acquire a new subscriber by the number of subscribers that churn. With average acquisition costs varying between $300 and $600 per subscriber, churn may cost the wireless industry billions of dollars every year (Berson et al., 2000). On the other hand, keeping an existing subscriber is generally much cheaper and easier. Mozer et al. (2000) showed that acquiring a new subscriber can be at least five times harder than retaining an existing one. Meanwhile, improving subscriber retention can help wireless carriers increase margins because existing subscribers are less sensitive to both price and sales referrals (Wei and Chiu, 2002). Therefore, churn represents "the biggest issue for all the carriers" and effective subscriber churn management becomes a priority for telecom managers as to ensure the sustainable growth of their companies (FCC, 2009).

Churn rate is also used by regulators to monitor the competitive dynamics within the wireless industry. For example, since 1995, U.S. Federal Communication Commission regularly released annual reports on competitive market condition with respect to commercial mobile services, as mandated by the The Omnibus Budget Reconciliation Act of 1993 (FCC, 2011). In these reports, churn rate has been adopted as the proxy measure of switching barrier that prevents subscribers from changing carriers. When regulators undertake policy interventions (*e.g.* mobile number portability that allows subscribers to retain their phone number when switching carriers) intending to ease the switching cost and facilitate the competition, how churn rate changes can be very indicative of the effectiveness of interventions (Cho et al., 2012).

Wireless carriers aim at controlling churn through active subscriber retention campaigns. For this purpose, they proactively identify subscribers with high propensity to churn, evaluate the underlying reasons for churn and devise strategies to prevent it. The perplexing nature of churn, however, makes it very difficult to explain and address churn in an efficient and comprehensive manner. Subscribers may churn for many different reasons (Braun and Schweidel, 2011; Liang et al., 2011). Gerpott et al. (2015) generally categorized these reasons into three streams. First, they may opt out due to the unsatisfication with the service quality (Eshghi et al., 2007; Kim and Yoon, 2004; Liang et al., 2011). Second, they may get induced by competing carriers that provide more attractive service offerings (Shaffer and Zhang, 2002) or decide to acquire a new handset or service that is either not compatible with or not provided by their carrier (for example, Apple's iPhone had exclusive arrangement with only one carrier in one country until 2011, see Cho et al. (2016) for details). Third, changes in subscribers' personal communication needs may lead their valuation of existing service to become not attractive anymore. For example, they may be persuaded by close friends to switch to another carrier, simply because they need to maintain the communications with the friends, while also ensure their current arrangement meet their needs. Thus, wireless carriers can hardly provide one single solution to prevent all potential churners from leaving. Therefore, understanding the complexity of the churn problem and disentangling the role of several factors that can trigger it is fundamental to design sound retention strategies.

In this paper, we look at one such complexity associated with churn, which largely is related to the third factor. We study the effect of peer influence on churn. If churn is contagious then churn can snowball quickly leading many subscribers to leave the carrier, specially when social networks are dense, as they locally tend to be in the case of wireless services. We apply survival analysis and generalized propensity score to separate peer influence from homophily and measure the effect of the former. We perform our empirical analysis on a massive dataset from a major European wireless carrier that shared call detail records (CDRs) with us. We show how churn increases with friends' churn, a result suggesting that churn managers should prevent group churn instead of looking at churn on an individual basis.

2.2 Literature Review

2.2.1 Customer Lifetime Value (CLV)

Many researchers have noted that understanding the value of existing subscribers is imperative because retention campaign can incur additional costs and investing in all subscribers (e.g. loyal subscribers) will be inefficient (Ahn et al., 2006; Hung et al., 2006). Rosset et al. (2002) introduced the CLV model in the wireless industry to estimate the effect of retention campaigns based on CLV. CLV models provide wireless carriers insights to build cost-effective retention strategies targeted to subscriber segments with high CLV and low loyalty. Both Gupta et al. (2004) and Hwang et al. (2004) proposed that the profit of the wireless carrier is a function of the total subscriber lifetime value. In their theoretical CLV models for the wireless telecommunication industry, the probability of a subscriber to churn is treated as a parameter in the CLV function used to determine how long the subscriber will stay in the network generating future cash flows. Another generalized CLV model proposed by Glady et al. (2009) identified churners as subscribers with decreasing CLV. Therefore, the probability of churn is central to the notion of CLV and thus understanding churn is paramount to correctly measure CLV. When carriers start to apply retention efforts to subscribers, the effect of churn reduction can be straightforwardly represented as the accrued incremental CLV (Braun and Schweidel, 2011) .

2.2.2 Churn Prediction

Today, wireless carriers gather wealthy data deemed useful to perform churn analysis (Blondel et al., 2015; Huang et al., 2015). Numerous data mining techniques have been applied to transform these raw data into useful knowledge. Hung et al. (2006) described the general framework for this purpose: i) identify discriminatory features that can differentiate a subscriber with high risk of churn from other loyal subscribers; ii) extract and transform data from identified features; iii) select the appropriate data mining techniques to build descriptive or predictive models; iv) evaluate the performance of these models according to specified criteria, e.g. lift curves. Extensive research has been empirically performed on churn prediction in wireless network. Refer to Ngai et al. (e.g. 2009); Verbeke et al. (e.g. 2011) for comprehensive reviews on the application of different algorithms to churn prediction using the collected mobile data.

Three disadvantages of pure data mining techniques are worth mentioning though. First, although many models and algorithms seem to provide satisfactory accuracy in identifying churners, the results obtained dependent not only on method but also on the data used and the features considered by researchers. For example, both Mozer et al. (2000) and Hwang et al. (2004) used logistic regressions and neural networks to predict churn. The former found that neural networks outperformed logistic regression. However, the latter concluded otherwise. Neslin et al. (2006) showed that the differences observed in accuracy of churn prediction algorithms may significantly change the profitability of a churn management campaign. Second, many data mining algorithms (e.g. ensemble methods used in Lemmens and Croux (2006)) are like a “black box”, which lack interpretability precluding us from understanding the true determinants of subscriber churn. As a result, an agent in a call center might be asked to call a certain subscriber because she is likely to churn. However, very little might have been said to this agent about the underlying reasons that may lead the subscriber to churn, which clearly difficulties the interaction with her. Third, a number of statistical based benchmarking measures of performance for data mining algorithms do not directly yield optimal results in terms of profit maximization from the practitioners’ perspective. Verbeke et al. (2012) proposed a profit-centric performance criterion focusing on the fraction of subscribers that generate the highest profits and showed that this approach yielded outcomes different from the ones resting on the best approaches as evaluated by statistically based performance measures. Lemmens and Gupta (2013) further formulated a predictive model with profit-based loss function and combined it with the optimal target size selection in terms of profit maximization. They showed significant improvement in profits compared to current methods.

2.2.3 Social Influence

Advances in studying the effect of social influence on subscriber churn in wireless networks have received considerable attention in recent times. Dasgupta et al. (2008) tried to learn whether the propensity of a subscriber to churn depends on the number of friends that have already churned.

This hypothesis is based on premise that a few key individuals may lead to strong "word-of-mouth" effects. These individuals may influence their friends to churn, who in turn spread the message to others. So they identified likely churners as those subscribers whose friends have already churned, using a spreading activation-based technique. A set of churners iteratively diffuse the message to other subscribers. Then a subscriber churns once the accumulated level of influence reaches a certain threshold. Both Nitzan and Libai (2011) and Haenlein (2013) showed that the propensity for focal subscriber to churn is positively associated with other individuals to whom are socially connected and have previously churned. Dierkes et al. (2011) used Markov Logic Networks and propositionalization to develop a predictive model for churn. They also confirmed that "word-of-mouth" has a significant impact on subscriber's churn decisions. Phadke et al. (2013) demonstrated that by integrating social factors such as influence from churners into machine learning models can greatly enhance the prediction performance.

However, correlation in the behavior among people who share social ties can be explained by both peer influence and their inherent similarities – homophily (McPherson et al., 2001). Work that identifies contagious churn separating it from confounding effects such as homophily (friends tend to exhibit similar behavior) remains unexplored at large, although a number of studies on the identification of peer influence in the presence of confounding effects in other networked contexts have been proposed (e.g. Bramoullé et al., 2009; La Fond and Neville, 2010; Steglich et al., 2010; Lewis et al., 2012; Matos et al., 2014; Han et al., 2015). One important contribution of this chapter is to analyze contagious churn avoiding misattributing homophily to contagion, or/and vice versa, which typically leads to overestimate the latter.

2.3 Data and Descriptive Statistics

We partnered with a major European wireless carrier, hereinafter called EURMO, which gave us access to its Call Detailed Records (CDRs) between August 2008 and May 2009. For each call we know the caller and the callee, the duration and time of the call and the id of the cell tower used to route the call. Subscribers are identified by their anonymized phone number. For each subscriber, we know their provider and tariff plan at all times. There are roughly 4 million EURMO subscribers in our dataset.

Understanding subscriber churn with prepaid plans is quite different from working with postpaid subscribers (e.g. Gerpott et al., 2015). First, we have very limited socio-demographic information on prepaid subscribers. Second, the usage pattern of prepaid subscribers is more irregular than that of postpaid subscribers. Third, prepaid subscribers churn by ceasing usage whereas postpaid subscribers explicitly inform the carrier when they intend to do so. We use the standard in the industry, which is also followed by EURMO and assumed that a prepaid subscriber churns if she places no calls for three of consecutive months.

We use a random sample of 10 thousand subscribers. Two subscribers are called friends if they exchange at least one call in each and every calendar month. We trim from our random sample subscribers with very high degree, which are likely to represent customer service and PBX machines, and with no degree (some subscribers purchase a SIM card but never use it to make calls). We are left with 8,345 subscribers in our sample. We observe network dynamics over time, namely new subscribers join EURMO and existing subscribers leave EURMO every month. Moreover, subscribers call and/or text different friends over time. We aggregate individual subscriber usage and structural properties at the monthly level in our analysis. Table 2.1

shows descriptive statistics for covariates used in our study. Over the period of analysis, the subscribers in our sample placed 3.75 million calls and 1,191 of them churned, which amounts to an average monthly churn rate of 2.04%. On average, users in our sample have 8.4 friends. Number of friends is the only time-independent covariate used in our study, which we compute across the whole panel of data available to us. This allows us to identify ”sticky” friends purging spurious short-lived connections.

We find that churners have much less usage, both in terms of number of calls and airtime, and fewer friends than non-churners. These differences are statistically significant as shown by a t-test in column [7] of this table. Since we know the tariff plan each subscriber holds, we are able to calculate estimated monthly expenses (in Euros), commonly *a.k.a* ARPU in wireless industry. Churners contribute with 11% of the revenues at EURMO. On average, the ARPU derived from the full sample is about 32 euros and that for churners is 7.5 euros less per month than non-churners. We observe that all subscribers have much more usage within the network. This is because EURMO operates under a Sender-Pays-All regime and interconnection charges are added to every call, which EURMO passes, partially, to consumers making calls outside the network more expensive (Hoernig, 2007).

We are particularly interested in the association between number of friends who churn and the propensity to churn. Therefore, we use $n\text{-}call\ frd\_churn$, the friends who churn that exchange at least n calls with the ego in the same calendar month. In this way, we can compare the extent to which the strength of a focal subscriber's social ties with churned friends contributes to that subscriber's churn probability (Nitzan and Libai, 2011). We find that 343 out of 1,191 (29%) churners and that 3,197 out of 7,154 (45%) non-churners have at least one friend that churned

Covariates	Description	Full Sample n=8,345 Mean [1]	Std [2]	Churner n=1,191 Mean [3]	Std [4]	Non-Churner n=7,154 Mean [5]	Std [6]	t-stat [7]
Monthly usage								
n_calls_out	Calls made	18.99	29.74	6.51	12.80	21.07	30.89	$p < 0.001$
n_calls_in	Calls received	22.96	29.45	7.84	16.44	25.48	30.70	$p < 0.001$
$airtime\_out$	Airtime made	44.87	116.44	13.55	40.87	50.08	123.88	$p < 0.001$
$airtime\_in$	Airtime received	55.54	116.63	17.71	50.15	61.83	123.17	$p < 0.001$
$expenditure$	Expenditure	31.78	53.36	25.33	55.31	32.85	52.95	$p < 0.001$
Structural properties								
frd	Number of friends	8.40	9.16	5.76	8.91	8.84	9.13	$p < 0.001$
$\%calls\_out\_other$	Ratio of calls made to other networks	0.24	0.23	0.27	0.27	0.24	0.23	$p < 0.001$
$\%calls\_in\_other$	Ratio of calls received from other networks	0.25	0.23	0.26	0.26	0.25	0.22	$p < 0.05$
$tenure$	Time with EURMO	48.10	39.32	19.91	25.66	52.79	39.24	$p < 0.001$
Churner friends								
1-*call* frd_churn	Number of 1-call churner friends	1.13	2.23	0.81	2.05	1.18	2.25	$p < 0.001$
3-*call* frd_churn	Number of 3-call churner friends	0.56	1.25	0.40	1.15	0.59	1.26	$p < 0.001$
5-*call* frd_churn	Number of 5-call churner friends	0.20	0.60	0.14	0.52	0.21	0.61	$p < 0.001$

Table 2.1: List of covariates extracted from the EURMO network. Descriptive statistics are performed for the our random sample (columns [1] and [2]), churners (columns [3] and [4]) and non-churners (columns [5] and [6]), respectively. Column [7] tests the hypothesis that the means between churners and non-churners are similar.

during the period of analysis. Table 2.1 also shows that on average subscribers see 1.13 friends churn, 0.56 and 0.20 friends churn that exchange at least 3 and 5 calls, respectively, with the ego in the same calendar month.

2.4 Churn Dynamics

As stated earlier, we organize data into a panel such that each unit of observation is a subscriber and each time period is a calendar month. As subscriber churn can be considered to be a time-to-event outcome so that a subscriber will no longer reappear in the sample once churned. This results in an unbalanced panel with the size of 51,488 observations. The panel structure of our data naturally allows us to determine the impact of time-varying covariates on churn using survival analysis (Haenlein, 2013). Survival analysis also allows for controlling for some unobserved individual-level heterogeneity, i.e., some subscribers are more prone to churn for reasons that are not captured in our data (e.g. marketing campaigns by competitors). We employ a Cox Proportional Hazard (PH) model with frailty terms to estimate the churn hazard rate as:

$$h(t, x_i, x_i^{(t-1)}) = \alpha_i h_0(t) \text{exp}[\beta x_i + \delta x_i^{(t-1)}] \tag{2.1}$$

where $h(t, x_i, x_i^{(t-1)})$ is the churn hazard for subscriber i at time t, α_i is the Gamma distributed frailty that represents the individual-level random effect, $h_0(t)$ is the non-paramatric baseline hazard function, x_i and $x_i^{(t-1)}$ include time-independent and time-varying covariates at time t, such as usage and friends' churn. We introduce a time-lag on the latter covariates because, similarly to epidemiology cases, there is an induction and latency period between a friend's churn and the ego's churn. Moreover, using lagged covariates obviates simultaneity problems with our estimation. Finally, we also introduce monthly dummies to control for seasonal effects on churn, such as promotions offered by competitors over the summer or during Christmas.

We count the number of friends who churn that exchange at least n calls with the ego, as explained in the previous section. We vary n in $\{1, 3, 5\}$ and use frd_churn to denote the covariate of interest. Then we apply the model across two specifications that mainly differ between number of calls and duration of airtime, respecively, because these two variables are functionally related measures of subscriber usage. Table 2.2 shows the results obtained from six regressions. The coefficients on frd_churn are statistically significant across all model specifications. Note that coefficients for $n = 5$ are larger but this can only mildly hint at the fact that churn from stronger friends is more relevant for the ego because standard errors are high. The other covariates behave as expected, in particular, higher expenditure leads to more churn. Yet, we recall that our goal in this paper is not to develop a predictive model of churn. Instead, we are interested in identifying the role of peer influence on churn. The additional covariates in our study are not necessarily used to predict churn. Instead, they are used to characterize the behavior of consumers with respect to how they use mobile service so that we can, in the next section, compare similar users, aiming at isolating the effect of peer influence.

These results show that the more friends churn the more likely the ego will churn. For example, 0.375 for the first model shown in this table indicates that if one more friend churns then the ego's likelihood of churn vs. no churn increases by $exp(0.375)$ or 45%. The 0.215 (0.458) in the second (third) model indicates that if one more friend with whom the ego exchanges 3 (5) or more calls in the same calendar month churns then the ego's likelihood of churn vs. no-churn increases by $exp(0.215)$ ($exp(0.458)$) or 24% (58%). To better understand the role of friends' churn on the ego's churn we perform Monte Carlo simulations for the relative hazard of churn vs. no-churn using the coefficients and variances of frd_churn estimated using the six Cox PH models in this table. Figure 2.1 shows how the relative hazard increases with the number of

Covariates	1-call [1]	3-call [2]	5-call [3]	1-call [4]	3-call [5]	5-call [6]
Influence factor						
frd_churn	0.375***	0.215***	0.458***	0.369***	0.333***	0.449***
	(0.104)	(0.0787)	(0.156)	(0.0978)	(0.120)	(0.167)
Usage Metrics						
n_calls	-0.00646***	-0.00621***	-0.00625***			
	(0.00159)	(0.00134)	(0.00148)			
n_calls^2	0.0000073***	0.0000071***	0.0000072***			
	(0.0000021)	(0.0000018)	(0.0000019)			
$airtime$				-0.000829**	-0.000801***	-0.000819*
				(0.000382)	(0.000297)	(0.000440)
$airtime^2$				0.00000013	0.0000001	0.00000014
				(0.0000001)	(0.0000001)	(0.00000011)
$expenditure$	0.00457***	0.00488***	0.00515***	0.00565***	0.00554***	0.00584***
	(0.00159)	(0.00127)	(0.00150)	(0.00158)	(0.00120)	(0.00184)
Structural Metric						
frd	-0.0746***	-0.0659***	-0.0671***	-0.0898***	-0.0803***	-0.0827***
	(0.0161)	(0.0143)	(0.0151)	(0.0150)	(0.0115)	(0.0184)
$month\_dummy$	Yes	Yes	Yes	Yes	Yes	Yes
Observations	51,488	51,488	51,488	51,488	51,488	51,488

***$p < 0.01$, **$p < 0.05$, *$p < 0.1$

Table 2.2: Parameter estimates for Cox PH frailty models on frd_churn at three thresholds: 1-$call$ (columns [1] and [4]), 3-$call$ (columns [2] and [5]) and 5-$call$ (columns [3] and [6])

friends that churn.

2.5 Effect of Friends' Churn

Propensity Score Matching (PSM) is a widely applied method to evaluate the effect of treatments on outcomes of interest with observational studies (Rosenbaum and Rubin, 1983; Mogan and

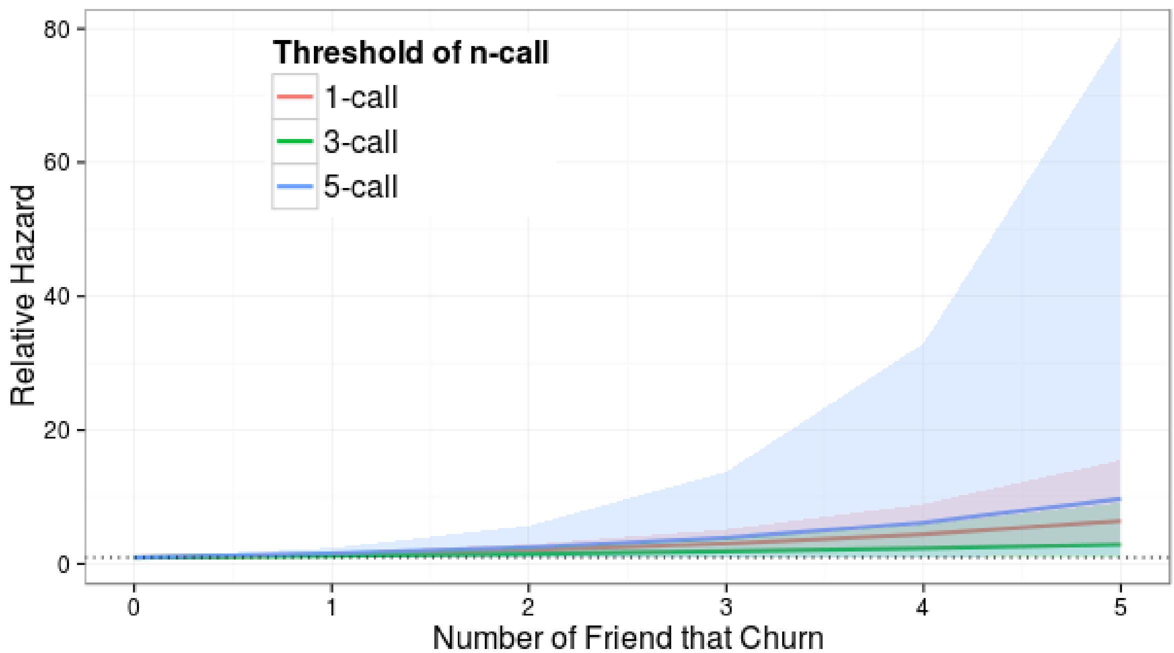

Figure 2.1: Simulated average churn hazards for models 1-3 in Table 2.2. Ribbons represent 95% confidence intervals. 10,000 simulations were run per value of frd_churn

Winship, 2014). In our setting, peer influence is associated with the presence of churner friends in subscriber's local social network (the treatment). As the assignment of treatment is not random due to homophily between connected subscribers, selection bias arise because the characteristics of treated and untreated units can differ. The propensity score, defined as the probability of receiving treatment conditional on observed confounding covariates that correlate with both the outcome and treatment, summarizes all the relevant information available to the researcher into a single scalar value. Conditioning on this value, the distributions of the observed characteristics across treated and untreated units become more similar and, therefore, unlikely to drive differences in outcomes. However, PSM still fails to provide full causal interpretations because it does not control for unobserved effects (Shalizi and Thomas, 2011). In any case, and when randomized experiments are unavailable, using PSM increases our confidence when reporting effects, in particular, when one controls for the most important characteristics of the units under analysis and so are latent characteristics that are correlated with these observed characteristics. For example, Aral et al. (2009) used PSM method to estimate the contagion effect in the adoption of an online service by analyzing a community of instant messenger users. They dichotomize the treatment level (number of adopter friends) to explore the heterogeneity in treatment effects, e.g. they compared the users having more than three adopter friends with those otherwise. However, the results obtained at one treatment level may "absorb" the effect at the next treatment level, and thus complicates the evaluation of cumulative and marginal peer influence.

Extensions of PSM to cases with continuous treatments (Imai and Van Dyk, 2004; Hirano and Imbens, 2005) have been proposed in the literature under the umbrella of Generalized Propensity Score (GPS). Essentially, GPS provides a dose-response function (DRF) that measures the relationship between the outcome of interest and the intensity of treatment. In our case, friends' churn (the treatment) is not binary but rather an integer. Egos can be subject to cumulative amounts of treatment as they see more friends churn. Different treatment intensities can have different effects on the ego's churn (the outcome). Therefore, we resort GPS to explore contagion effects from friends' churn in more detail.

Details of the GPS model setup and estimation procedure for contagious churn can be found in Appendix 2.1. For brevity, we only note the following modeling options: i) the distribution of friends' churn is far from normal, which violates the normality assumption when estimating the score through the maximum likelihood, therefore, we followed a more flexible parametric solution using the GLMDOSE procedure in Stata, as proposed in Guardabascio and Ventura (2014), and allow the intensity of treatment to be skewed towards zero, using an exponential functional form (with Poisson distribution and canonical logarithm link function), whose parameters are estimated by generalized linear model; ii) we use a polynomial approximation of order two to regress outcomes on treatments and propensity scores, from which we compute the average conditional expectation for the effect of treatment. This polynomial approximation of degree two, with an interaction term between treatment intensity and propensity score, allows for a better understanding of the non-linear effects of treatment than linear regression.

2.5.1 Description of Matching Panels

The panel structure of our data poses several extra challenges when applying matching techniques that have been developed for cross-sectional data (Nielsen and Sheffield, 2009). First, although our data is a random sample, standard matching routines on panel data typically ignore the time dimension in the panel and pair observations of the same unit in different time-periods (i.e. subscriber-month). Thus the systematic within-panel dependence may violate the independence assumption between matched observations (a.k.a stable unit treatment value assumption (SUTVA) in PSM setting). Second, standard matching routines would discard unmatched observations from the middle of some panels that may cause missing data problems. Third, compressing the data to one observation per panel (e.g. averaging covariates across time) may help alleviate the concerns of matching the same unit at different periods, however, important information regarding the dynamic subscriber behavior is inevitably eliminated. In our case, a downward trend in usage may be a signal of eventual churn. Mismatch of subscribers with similar average usage but divergent trend may lead to the biased conclusion. Therefore, analogous to (Nielsen and Sheffield, 2009), we choose to estimate the panel-level GPS as the unit of analysis that accounts for both the systematic dependence between observations of single subscriber and dynamic behavioral trend across time periods. Specifically, we include the lagged values of covariates prior to treatment in the GPS model, such that treatment assignment is applied to the panel that include all observations having the same subscriber identifier, rather than the individual observation. As discussed earlier, contagious churn is studied after the exposure to friends' churn during an induction and latency period. Therefore, we split the period of analysis into three intervals: i) the Pre-Treatment Period (PTP), during which we observe the important subscriber characteristics; ii) the Treatment Exposure Period (TEP), during which egos observe, and count, their friends churning, represented by frd_churn; iii) the Churn Observation Period (COP), during which we observe whether the ego churns. This definition of intervals can ensure the best possible match between treated and control panels because the matched pairs correspond to the same duration of intervals rather than the same calendar month. For example, a treated panel spanning the first three months may have a control panel spanning the last three months as its best match, such that the selection bias on observables is reduced to the most extent.

Our dataset spans 10 months of data. However, we need to observe subscribers for 3 months to determine whether they churn. So, in fact, we are limited to a panel with 7 time periods. There are several options to define intervals on these time periods and it is preferable that each of these intervals should be sufficiently large. PTP and TEP needs to be sufficiently large so that we can observe trends of covariates prior to treatment as well as several intensities of treatment. COP also needs to be sufficiently large so that we can observe outcomes and, in particular, outcomes triggered by treatment. Therefore, one natural choice is to include all available yet well-balanced time periods as: $\Big\{\text{PTP:}\{1,2,3\}, \text{TEP:}\{4,5\}, \text{COP:}\{6,7\}\Big\}$. Below we show results using this definition but certainly other options to divide intervals should work without further clarifications.

2.5.2 Description of GPS Analysis

The key for any GPS analysis to identify believable effects is to control for the important covariates. Given the richness of our dataset, we are fortunate to observe covariates that, allegedly, capture most of what is important to control for to study contagious churn. First, we control for number of calls placed and for the percentage of calls to other networks. These covariates together capture well the level of cell usage. Otherwise, it would be unreasonable to compare subscribers with low usage to subscribers with high usage as their engagement with cell phone service is likely very different. Second, we control for expenditure. While this covariate is highly correlated to usage, controlling for expenditure allows for reducing selection bias introduced by having subscribers choose their tariff plan. Indeed, consumers with different levels of income, and different tariff plans, can pay different amounts of money for the same level of call usage. Such differences could, therefore, be attributed to usage of services other than calling, such as data. Ensuring that both number of calls and expenditure are similar across subscribers allows us to be more confident that we are comparing subscribers that use their cell phone similarly, even for services whose usage we do not observe in our dataset. Note that both cell usage and expenditure are matched dynamically, namely we also control for lagged values in order to capture the trends of covariates.

Third, we control for tenure with EURMO, thus making sure that we take the life cycle of the subscriber within EURMO into account. Otherwise, it would be unreasonable to compare new consumers to old consumers as the latter typically have developed a different level of trust with EURMO, enjoy different prices and, most importantly for the issue of churn, experience different switching costs because their lock-in periods are more likely to have expired. Finally, we control for the number of friends to make sure that we compare subscribers with similar social circles. Otherwise, it would be unreasonable to compare subscribers that can receive many signals (friends churning) from many friends with subscribers with very few friends can only obtain very limited signals.

We consider 3 levels of treatment intensity: i) T1: at most one friend churns; ii) T2: 2 or 3 friends churn; iii) T3: more than 3 friends churn. We compute a propensity score for each panel and for each treatment intensity. Then we investigate the balancing property for each controlled covariate, by testing whether the conditional means of the subscriber's covariates given the propensity score are different for subscribers with each treatment intensity and subscribers with other treatment intensities. If the latter are similar then subscribers with different treatment intensities are not different in other aspects of their behavior, which allows us to better associate changes in the outcome (the ego's churn) to changes in the treatment (friends' churn). Again, we choose the treatment as $n\text{-}callfrd\_churn$ for $n = 1, 3, 5$ to denote different levels of tie strength. The implementation procedures and balancing results for GPS-adjusted covariates are detailed in Appendix A1.3. Generally, adjustment for the GPS improve the balance of each covariate significantly. More importantly, matching on panels allows us to ensure subscribers with divergent trends are not comparable, whereas matching on averaged covariates may ignore such property.

2.5.3 Empirical Results

We use GPS to estimate the effect of treatment, friends' churn, on the outcome, the ego's churn, as the output of dose response function. We do not report the regression results, because the estimated coefficients have no direct meaning (Hirano and Imbens, 2005). We estimate the dose response function relative to having no friends that churn and we report results for up to five friends churning, which covers 99.9% of the observations in our dataset.

Figure 2.2 shows the results obtained for $n = 1, 3, 5$. We observe that having more friends churn increases the likelihood of churn for any n considered in our analysis. Also, we see that when considering the churn from the marginal effect of treatment, that is the effect of having one more of these friends churn, remains nearly positive with the number of friends that churn. This provides evidence of peer influence in churn in wireless networks. Furthermore, for treatment intensity T3 (that is, more than 3 friends churn), the churn likelihood for 5-call increases well beyond the 1-call and 3-call, which provides some evidence that churn from stronger friends might be more important. This is a sensible result showing that enough strong friends churning makes a significant difference on the ego's probability to churn when enough friends churn.

Finally, note that the treatment effect from the GPS analysis is not directly comparable to the relative hazard from the survival analysis. Yet, we can see that the effect of having five strong friends churning is barely over 10% in the GPS results. This means that homophily is likely to play a significant role in the correlation of churn across friends that the survival analysis confounds with peer influence. GPS reduces the bias in estimating the effect of peer influence by comparing across similar subscribers. This takes away the effect of all possible homophily captured by the covariates that we control for in computing the propensity score.

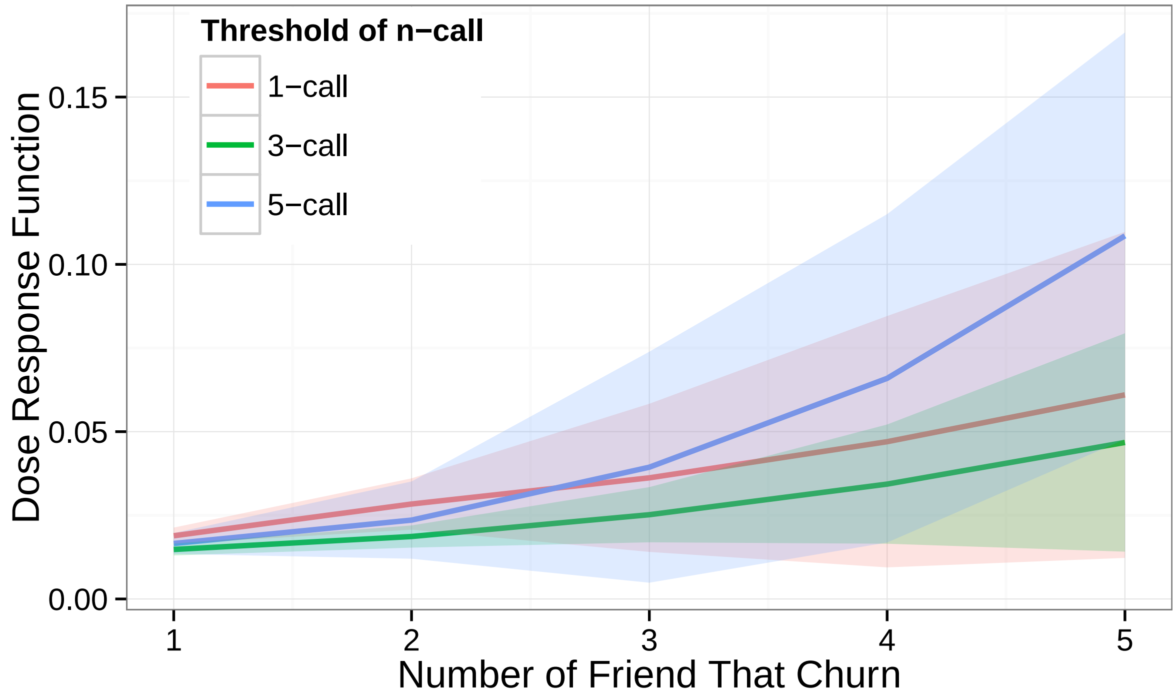

Figure 2.2: Estimated dose response function with PTP:{1,2, 3}, TEP:{4,5}, COP:{6,7} relative to having no friends churning. Ribbons represent the 95% confidence intervals. Standard errors are obtained via bootstrapping (100 repetitions)

2.5.4 Robustness Check

As shown in the preceding section, GPS method may account for homophily on observed characteristics and those unobserved characteristics that are correlated with observed ones, such that selection bias on observables can be eliminated. However, we are still careful to establish full causal interpretation of our results, because latent characteristics may contribute to both the treatment assignment and outcome. Therefore, we perform the GPS analysis using alternative functional forms of treatment to check the robustness of the estimated peer influence. Specifically, we choose the fraction of friends that churn as the treatment. The rationale rests upon the classical threshold model proposed by (Granovetter, 1978) and then the cascade model by (Kempe et al., 2003), that subscriber churn occurs when a particular threshold i.e., fraction of friends churn is exceeded. With the same amount of cumulative exposure of friends churn, the outcome of ego churn may vary, because each subscriber selects her own threshold to be "activated" (Dasgupta et al., 2008). Guardabascio and Ventura (2014) described a flogit estimator for GPS analysis using fractional treatment and mechanics about estimation techniques are detailed in Appendix A.1.4.

Figure 2.3 shows the results obtained using fraction of friends churn as the treatment (99% observations have up to 30% of friends churn and 99.9% observations have up to 50% of friends churn). Similar to previous results, subscribers with higher fraction of friends churn, regardless of tie strength, are more inclined to churn. In particular, the effect of 5-$call$ friends churn is stronger than both 1-$call$ and 3-$call$ friends churn, which is not so clearly when we simply count the number. This implies that subscribers typically have fewer strong friends and thus have low threshold level to churn. To the contrary, the effects of weak and intermediate friends churn are very close, indicating that subscribers have similar threshold level to churn with this definition of tie strength.

2.6 Discussion

Retaining existing subscribers is of vital importance for wireless carriers to survive in today's dynamic and competitive mobile market. As such, understanding the determinants of churn becomes a priority. Carriers need to confidently identify potential churners to apply appropriate retention strategies aimed at reducing subscriber loss. However, the perplexing and evolving nature of churn still poses significant challenges to churn managers. In this paper, we look

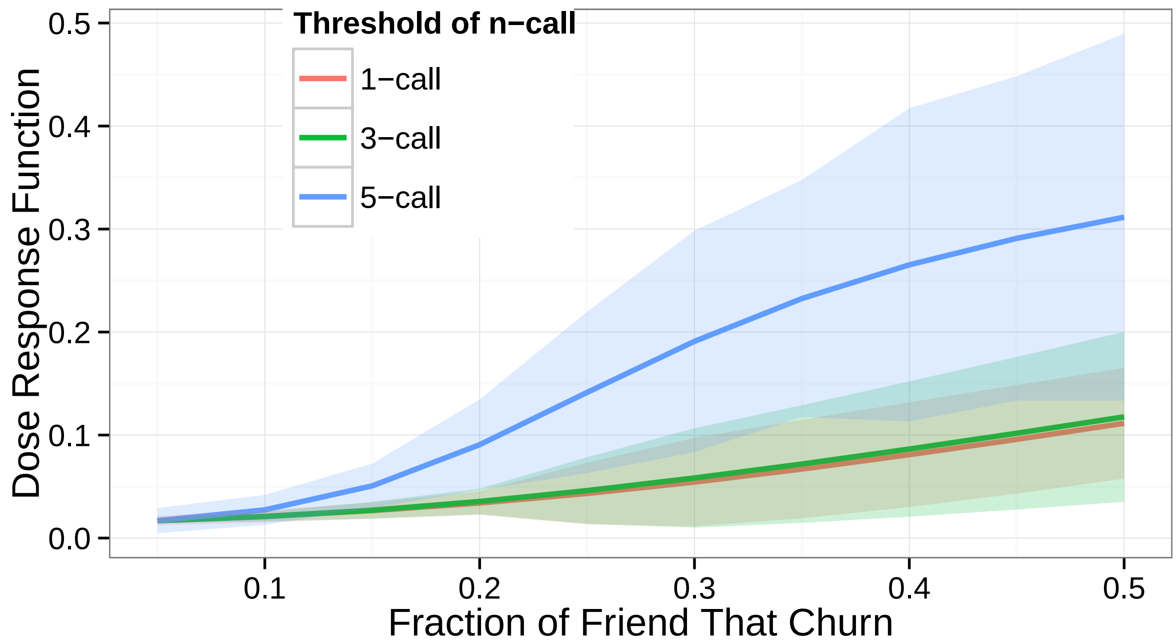

Figure 2.3: Estimated dose response function with PTP:{1,2, 3}, TEP:{4,5}, COP:{6,7} relative to having no friends churning. Ribbons represent the 95% confidence intervals. Standard errors are obtained via bootstrapping (100 repetitions)

at one of such complexities. We examine the effect of peer influence on churn. We do so empirically using a real world dataset of cell phone activity, from which we extract a set of covariates measuring cell phone usage, tenure with the carrier and network structural properties such as number of friends. We analyze a random sample of roughly 10 thousand subscribers. For each subscriber in our sample we define the number of friends that churn (and later the fraction of friends for robustness check) with whom the subscriber exchanges 1, 3 or 5 calls in the same calendar month. These different definitions allow for estimating the effect of churn from strong vs. weak friends.

In a preliminary analysis, we fit survival models to our data to correlate the cumulative number of friends that churned up to the previous month and the ego's likelihood of churn in the current month. An additional friend that churns increases the hazard of churn in at least 24%.

In a deeper analysis, we use Generalized Propensity Score (GPS) matching to help reduce the potential selection bias across subscribers in our sample and estimate the effect of contagion. With GPS we can compare subscribers that are similar in a number of relevant covariates and differ only in the intensity of their treatment. The latter is the number of friends that churned. We control for the cell phone usage, monthly expenses, the size of the social network and subscribers tenure with the carrier. These are the relevant covariates to capture the subscribers behavior with respect to cell phone service and choice of provider. Due to the panel nature of our data, we perform the matching at the panel level rather than individual level, because we intend to include the dynamic subscriber behavior. In practice, we do so by controlling for the lagged value of cell phone usage and monthly expenses. In this framework, we associate differences in the intensity of treatment to differences in the outcome of interest. In our case, the latter is the ego's propensity to churn.

We find that the ego's propensity to churn increases with the number of friends that churn and in particular with increased numbers of strong friends churning. With GPS, we find that the cumulative effect of up to 5 friend churning is at best 10%. To check the robustness, we also adopt the fraction of friends churn as the treatment, and find similar results. In particular, the peer effects that arise from strong friends churn are higher compared with weak friends churn, which is very indicative of role of tie strength in moderating peer influence. With GPS we show that the role of peer influence is still significant in wireless networks: when someone churns from the our carrier, other people may do so due to peer influence. This result, which may well represent a loss of 10% in revenues for the carrier due to peer influence in churn, suggests that churn managers should consider strategies aimed at preventing group churn instead of looking at subscribers on an individual basis. Although the results is not directly comparable between survival analysis and GPS approach, our results may still imply that homophily is likely to drive a significant amount of the correlation between churn and friends' churn that a simple survival analysis is unable to disentangle.

However, our method and results are not without limitations. We may still fail to capture all possible homophily that we do not include in our analysis. Unobserved reasons that explain both the churn and friends churn may still play a role here and GPS is inherently not able to deal with this concern. After all, randomized experiments is ideal to identify the peer influence on churn. For example, Matos et al. (2015) performed proactive retention strategies by targeting on group of connected subscribers and they showed that churn probability is reduced and CLV is increased. Also, with only 7 months of the panel, we have limited option to split the time periods into pretreatment, treatment exposure and outcome observation, so that each period can be sufficiently large. With the alternative definition of treatment, the results combined may provide evidence of robustness.

Moreover, we also note several limitations that caution us to interpret the results in more general settings. As mobile usage patterns have drastically changed over the past years due to the technological advances in smartphones (Han and Cho, 2016), using mutual call communications to identify the social relationship between subscribers may not be able to fully reflect the latest facts. For example, subscribers increasingly resort to various messenger apps to exchange communications and thus make the mobile social network built upon call graph less relevant. Meanwhile, as our analysis only focus on the prepaid consumer segment, we do not expect our results would equally apply to postpaid consumers who clearly face different decision-making process due to the contractual obligations without further analysis. Thus more empirical studies that can leverage the dataset of up-to-date mobile phone usage to study the subscriber churn should be performed, as it still remains as the topmost challenge in wireless industry.

Finally, we stress that our work puts together a methodology to measure peer influence in social networks that can be used in other contexts such as learning about the effect of word-of-mouth in the dissemination of new products or services.

2.7 Appendix

2.7.1 GPS model for contagious churn

Model Setup

Formally, following (Hirano and Imbens, 2005), we consider a set of N subscribers and let i denote a single subscriber. Let $P = \{1, \ldots p\}$ represent a set of time periods. We observe a vector of covariates X_{ip} at each time period. We define the treatment at each period for each subscriber as the number of churner friends in the last time period t_{ip-1}. Therefore, we decided to take value of treatment in the unit interval: $T \in \{0, 1, 2, 3, \ldots\}$, to indicate 0, 1, 2 or 3 or more friends' churn, respectively. The outcome of interest is whether subscriber i churns in time period p: $Y_{ip} \in \{0, 1\}$.

GPS method requires weak unconfoundedness assumption to hold: $Y(t) \perp\!\!\!\perp T|X$. This means potential outcomes are conditionally independent for each value of the treatment. Meanwhile, as explained in the section 5.1, we estimate the panel level GPS rather than individual observations, i.e. $X = \{x_1, \ldots, x_p\}$. Dropping the index term for simplicity, propensity function $r(t, x)$ is defined as the conditional density of treatment given the covariates $X = x$ evaluated at $T = t$:

$$r(t, x) = f_{T|X}(t|x) \tag{2.2}$$

Then GPS is just the corresponding random variable: $R = r(T, X)$. Note that GPS has similar balancing property to the propensity score for binary treatment. That is, within strata with the same value of $r(t, X)$, T and X are conditionally independent:

$$X \perp\!\!\!\perp 1\{T = t\}|r(t, x) \tag{2.3}$$

Together with the weak unconfoundedness, Hirano and Imbens (2005) proved that the assignment to treatment is also unconfounded given GPS. Thus we can use GPS to remove the bias associated with differences in the covariates. Similar to the property of dimentionality reduction to classical propensity score, we denote the dose response function as the average of a set of potential outcomes given the treatment level t: $\{E\big[Y(t)\big]_{t\in T}\}$ where T is the set of potential treatment values. Then the conditional expectation of churn is a function of two scalar variables: number of churner friends T and of the GPS R:

$$\lambda(t, r) = E[Y(t)|r(t, X) = r] = E[Y|T = t, R = r] \tag{2.4}$$

Therefore, the dose response function of churn is the average conditional expectation over

GPS given number of churner friends:

$$\mu(t) = E[Y(t)] = E[\lambda(t, r(t, X))] \tag{2.5}$$

Estimation

Implementing GPS method is far from trivial. In practice, there are three functions needed to be estimated: (i) the score function $r(t, x)$; (ii) the conditional expectation of outcomes given treatment and GPS $\lambda(t, r)$; (iii) the dose response function $\mu(t)$. Hirano and Imbens (2005) assumes the treatment conditional on covariates follows the normal distribution, so that the maximum likelihood can be used to estimate the GPS. However in our case, due to the count nature of the frd_churn, the distribution of treatment may be explicitly non-normal. Therefore, we follow the parametric solution proposed by Guardabascio and Ventura (2014), to replace the linear regression by the generalized linear model (GLM) and estimate the GPS from the exponential family function. Namely, the distribution of T is selected from exponential family:

$$f(T) = c(T, \phi)\text{exp}\Big\{\frac{T\theta - a(\theta)}{\phi}\Big\} \tag{2.6}$$

Then the transformation of the mean of the treatment is represented as a link function $g(\cdot)$ that linearly relates to the covariates:

$$g\big[E(T)\big] = X\beta \tag{2.7}$$

Having chosen the family and link option for distribution of T from exponential family, we can estimate θ, ϕ and β using quasi-maximum log likelihood from equation (5) and (6).

$$l(\beta) = \sum \log f(T;\beta) = \sum \Big\{ \log c(T,\hat{\phi}) + \frac{T\hat{\theta} - a(\hat{\theta})}{\hat{\phi}} \Big\} \tag{2.8}$$

After that we compute the GPS at the value of $\hat{\beta}$, given the covariates:

$$\hat{R}_i = r(t,x) = f(\hat{\beta}) = c(T,\hat{\phi})\text{exp}\Big\{ \frac{T\hat{\theta} - a(\hat{\theta})}{\hat{\phi}} \Big\} \tag{2.9}$$

If balancing property is satisfied after adjusting for GPS, we can estimate $\lambda(t,r)$ through a function $\varphi(\cdot)$, by taking number of churner friends T_i, and the GPS R_i as arguments of the function on subscribers' churn outcome Y_i. In practice, we use a polynomial approximation of order two to regress Y_i on T_i and R_i:

$$\begin{aligned} \varphi\big(\lambda(T_i,R_i)\big) &= \varphi\big(E(Y_i|T_i,R_i)\big) = \psi(T_i,R_i;\alpha) \\ &= \alpha_0 + \alpha_1 T_i + \alpha_2 T_i^2 + \alpha_3 R_i + \alpha_4 R_i^2 + \alpha_5 T_i R_i \end{aligned} \tag{2.10}$$

In the last step, to derive the dose response function $\mu(t)$, we plug the estimated score function $\hat{r}(t,x)$ into equation (4) and average the estimated conditional expectation over the GPS evaluated at each level of treatment:

$$\begin{aligned}\hat{\mu}(t) &= \frac{1}{N}\sum_{i=1}^{N}\hat{\lambda}(t,\hat{r}(t,X_i)) \\ &= \frac{1}{N}\sum_{i=1}^{N}\varphi^{-1}\big(\hat{\psi}(t,\hat{r}(t,X_i);\hat{\alpha})\big) \qquad (2.11)\end{aligned}$$

In our case, the marginal treatment effect can easily be derived as the difference between the outcome at a particular level of frd_churn compared to the outcome of a lower reference level. More specifically, having l more churner friend will cause subscriber's propensity to churn to change by $\hat{\mu}(t+l) - \hat{\mu}(t)$.

2.7.2 Balancing Covariates

As explained in Hirano and Imbens (2005), we divide the sample into 3 groups according to the distribution of treatment intensity, compute the GPS for individuals in each group and further divide the obtained GPS into 5 blocks. Within each block, we calculate the mean difference of each covariate between subscribers who belong to one group and subscribers who belong to other groups. Specifically, we compare subscribers who are treated with subscribers who are in the same block but not treated at the same level, and calculate the weighted average of the differences, with weights given by the number of observations in each GPS block.

Table 3-5 shows how the adjustment by conditioning on the propensity score balances covariates listed in section 5.2 for n-$callfrd\_churn$ and $n = 1, 3, 5$. We can see that most covariates are different before adjustment but become statistically similar at the 5% level after adjustment.

2.7.3 Fractional Treatment

As specified in section 3.1 of (Guardabascio and Ventura, 2014), when we define the fraction of friends that churn as the treatment, i.e., $T \in [0, 1]$, following (Papke and Wooldridge, 1996), we specify the functional forms for $E(T_i|X_i)$ for all i as:

$$E(T_i|X_i) = F(\beta' X_i) \tag{2.12}$$

where $F(\cdot)$ is typically a logit or probit function. In practice, we choose a binomial family and logit link function for GLM estimator and thus the Bernoulli log-likelihood function to feed into equation (8) is:

$$l_i(\beta) = T_i \log\{F(\beta' X_i)\} + (1 - T_i) \log\{1 - F(\beta' X_i)\} \tag{2.13}$$

After estimating $\hat{\beta}$ by maximizing the sum of $l_i(\beta)$ over all N, the GPS is computed exactly in the same way as treatment is continuous.

Covariates	Before adjustment			After adjustment		
	[T1]	[T2]	[T3]	[T1]	[T2]	[T3]
n_calls month -1	**-6.07**	**4.85**	**3.66**	-0.99	-0.16	0.85
n_calls month -2	**-5.50**	**4.70**	**2.19**	-0.22	0.13	-0.02
n_calls month -3	**-5.80**	**4.98**	**2.77**	-1.56	0.65	1.43
$expense$ month -1	**-4.92**	**3.57**	**3.65**	1.67	-1.36	-1.37
$expense$ month -2	**-4.65**	**3.38**	**3.45**	1.82	-1.92	-1.61
$expense$ month -3	**-4.88**	**3.70**	**3.28**	1.21	-1.47	**-2.33**
$\%call\_other$ month -1	**2.27**	-1.42	**-2.22**	-0.71	0.40	0.12
$\%call\_other$ month -2	**2.60**	-1.75	**-2.28**	-1.29	1.13	1.17
$\%call\_other$ month -3	**1.99**	-1.23	**-2.02**	-0.60	0.84	0.96
frd	**-14.01**	**9.86**	**10.60**	**-2.39**	**-3.99**	**2.03**
$tenure$	**6.62**	**-4.75**	**-5.20**	1.35	1.86	0.75

Table 2.3: Balance in subscribers' covariates for 1-$call$ frd_churn introduced by conditioning on the generalized propensity score. **Bold values** indicate covariates that are statistically different (at the 5% level, i.e., t-stat is larger than 1.96) between that treatment intensity and the other treatment intensities.

Covariates	Before adjustment			After adjustment		
	[T1]	[T2]	[T3]	[T1]	[T2]	[T3]
n_calls month -1	**-5.96**	**4.75**	**5.85**	1.31	-1.09	-1.69
n_calls month -2	**-4.54**	**5.19**	**4.48**	1.21	-1.43	-1.23
n_calls month -3	**-4.16**	**5.42**	**4.11**	1.12	-1.75	-1.42
$expense$ month -1	**-5.21**	**3.85**	**5.25**	1.19	-1.60	-0.68
$expense$ month -2	**-5.11**	**4.09**	**5.17**	1.50	-1.67	-0.70
$expense$ month -3	**-4.77**	**4.10**	**4.81**	1.95	-1.62	**-2.93**
$\%call\_other$ month -1	**4.16**	-1.79	**-4.16**	-0.95	0.68	-0.66
$\%call\_other$ month -2	**4.36**	-1.85	**-4.37**	-0.96	1.59	-0.94
$\%call\_other$ month -3	**3.83**	**-2.34**	**-3.85**	-0.05	0.27	0.01
frd	**-13.17**	**10.97**	**13.14**	**2.88**	**-3.08**	**-3.61**
$tenure$	**6.61**	**-4.96**	**-7.01**	0.74	1.32	1.21

Table 2.4: Balance in subscribers' covariates for 3-$call$ frd_churn introduced by conditioning on the generalized propensity score. **Bold values** indicate covariates that are statistically different (at the 5% level, i.e., t-stat is larger than 1.96) between that treatment intensity and the other treatment intensities.

Covariates	Before adjustment			After adjustment		
	[T1]	[T2]	[T3]	[T1]	[T2]	[T3]
n_calls month -1	**-6.33**	**5.11**	**3.75**	0.80	-1.08	-1.52
n_calls month -2	**-5.48**	**4.64**	**2.77**	**2.00**	-1.24	-1.75
n_calls month -3	**-5.46**	**5.00**	**2.84**	0.48	-1.09	-0.95
$expense$ month -1	**-4.11**	**3.62**	**3.91**	1.38	-1.82	-1.45
$expense$ month -2	**-3.63**	**3.30**	**3.05**	1.92	-1.91	-1.67
$expense$ month -3	**-3.92**	**3.53**	**3.40**	1.68	**-2.08**	**-2.47**
$\%call\_other$ month -1	**3.56**	**-3.24**	-1.39	0.75	0.42	0.78
$\%call\_other$ month -2	**3.73**	**-3.61**	-0.99	0.67	0.09	0.53
$\%call\_other$ month -3	**2.85**	**-3.06**	-0.35	1.37	-1.57	0.10
frd	**-11.83**	**10.02**	**5.94**	**2.38**	**-2.50**	**-2.17**
$tenure$	**5.83**	**-5.52**	**-3.17**	-1.57	1.11	1.56

Table 2.5: Balance in subscribers' covariates for 5-$call$ frd_churn introduced by conditioning on the generalized propensity score. **Bold values** indicate covariates that are statistically different (at the 5% level, i.e., t-stat is larger than 1.96) between that treatment intensity and the other treatment intensities.

Chapter 3

The Effect of Product Placement Recommendations on Shopping Behavior: Evidence from Randomized Experiment in a Physical Bookstore

3.1 Introduction

Physical retailers nowadays are struggling to compete with online retailers in improving both the sales and customer satisfaction. One of the many challenges facing physical retailers is how to optimize the allocation of valuable shelf space to accommodate the increasing product variety (Dreze et al., 1994). Typically, physical retailers only display a carefully selected subset of items on store shelves that appeal to consumer preferences due to the fundamental shelf space constraints. Even so, consumers choice may be limited by search costs over all of available items in the inventory during the shopping process, as this requires significant time commitment and cognitive load. As such, shoppers may have difficulty in finding items that they like and leave the store without buying anything, even when retailers carry these products in the store. Therefore, physical retailers seek to have effective shelf management strategy, *e.g.* through adjusting the product assortment and placement in ways that raise shopper's awareness and thus induce their purchase decision.

Moreover, many studies have documented that the majority of consumers make unplanned purchase decisions once they enter the store (Gilbride et al., 2015; Sam Hui and Suher, 2013; Inman et al., 2009; Stilley et al., 2010). Court et al. (2009) report up to 40 percent of customers change their minds due to the visual-dimension factors at an in-store touch point: packaging, position, signage, etc. For retailers, this offers substantial opportunities for the planning and execution of effective in-store merchandising practices to stimulate incremental revenues and profits by converting demand to purchase. To this end, retailers are increasingly shifting their focus from traditional marketing to shopper marketing, i.e., to engage with consumers at each touch point and trigger their intention to purchase throughout the shopping cycle, which comprises different stages such as search, evaluation, item choice, purchase, and post-purchase, and so on (Shankar et al., 2011). In particular, understanding what happens at the point of purchase is rapidly growing in importance, as it may be informative about the final stage of consumer buying decision process, e.g., product consideration, termed by Procter & Gamble as the "first moment of truth" (Lofgren, 2005; Hui et al., 2013).

Understanding shopper's decision-making process in physical retailing is still challenging because of measurement and analytics hurdles. Typically, retailers and marketers alike manipulate various marketing mix factors such as pricing, promotion, and product position, etc., and then examine the resulting impact on sales (Gaur and Fisher, 2005; Nordfalt et al., 2014; Valdimar Sigurdsson and Foxall, 2009). However, retailers clearly need to further investigate causal drivers along the entire path to purchase that goes well beyond survey and scanner data (Shankar et al., 2011). Although continuing innovations in retail environment have emerged to capture metrics that are indicative of shopper behavior (e.g. Shankar et al., 2011), there is a dearth of research that performs the controlled experiment while accurately recording shopper actions inside the store (Burke, 2006; Shankar et al., 2011).

In this chapter, we implemented an *in vivo* experiment that ensures random product placement in a physical bookstore. As product placement can be considered as the display recommendation mechanism that may affect consumers' purchase behavior (Seiler, 2013), we leveraged video tracking technologies to monitor how shoppers respond to random book placement, which induces random search costs. More specifically, we randomized the position of newly released books on the top of a large display table across several rows and columns, such that each book's search cost becomes independent of the book's characteristic, such as unobserved quality. As such, we are able to assign incoming books into a treatment group, in which books are placed at the edge of the table, and a control group, in which books are placed at the center of the table. To overcome the large costs associated to collecting and processing large-scale video data, we used advanced 3D cameras and vision-understanding algorithms that can track human motions in real-time. This allows us to significantly reduce the cost of encoding shopper activities.

We used in total 90 news books are used in the 6-week experiment and our randomization procedure proved to well the books in treatment and control group. We capture shopper behavior that is indicative of decision process, in terms of she picking a book after browsing over the display table and taking the book after reading through its content. Regression results show that books in the treatment group are 102% more often to be picked and 77% more often to be taken than those in the control group. This indicates that book placement positively affect the shopper's search and consideration process. Moreover, we further evaluate whether shoppers tend to take the book conditional on they have picked the book. We find that books in the treatment group are equally likely to be taken, compared to those in the control group. This indicates that book placement affects shopper's intention to purchase only through the search process.

This chapter is structured as follows. Chapter 3.2 provides extensive literature review on product placement and shopper's path to purchase in both online and offline setting. Chapter 3.3 explains the research context that our experiment is based on. Chapter 3.4 provide the experiment design details. Chapter 3.5 explains data collection process using video tracking technology. Chapter 3.6 lists the econometric model specifications at both book level and visit level. Chapter 3.7 presents empirical results with the data we have collected through the experiment. Chapter 3.8 provides discussion of our results and Chapter 3.9 proposes the future work.

3.2 Literature Review

Given its interdisciplinary nature, this chapter potentially combines several streams of research together. First, our work adds to the marketing literature on how technological changes transform retail industry. Second, our work is closely related how changes in product placement affects retail performance. Third, more broadly, our work is within the scope of understanding the role of product placement in influencing consumer's path to purchase in both online and offline channels.

3.2.1 Technological Change in Physical Retailing

In general, Burke (2006) explained the evolution of marketing intelligence in retail industry in terms of three waves during the past decades. The first wave occurs when point-of-sale (POS) systems and UPC barcode scanners are in widespread use. Retailers can acquire real time transactional data that allows them to estimate the product sales and market share. Since the pioneering work by (Guadagni and Little, 1983), both academic researchers and industry practitioners heavily rely on the scanner panel data for supporting business decisions, such as category management, shelf space allocation and inventory management (see (e.g. Bucklin and Gupt, 1999) for a overview of the use of scanner data from both industry and academic perspectives). However, scanner data are typically aggregated at either brand/category or store/firm level, which lack interpretability of the individual choices (Chen and Yang, 2007).

The second wave occurs when retailers launch the customer loyalty program to log transactions at individual level. As customers present loyalty cards at the point of sale in exchange for discounts and rewards, retailers can recognize customer's repeat purchase and further integrate with external data sources (e.g. geo-demographic characteristics and credit card information) to create a profile and purchase history for each customer/household (see (e.g. Kumar and Shah, 2004) for a review on common practices using loyalty card in retail industry). Together with sales data collected from UPC scanning, retailer managers can analyze customer purchase patterns using data mining techniques to derive customer lifetime value and formulate personalized marketing activities (Linoff and Berry, 2011). However, loyalty program data reveals little information on decision-making or customer experience (Uncles et al., 2003). For example, customers who do not make any purchase will not be recorded in the database, thus retailers who feel compelled to strengthen customer relationship would have limited evidence on customer loyalty and retention.

Court et al. (2009) described the shopper's decision journey as four major circular phases that during each phase marketers can engage with shoppers and impose influence: initial consideration, active evaluation, purchase and post-purchase. Thus the third wave of changes enables the real-time tracking of customer's actual shopping activities in different phases, which is traditionally difficult because researchers have to hire research assistants to physically follow shoppers in the store (e.g. Farley and Ring, 1966; Granbois, 1968). More recently, technology innovations in both computer hardware and software have emerged to address shopper's in-store activities throughout the shopping cycle (Shankar et al., 2011), including RFID path tracking (e.g. Larson et al., 2005; Hui et al., 2009; Hui and Bradlow, 2012; Sam Hui and Suher, 2013), eye-tracking (e.g. Chandon et al., 2009; Wedel and Pieters, 2008), wearable video tracking device (Hui et al., 2013), surveillance video (Zhang et al., 2014).

Both Larson et al. (2005) and Hui and Bradlow (2012) used the shopping path data collected from RFID tags affixed under the grocery carts and find that most shoppers only visit selected aisles of the supermarket during the trip. Similarly, Hui et al. (2009) showed that shoppers tend to deviate from the most efficient shortest path connecting all their purchases. Furthermore, eye-tracking technology allows marketers to measure customer's visual attention through eye-movement data. Wedel and Pieters (2008) reviewed the extensive eye-tracking studies in marketing literature and show the value of visual marketing in practice. For example, both Chandon et al. (2007) and Chandon et al. (2009) empirically examined the customer's visual attentions and considerations towards shelf displays in retail setting, measured as the position and duration of eye-fixation. However, neither RFID-based path data nor eye-tracking data can capture the phase of active evaluation and product accessibility, e.g., trial and touch that is critical to purchase conversion (Underhill, 2009, section iv).

Video tracking provides a holistic understanding of the shopping experience through recording how customers shop in the store. Burke (2006) used video data from surveillance cameras mounted in consumer electronics stores during the holiday shopping season. Hui et al. (2013) recruited shoppers to wear portable camera devices that can follow their range of vision and shopping path. Zhang et al. (2014) collected data from a panoramic video camera in a retail apparel store and track shopper movement and crowding conditions. In general, video tracking solutions not only help marketers to collect direct observations of in-store shopper behavioral patterns, but also identify important opportunities in managing shopping process. Thus the next step would be to validate the findings from observational studies and estimate the interaction effects of environmental drivers on shopper behavior through experimental approach (Shankar et al., 2011).

However, none of the existing research is suitable to support the planning of in-vivo experiment in real world retail setting for the following reasons. First, either eye-tracking or portable video tracking device may require participant's attention and this can distort the experiment results due to the "Hawthorne effect" (Adair, 1984). Second, product consideration and evaluation, a key prerequisite to purchase is not applicable to be measured using RFID-based path tracking. Third, although video-tracking solutions are seemingly less intrusive and more comprehensive to capture shopper behavior, manually extracting information from videos has been noted to incur prohibitively high costs. Hui et al. (2013) reported that encoding each hour of video data took more than four hours. Zhang et al. (2014) used only three hours of coded video data, which should not be sufficient for typical controlled experiments lasting weeks or even months at a time. Fourth, privacy concerns still prevail when retailers consider performing field experiment in which shopper identities may be revealed (Burke, 2006). Therefore, researchers continue to call for technological innovations to resolve these issues, e.g. vision understanding techniques to lessen the efforts of the video encoding process.

3.2.2 Product Placement and Retail Performance

Providing the product display at a place for consumers to conveniently access has long been acknowledged as an important component of marketing mixing instruments in marketing literature (McCarthy, 1964). In general, research on the allocation of retail products to shelf space can be decomposed into two dimensions: 1) amount of space measured as number of shelf facings, and 2) spatial shelf layout measured as vertical and horizontal shelf positions (Frank and Massy, 1970; Dreze et al., 1994). On one hand, studies on the shelf space have mostly shown consistent findings: increasing the number of facings has a positive but marginally diminishing effect on store sales (Cox, 1964; Curhan, 1972; Dreze et al., 1994). On the other hand, how location of products impacts on retail performance remains inconclusive. Although early work that considers only vertical shelf levels found very modest effect (Frank and Massy, 1970), more recent studies that took into account of both vertical and horizontal dimensions showed that products placed between slightly above the eye level and hand level achieve better sales (Dreze et al., 1994; Philips and Bradshaw, 1993; Valdimar Sigurdsson and Foxall, 2009; Underhill, 2009). Moreover, retailing literature supports mixed evidences about the best position on the horizontal dimension, such as products at the end-of-aisle display (also known as endcaps) (Underhill, 2009, p.85), products placed in the middle shelf (Christenfeld, 1995), or contingent on the product category (Dreze et al., 1994).

After all, the saliency of product placement on retail performance presses both manufacturers and retailers to address the shelf management challenges. Manufacturers compete for scarce shelf space and favorable shelf placement and are even willing to pay for the slotting allowance (e.g. Klein and Wright, 2007). Retailers endeavor to design the optimal shelf space allocation paradigms. These paradigms typically involve mathematical models with number of facings, vertical and horizontal position as decision parameters, and total shelf space and space for each aisle as constraints and then solve for combinatorial optimization problem (e.g. Yang and Chen, 1999; Yang, 2001; Lim et al., 2004). Recent shelf-allocation models further incorporate elements from inventory levels (Hwang et al., 2005), marketing mix instruments (Murray et al., 2010; van Nierop et al., 2008), product assortment within the category (Russell and Urban, 2010) and cross-category (Bezawada et al., 2009), among others.

Still, marketers eagerly aim at understanding the underlying mechanisms through which shoppers interact with in-store shelf layout and tend to choose retail products with position advantages. Existing literatures proposed two competing theories explaining such shopper preferences. On one hand, researchers suggested that position and number of facings affect the visibility and accessibility of products, such that favorable position and more facings would incur lower search cost to draw shopper's attention. Numerous studies employing eye-tracking techniques support the idea of "unseen is unsold". Chandon et al. (2007) found that products located near the center receive more attention. Wedel and Pieters (2008) found that product display size strongly influences the visual attention. Using experiment approach that randomize the position and display size of products, Chandon et al. (2009) also revealed that products placed at top and middle shelf position gain more attention than those at low-shelf positions, and horizontal position generally makes no difference to attention except the position that is near the center. Moreover, they further noted that only number of facings, vertical and horizontally center position have positive effect on evaluation through the attention. Atalay et al. (2012) explained that products in the horizontal center positions are more likely to be chosen because human tend to look first at the center in terms of eye-fixation (centrality bias in visual attention), and then this would reinforce the gaze cascade effect later which is related to the choice.

On the other hand, consumer psychology literature argues that position effects are not mediated by attention but rather by inferred quality, such that favorable positions and more facings are believed to provide shoppers with higher expected utility. Several studies using lab experiments found that consumers perceived that i) products in central position are more preferred over those at either end of the array (Valenzuela and Raghubir, 2009); ii) retailers place premium products on the top shelves and popular products on the middle shelves (Valenzuela et al., 2013); iii) products placed on the right hand side tend to have better quality compared to products placed on the left side (Valenzuela and Raghubir, 2015). Nordfalt et al. (2014) performed experiments in the retail stores by placing the same new shampoo on different levels of shelf, and survey results showed that customers perceived the shampoo in the middle level to be the most expensive, followed by upper level and then bottom level.

3.2.3 Product Placement and Path to Purchase

Hui et al. (2009) argued that although from seemingly unrelated marketing domains, how shoppers interact with their shopping environment and make dynamic decisions, regardless of online or offline, can essentially be modeled as the same type of path data. As data collection technologies evolve, path data increasingly plays a central role in marketing research. For example, consumer's web browsing patterns are tracked as sequences of the page viewed when they navigate through an E-Commerce website(Montgomery et al., 2004). Also, as explained earlier, RFID-enabled shopping path and visual movement using eye tracking are two examples of path data in retail environment. However, data that capture the shopper's in-store consideration and evaluation that lead to their choices are still very limited, inasmuch as this would be informative about their decision process and goal orientations. For example, one carrying a shopping list may behave quite differently compared to a casual shopper, even though the two may purchase the same product.

Branco et al. (2012) presented a theoretical framework that models the search stage before making purchasing decisions. In general, consumers tend to gather information about product by examining the product, then update their beliefs on how much they value the product and up to a point they decide whether to purchase the product or not. In this model, both search cost and valuation represent as important factors in the search process that lead to the purchase. More specifically, consumer's valuation of the product is the sum of the utility from all the information she have searched, and search cost moderates the extent of the search process, so as to influence the purchase likelihood. This implies that the retailers can influence the consumer's decision to purchase by changing the search cost, contingent on the expected utility of products from consumers.

When consumers search for the product information in online channels, they are often presented the resulting list of ranked potential choices. Generally, position effects represent the search cost well. For example, the position effect of search on mobile phones is more significant compared to on personal computers, simply because mobile phones have smaller screen size and thus increase the search cost (Ghose et al., 2013). Studies of position effect on consumer choices have been prevalent in a myriad of online applications, such as shopbots (Brynjolfsson et al., 2010), sponsored advertisement (Agarwal et al., 2011), search engines (Ghose et al., 2014), reviews (Matos et al., 2016), etc. and consistent results are reported as consumers are more likely to click and purchase from the products at the top position. However, the position examined in previous studies is endogenous, chosen by e-commerce retailers to maximize their profits. Ursu (2016) used a unique dataset from an online travel agent, that consumers observed hotels positioned randomly, such that the hotel's characteristics are not related to the display layout. She found that although top positioned hotels still receive more considerations in terms of clicks, these hotels conditional on a click, the purchase likelihood remain constant across positions. This implies that product placement affects only the search cost during the search stage, but hardly carry over to affect consumer's expected utility during the valuation stage.

3.3 Research Context

We partner with an independent bookstore located in the center of a European capital city (referred to as IndieBookstore hereafter). IndieBookstore has a 60-year history of bookselling and is one of the largest independent bookstores in the country. It operates 7 days a week from 9AM to 11PM between Monday and Saturday and from 9AM to 7PM during Sundays and holidays. The store has two floors (ground floor and basement), with sales areas over 200 m^2 at each floor. The ground floor area mainly sells books, magazines, music CDs and the basement area mainly sells children's books and school supplies.

The store layout of IndieBookstore is designed to stock and display newly released books (also known as *frontlist books*) and older titles (also known as *backlist books*) separately, as shown in Figure 3.1. More specifically, IndieBookstore receives frontlist books from publishers and places them face-out on the top of several large tables in the center of the store, and places backlist books spine-out on the surrounding wall shelves. In particular, publishers deliver the latest released titles to the IndieBookstore a few days before the publication date and all incoming books are first placed on the front most table near the entrance on the publication date. Later, these books are replaced by the most recent released books and redistributed to other tables according to the own section code throughout the bookstore categorized by themes, such as literature, history, sports, etc. Once shoppers enter the store, they first visit the tables in the section that showcases the latest frontlist books before they walk around to visit other tables.

The book placement may be subject to mixed factors. First, as publishers independently schedule publication date of newly released titles and deliver to the bookstore over time, the rapid turnover leads each book to have limited lifetime on the front table before it gets redistributed to

Figure 3.1: Snapshot of store layout inside IndieBookstore

other places. Second, books have heterogenous attributes (e.g. the popularity of previous books from the same author) that may have impact on the sales. Thus, the bookstore manager typically makes decision on book placement based on the own experience. For example, she may put selected popular books at more accessible positions for longer period because she believes that this will increase sales. Also, it is not uncommon that publishers pay placement allowance to the bookstore for prominent location[1]. As such, book placement may be correlated with unobserved factors such as book characteristics, display arrangement between publishers and bookstore, etc.

3.4 Experimental Design

We use the table for latest frontlist books to perform this experiment because unlike branded products, consumers may have limited prior knowledge about the new books and thus tend to rely on the information they gather in the bookstore to make purchase decisions (Chevalier,

[1]Anecdotal evidence showed that publishers spend significant amount of money on pay-for-display program that allows their books to be placed at the front-of-store promotional tables, see: Kennedy, R. (2005, June 5), Cash Up Front, *New York Times*. Retrieved from http://www.nytimes.com/2005/06/05/books/review/cash-up-front.html

1975). As shown in Fig 3.2, 30 books are placed on the top of the table across 5 rows and 6 columns. At the beginning of the experiment we randomly shuffled the books already on the top of the table. Right after we started the first cycle of our experiment. At the beginning of this cycle we randomly shuffled the integer numbers between 1 and 30, where 1 denotes the position in the front most row and left most column of the top of the table and 30 denotes the position in the back most row and right most column of the top of the table. Then, we placed incoming books in the slots identified by the ordered sequence of integer numbers obtained from the random shuffle. The first incoming book replaced the book on the top of the table in the position indicated by the first number in the shuffled sequence, the second incoming book replaced the book on the top of the table in the position indicated by the second number in the shuffled sequence and so on. The first experimental cycles ended when 30 new books were placed on the top of the table, therefore exhausting all available positions. Our experiment included two additional cycles similar to the first cycle.

We recorded characteristics of every book used during our experiment, such as ISBN, title, rating at GoodReads.com, price, number of pages and the date when the book was placed on the top of the table. We also recorded lifetime as the time that the book was on the top of the table until it is replaced with another book. Appendix lists the details of the placement schedule and book characteristics used throughout the experiment. In order to ensure that all shoppers face a similar table top, bookstore staffers were instructed to keep each pile of books at the same height[2] and to restore the table layout whenever it was changed due to the shopper's activity[3].

[2]IndieBookstore manager orders different amount of copies of frontlist books from different publishers. Typically, she may order more copies of the book that he believes to have higher sales and/or he can return unsold copies back to the publisher. Thus we instruct IndieBookstore staffers to put backlist books underneath the piles of books with fewer copies to equalize the height with other piles.

[3]Shoppers may pick up a book and put it back to the wrong position that may cause confusion to other shoppers.

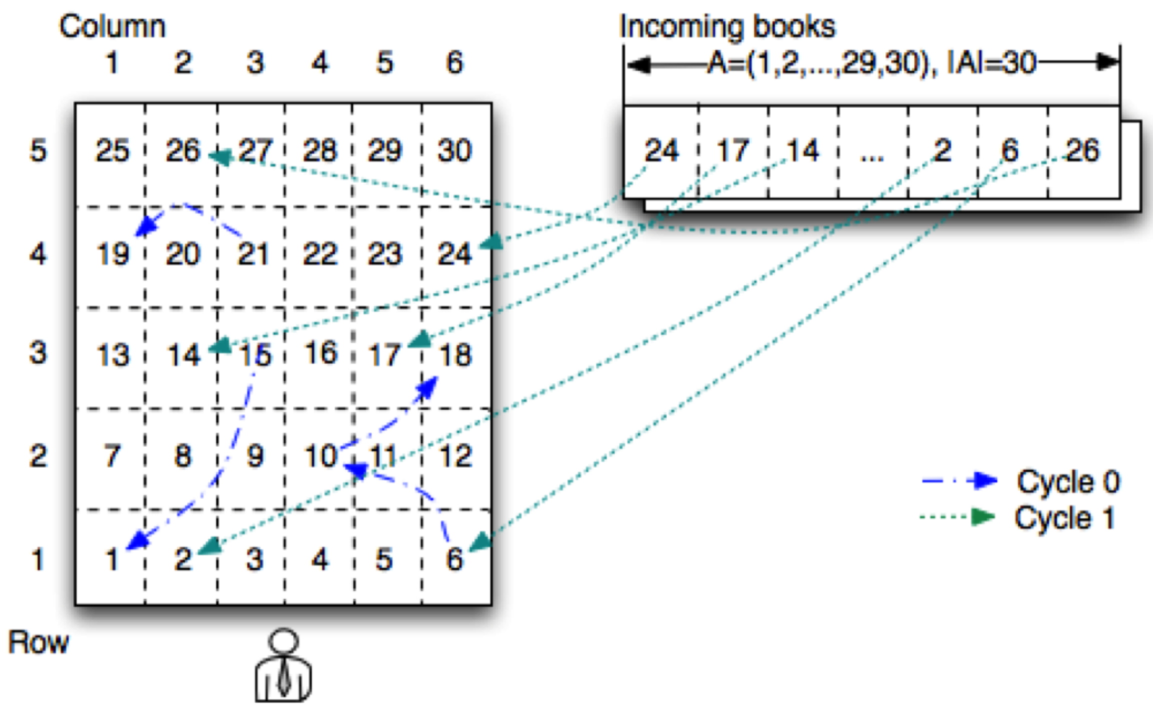

Figure 3.2: Experiment design on random book placement across cycles

Although our experiment setting provides shoppers with similar impressions during their visits to the front table, varying search costs may be incurred when they browse for books that are placed at different slots. For example, books that are closer to the shopper should be more visible and may require less effort to pick up for further valuation. Therefore, we denote the front most row and left/right most columns of the table as edge positions and the other ones as center positions, such that treatment group includes books placed at 14 edge positions and control groups includes books placed at 16 center positions. This would allows us to assign treatment of placement with almost equal probability to all incoming books.

Meanwhile, when shoppers decide to purchase a book and take it with them, the resulting lower height of that pile may signal the book quality to other shoppers.

3.5 Data Collection via Video Tracking

We aim at understanding shopper behavior along the path to purchase by capturing metrics that are indicative of the shopper's decision process. We use video tracking over other technologies to capture the shopper behavior in our experimental setting because the former is much less intrusive, whereas either eye-tracking or wearable cameras may require the participants' attention. For example, shoppers who agree to participate and wear tracking equipment for the experiment may be systematically different from shoppers in general (Hui et al., 2013). Moreover, whether intercepting shoppers would subsequently change their shopping behavior remains an unanswered question.

In order to overcome the large costs associated to large-scale video data and improve detection reliability, we used Microsoft Kinect, an advanced 3-D camera and motion capture input device and implemented vision-understanding algorithms to lessen the effort of the video encoding process. Essentially, Microsoft Kinect provides an extra high-resolution depth 3-D sensor that can complement regular video to address fundamental problems in human tracking and recognition (Han et al., 2013). With the aid of depth information, we are able to build a robust indoor sensing infrastructure that is invariant to environment changes such as lighting conditions.

As Fig 3.3 shows, we mounted the camera on the ceiling of the IndieBookstore to monitor the area surrounding the front table of new books. The camera was connected to the local server that enabled the real-time tracking of shopper activities. Fig 3.4 shows both the regular video and depth video captured using the camera. The camera was programmed to detect and track shoppers using the vision-understanding algorithm detailed in (Carvalho et al., 2016). More specifically, we obtained depth background images and constructed a 3-D spatial representation of the front table area from the depth data. We detected shoppers and started to record videos when they entered the scene by subtracting the background from the input images. As Fig 3.5 shows, we extracted foreground and segmented shoppers out of the scene using 3-D point clouds. Then, we continuously tracked each shopper in successive frames and acquired new background images after she exited the scene, such that we recorded video only when she approached the table, picked or took books and moved away from the table. Our sensing infrastructure is robust enough to handle scenarios that include multiple shoppers. Last, the camera was also programmed to send a snapshot of the monitored area to one of the author's mailbox every hour, so that we can regularly check the experimental setting in case of unexpected errors.

Figure 3.3: The installation of camera inside the IndieBookstore to monitor the front table area

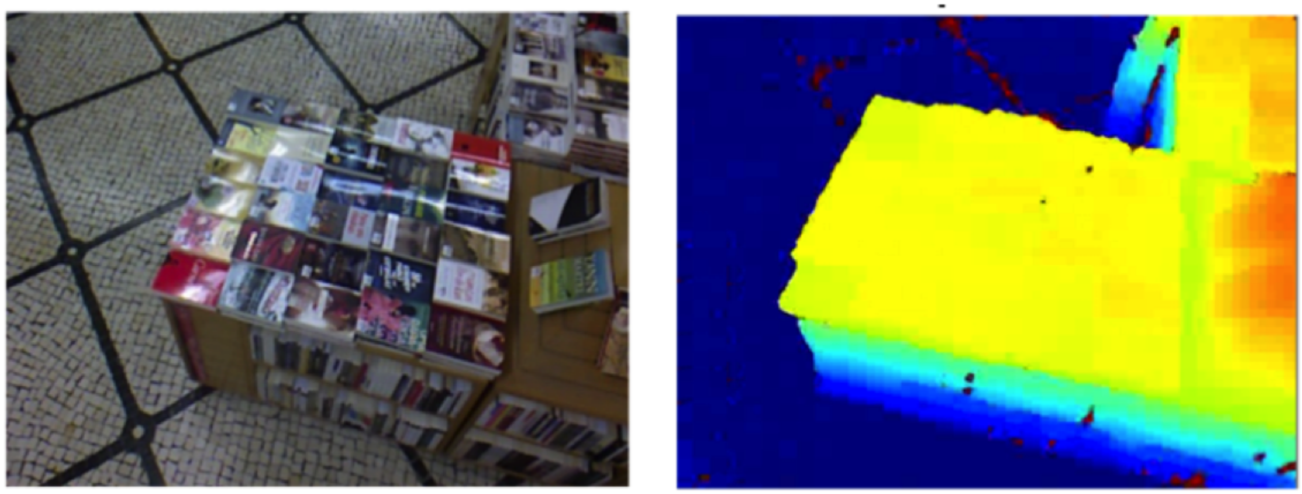

Figure 3.4: RGB image (left) and depth image (right) of the monitored front table area captured using the camera mounted on the ceiling of IndieBookstore. The depth image is calibrated with filtered signals to reduce random error of depth measurement.

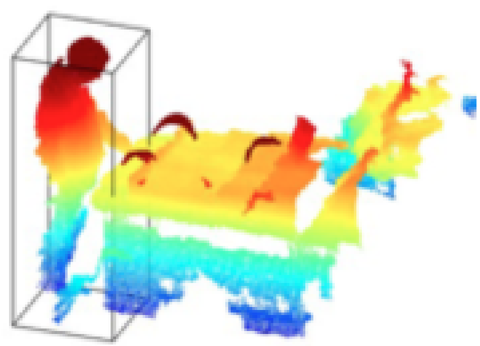

Figure 3.5: 3-D spatial representation of the front table area in a scene that a shopper (identified in a bounding box) is picking up a book

3.6 Model

Typically, the average treatment effect (ATE) can be estimated without bias by simply comparing the differences in mean outcomes between units assigned in treatment group and those assigned in control group, as long as the randomization procedure can balance both observed and unobserved factors that are correlated with the outcome. Still, using regression to adjust for baseline covariates tends to improve the precision and power of ATE estimation (Athey and Imbens, 2017). In particular, covariate adjustment may be most helpful when covariates do not affect the assignment mechanism but are predictive of outcomes. Rosenblum and van der Laan (2009) further show that estimates from regression models for analyzing randomized experiment are asymptotically correct even for misspecified models.

In our setting, book characteristics are clearly pre-treatment covariates that influence the consumer's attention and choice. Previous literature shows that reputations of author and publisher, attractiveness of book cover have significant impact on shopper's interest (d'Astous et al., 2006). In addition, prior studies find that price, genre, reputations of author and publisher, hardcover/paperback, whether translated from a foreign language, whether published previously and content quality are key determinants of book sales (Asai, 2016; Clerides, 2002; Shehu et al., 2014). Therefore, we include these covariates and/or book fixed effects in the model specifications.

3.6.1 Book-Level Specification

Because all books are randomly placed at different slots on the display table and stayed at the same slot until they are replaced, we evaluate the effect of book placement on shopper behavior by looking at how often shoppers pick and take each book, which are two key indicators of the shopper's consideration and intention to purchase processes. We normalize the number of times a book is picked and taken per day. Our variable of interest is whether the book is placed at the edge or center of the table. We also control for book characteristics, such as number of pages, price, and lifetime on the table. We use the following reduced form model to estimate the ATE of book placement on the aggregated measure of books being picked and taken:

$$y_i^p = \alpha_i^p + \beta^p x_i + \gamma^p D_i + c_i + \epsilon_i \tag{3.1a}$$

$$y_i^t = \alpha_i^t + \beta^t x_i + \gamma^t D_i + c_i + \epsilon_i \tag{3.1b}$$

where y_i^p and y_i^t represent the number of times that book i is picked and taken normalized by the lifetime, respectively; x_i represents the observed book characteristics such as $paperback$, $price$ and $goodreads$; D_i is a dummy variable that indicates whether the book is placed at edge positions; ϵ_i is an unobserved error term.

3.6.2 Visit-Level Specification

Shoppers may have different motives and time constraints to search for information about books and make their purchase decisions (Stokmans and Hendrickx, 1994). Analogous to online shopping environment, we view each shopper's visit as an "impression" to a list of 30 books presented on the table. Thus we also use visit-level model specification to account for individual heterogeneity from each impression:

$$z_{ij}^p = \alpha_{ij}^p + \beta^p x_i + \gamma^p D_i + c_i + v_j + \epsilon_{ij} \tag{3.2a}$$

$$z_{ij}^t = \alpha_{ij}^t + \beta^t x_i + \gamma^t D_i + c_i + v_j + \epsilon_{ij} \tag{3.2b}$$

where z_{ij}^p and z_{ij}^t represents the dummy variable indicating whether a shopper j picks up and takes a book i during the visit, respectively; x_i, D_i and c_i are the same as book-level specification;

v_j represents visit fixed effects; ϵ_{ij} is an unobserved error term.

Moreover, visit-level specification may help reveal the underlying mechanism of treatment effect by linking shopper's intention to purchase with her attention drawn from the impression. After all, a shopper has to pick up a book before she decides to take it with her. Thus we can use the same model specification but further restrict the sample that only contains observations where the book is picked:

$$z_{ij}^t[z_{ij}^p = 1] = \alpha_{ij}^t + \beta^t x_i + \gamma^t D_i + c_i + v_j + \epsilon_{ij} \tag{3.3}$$

3.6.3 Identification

Identification is achieved through our experimental setup. As the publication date of newly released books is exogenously determined by different publishers, and we generate the book position list before the experiment begins, all incoming books are treated at random to be placed on the top of the table. The randomization procedure essentially facilitates *masking* to both the incoming books and us researchers and thus eliminate selection bias in treatment assignment, i.e., the covariates x_i are uncorrelated with the treatment indicator D_i. Moreover, our robust and unobtrusive video tracking technology captures shopper's truthful response to the randomized design of book placement over the experimental period and this may help overcome the *attrition* problem. Finally, note that our randomized schedule also ensures lifetime of books (time on the top of the table) is also random. This eliminates any correlation between the book's lifetime and sales.

3.7 Empirical Results

Our experiment started on April 21st, 2016 and concluded 3 experimental cycles using a total of 90 books by May 24th, 2016. We collected about 91 hours of activity video, which represent about 20% of the working hours during 34 days (451 hours). In other words, our method helps us to eliminate over 80% of the encoding efforts. Once the shopper starts browsing books' front covers over the table, she may decide to pick up one book for further valuation, e.g. through reading the summary on its back cover and contents. An observation is added to our dataset with the timestamp when a shopper picks up a book and browses it for at least 10 seconds. This way we eliminate cases where there is no consideration process whatsoever. At a certain point, she may decide not to read any more and take the book away or put it back. Each observation thus updates with whether the book is taken and the duration between shopper's picking and taking/reverting. Then she can sequentially pick up as many books as she wants. In total, 1,276 customers picked up 1,751 books during our experiment and took 122 books with them during the experimental period. The exact book that each customer picked or took during our experiment was confirmed manually and a-posteriori from the video data obtained to avoid measurement error.

3.7.1 Descriptive Statistics

Table 3.1 lists the descriptive statistics of book characteristics. As explained earlier, treatment group contains 42 books that are placed at the edge positions and control group contains the rest 48 books that are placed at the center positions. $page\_counts$ is the number of pages of the book.

$paperback$ is 1 when book is softcover and 0 otherwise. $price$ is the list price of the book set by the publisher. $lifetime$ is the time that the book was on the top of the table until it is replaced with another book. $goodreads$ is the average rating at Goodreads.com of the book on a scale of 1-5. Note that 25 books do not have Goodreads rating records, with 12 books in the treatment group and 13 books in the control group, respectively[4].

We aggregate the shopper's picking and taking behavior at book level to measure the response to the treatment. Then we normalize the number of times books that are picked and taken by its lifetime on the table. The lifetime of books on the table varies from 1 day to 26 days, with an average of about 11.3 days. In general, books in the treatment group have 2.14 times more daily picks and 2 times more daily taken than those in the control group.

3.7.2 Covariate Balance

Although researchers routinely conduct hypothesis tests (e.g. t-test) to check the balance of the treatment and control groups, Imai et al. (2008) argue that this procedure is of little practical value. They show that balance is inherently a characteristic of the observed sample without reference to any broader hypothetical population. Thus test statistics are irrelevant for assessing balance. Instead, they suggest to check the balance from experiment data by directly comparing empirical distributions of covariates between treatment groups and control groups.

In our setting, trade books are marketed to a general readership and ones we used in the experiment can be assumed to be identical to the population of all trade books. However, given that we do not perform blocking due to the limited availability of new books, we need to exam-

[4]We do not include average ratings for books with less than 5 ratings to avoid unreliable measurement.

	Total n=90		Treatment group n=42		Control group n=48	
	Mean	Std	Mean	Std	Mean	Std
Covariates	[1]	[2]	[3]	[4]	[5]	[6]
Book characteristics						
$page\_counts$	332.12	150.59	328.19	176.46	335.56	125.51
$paperback$	0.92	0.27	0.88	0.33	0.96	0.20
$price$	17.29	2.30	17.65	2.55	16.97	2.03
$lifetime$	11.33	5.83	11.33	6.40	11.33	5.36
$goodreads$	3.89	0.34	3.84	0.34	3.94	0.33
Shopper behavior						
$npicks\_perday$	2.14	1.90	2.99	2.27	1.40	1.08
$ntaken\_perday$	0.15	0.20	0.20	0.25	0.10	0.13

Table 3.1: List of extracted covariates for the books used in our experiment. Descriptive statistics are performed for the total 90 books (columns [1] and [2]), 42 books placed at the edge positions (treatment group) (columns [3] and [4]) and 48 books placed at the center positions (control group) (columns [5] and [6]), respectively.

ine whether key covariates are imbalanced by chance in our sample (Athey and Imbens, 2017). Therefore, we employ side-by-side boxplots together with beeswarm plots to visualize and compare the underlying distribution and density of covariates in both the treated and control group[5].

Fig 3.6 display the empirical distribution of observed book characteristics in both the treated and control group. We show that there are no systematic differences in observed book charac-

[5]Imai et al. (2008) suggest alternative quantile-quantile plot (QQ plot) to compare the quantiles of each covariate in the treatment group against the corresponding quantiles in the control group. However, QQ plot does not yield reliable representation in small sample size, which is our case.

teristics between books placed at the edge of the table top and in the center of the table. This provides evidence of the good balance in observed covariates that our randomized procedure has achieved. We show similar findings in Appendix that balance properties also hold for these book characteristics in each cycle.

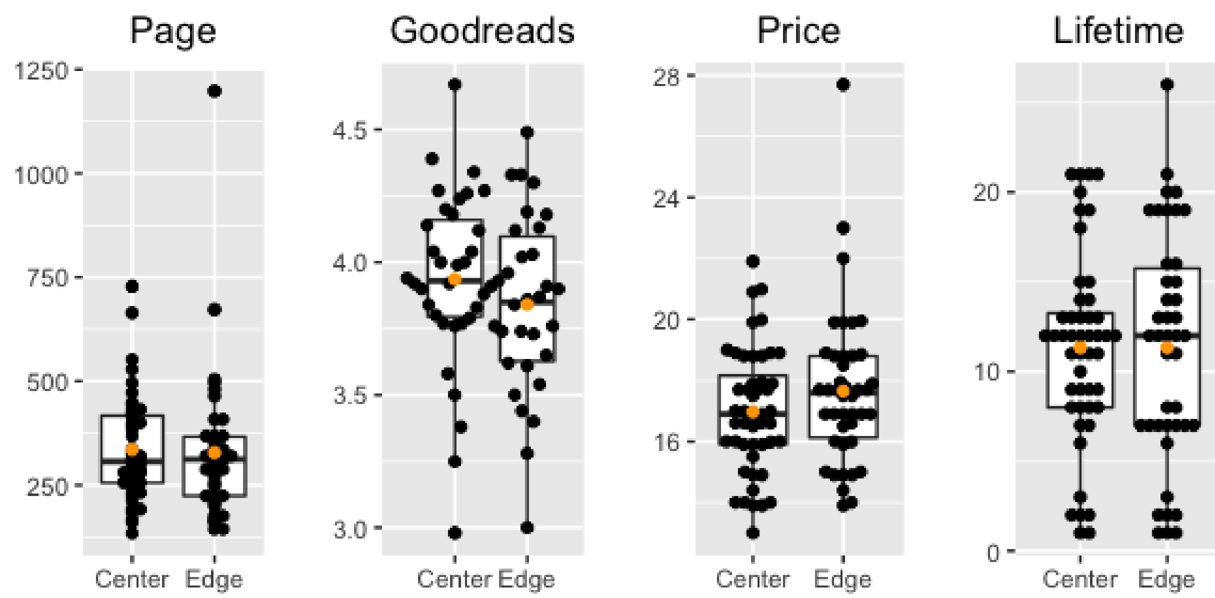

Figure 3.6: Comparing observed book characteristics between treated and control group. The boxplot and beeswarm plot show the empirical distribution and density of each covariate. Yellow dots indicate the sample means of covariates.

3.7.3 Effect of Book Placement on Daily Number of Picks and Taken

We compare how shoppers pick and take the books in treatment group and the ones in control group by examining the treatment effect on the number of times books are picked and taken per day. Given the nature of count data, we use both negative binomial and zero-inflated Poisson models with robust standard errors to control for excessive number of zeros in our dependent variables, in particular in the case of the number of times a book is taken[6]. We find that placing a book at the edge of the table positively affects the shopper's search process. Table 3.2 shows our regression results on number of times books are picked using each book as an observation and the treatment effect is very robust across model specifications. In particular, on average, books placed at the edge of the table are picked 102% more often per day than those placed at the center of the table.

Covariates Model	n_picked/day Negative binomial	n_picked/day ZIP	n_picked/day Negative binomial	n_picked/day ZIP
Edge Dummy	0.759^{***}	0.779^{***}	0.705^{***}	0.707^{***}
	(0.161)	(0.160)	(0.164)	(0.167)
Book Characteristics	No	No	Yes	Yes
Observations	90	90	90	90
Log pseudolikelihood	-152.59	-153.08	-151.10	-152.06
$p < 0.01^{***}, p < 0.05^{**}, p < 0.1^{*}$				

Table 3.2: Effect of edge vs. center on how many times a book is picked from the table (each observation is a book used during the experiment). Robust standard errors are in parentheses.

Table 3.3 further shows our regression results on number of times books are taken using each book as an observation and the treatment effect is also very robust across model specifications. We find that placing a book at the edge of the table positively affects the shopper's consideration process that may lead to their intention to purchase. In particular, on average, books placed at the edge of the table are taken 77% more often per day than those placed at the center of the table. Knowing this, the bookstore owner may maximize profit by placing books with higher margins in the edge of the table.

[6]Following the suggestion in Athey and Imbens (2017), we use the robust variance estimator implemented in standard statistical software Stata. Rosenblum and van der Laan (2009) also find that robust variance estimator for generalized linear models such as ones used in our study are asymptotically correct even when the model is misspecified.

Covariates	n_taken/day	n_taken/day	n_taken/day	n_taken/day
Model	Negative binomial	ZIP	Negative binomial	ZIP
Edge Dummy	0.640^{**}	0.640^{**}	0.572^{**}	0.572^{**}
	(0.261)	(0.261)	(0.265)	(0.265)
Book Characteristics	No	No	Yes	Yes
Observations	90	90	90	90
Log-likelihood	-33.69	-33.69	-33.36	-33.36
$p < 0.01^{***}, p < 0.05^{**}, p < 0.1^{*}$				

Table 3.3: Effect of edge vs. center on how many times a book is taken from the table (each observation is a book used during the experiment). Robust standard errors are in parentheses.

3.7.4 Effect of Book Placement on Individual Shopping Behavior

We now study our dataset from the perspective of each shopper's visit to the table of new books at the bookstore. Each shopper's visit to this table mimics an "impression" in the online world, allowing the shopper to browse for books and make purchase decisions among the 30 books on the top of the table when she visits. We use a Linear Probability Model (LPM) to estimate the effect of the random book placement on the top of the table on the shopper's search and purchase process. This type of analysis allows us to characterize the shopper's path to purchase. The dependent variables are whether the shopper picks a book, takes a book and whether she takes a book conditional on picking up the book. We also include day fixed effects to control for systematic differences over time, e.g. during weekends and holidays consumers purchase more books. Table 3.4 shows the results obtained, which are in line with the ones reported above. In particular, books placed at the edge of the table are more likely to be both picked and taken. However, conditional on being picked, shoppers are equally likely to take books placed at the edge and at the center of the table. This finding suggests that book placement positively affects consumer choice mainly through its effect on the search process and not through its effect on the consideration process. This is aligned with the theoretical framework proposed by (Branco et al., 2012) that consumers tend to gather information about products during the search process but then update their beliefs during the valuation process.

Covariates	Pick	Take	Take conditional on pick
	(a)	(b)	(c)
Edge Dummy	0.0386^{***}	0.0019^{***}	-0.0201
	(0.002)	(0.0004)	(0.0132)
Visit FE	Yes	Yes	Yes
Observations	38,280	38,280	1,735
Adjusted R^2	0.0093	0.0011	0.0202
$p < 0.01^{***}, p < 0.05^{**}, p < 0.1^{*}$			

Table 3.4: Effect of edge vs. center on whether a book is picked, taken and taken once picked from the table (each observation is a book that shopper sees at the table). Robust standard errors are in parentheses.

3.8 Discussion

Our research mainly contributes to understanding shopping behavior at the point of purchase in physical retailing in three aspects: 1) we perform a field experiment in a setting that resembles online recommender systems, where products are placed in different slots with varying degree of search costs; 2) we measure the shopper's in-store decision-making process as different stages, such as search stage, consideration stage and intention to purchase; 3) we install the 3D sensing infrastructure to monitor the shopping area and implement vision-understanding algorithms to capture shopper behavior in situational contexts in real-time.

Our empirical results provide several important findings on how product placement affects consumer choice in physical retail setting. We show that books placed at the edge of the table are more likely to be picked and taken than those placed at the center of the table. This is unsurprising as shoppers are conspicuous by books placed at prominent spots when browsing over the table due to the saliency effects. More interestingly, we also show that conditional on being picked, shoppers are equally likely to take books placed at the edge and at the center of the table. This suggests that book placement positively affect consumer choice mainly through its effect on the search process and not through its effect on the consideration process.

Moreover, we demonstrate that the book placement on the table essentially resembles the display recommendations that online bookstores are heavily in use today. Armed with the knowledge that recommending books at prominent spots may influence consumers choices by lowering the search costs, bookstore manager has incentives to place books with higher margin at the edge of the table for longer time, as she knows that this may incur higher sales. To this end, consumer welfare is likely to decrease because consumers tend to purchase books that do not give them the highest utility, as a result of search cost obfuscation.

Our empirical setting allows us to examine demand effect of newly launched experience goods in physical retailing when sellers manipulate recommender system design. As physical retailers need to manage scarce shelf space to maximize the profits, product manufacturers have to compete for prominent placement. Our results show that the success of new products at the introduction stage may critically depend on the placement assigned by retailers, as shoppers are poorly informed and tend to rely on in-store information to evaluate them. Moreover, given dramatically increasing number of new products hitting the market, underperformed products may be replaced very quickly simply due to their initial unfavorable placement. For example, in our empirical setting, bookstore manager may strategically favor some publishers that she can collect more slotting allowance to always recommend books from those publishers at better placement. Such profit-based recommender system design would mislead the consumers' preference towards the new product launch and thus hamper the competitive landscape of innovations.

Meanwhile, as non-personalized recommender systems are common in physical retailing by setting up a display that is independent of consumer, salient product placement affects consumer choice even more strongly than personalized recommender systems commonly used in online retailing. Thus the sales of products may become more concentrated and amplified towards those with better placement. Along these lines, retailers may also consider favoring products that match most consumers' preferences. This would benefits mainstream consumers for offering them wanted products with lower search costs but hurts niche consumers (Hervas-Drane, 2015). Therefore, retailers may enjoy market power over manufacturers through the design of recommendation scheme. Policy makers should be concerned about the potential welfare loss due to retailer's use of recommender systems.

Our sensing infrastructure setting indicates the great potential for physical retailers to capture holistic in-store shopper behavior and understand their preferences. Given the recent advances in retail tracking technologies, it is imperative for physical retailers to identify individual shopper's preferences during their shopping process, and provide personalized *in situ* recommendations. For example, Radhakrishnan et al. (2016) proposed in-store behavior analytics to identify the sequence of shopper activities through sensor data collected from personal smartphones. They suggested recommending newly launched products to shoppers based on their browsing behavior. Such personalized recommender system design brings immense opportunities for physical retailers to provide consumer with products that match their own preferences. However, privacy concerns would arise when consumer data is unknowingly collected and stored. Physical retailers need to find the "sweet spot" between offering personalized shopping experience and protecting consumer privacy.

3.9 Appendix

3.9.1 Random book placement list

As explained in section 3.4, we generate multiple lists with randomly shuffled integer number between 1 and 30. During each cycle, we put the incoming books according to the ordered sequence in the list until all slots on the front table have been replaced. Table shows the details of random sequences that books are placed accordingly and their characteristics.

3.9.2 Descriptive Statistics for Shopping Behavior

3.9.3 Covariate Balance

Figure 3.8 shows that there are no systematic differences in observed book characteristics between books placed at the edge of the table top and in the center of the table.

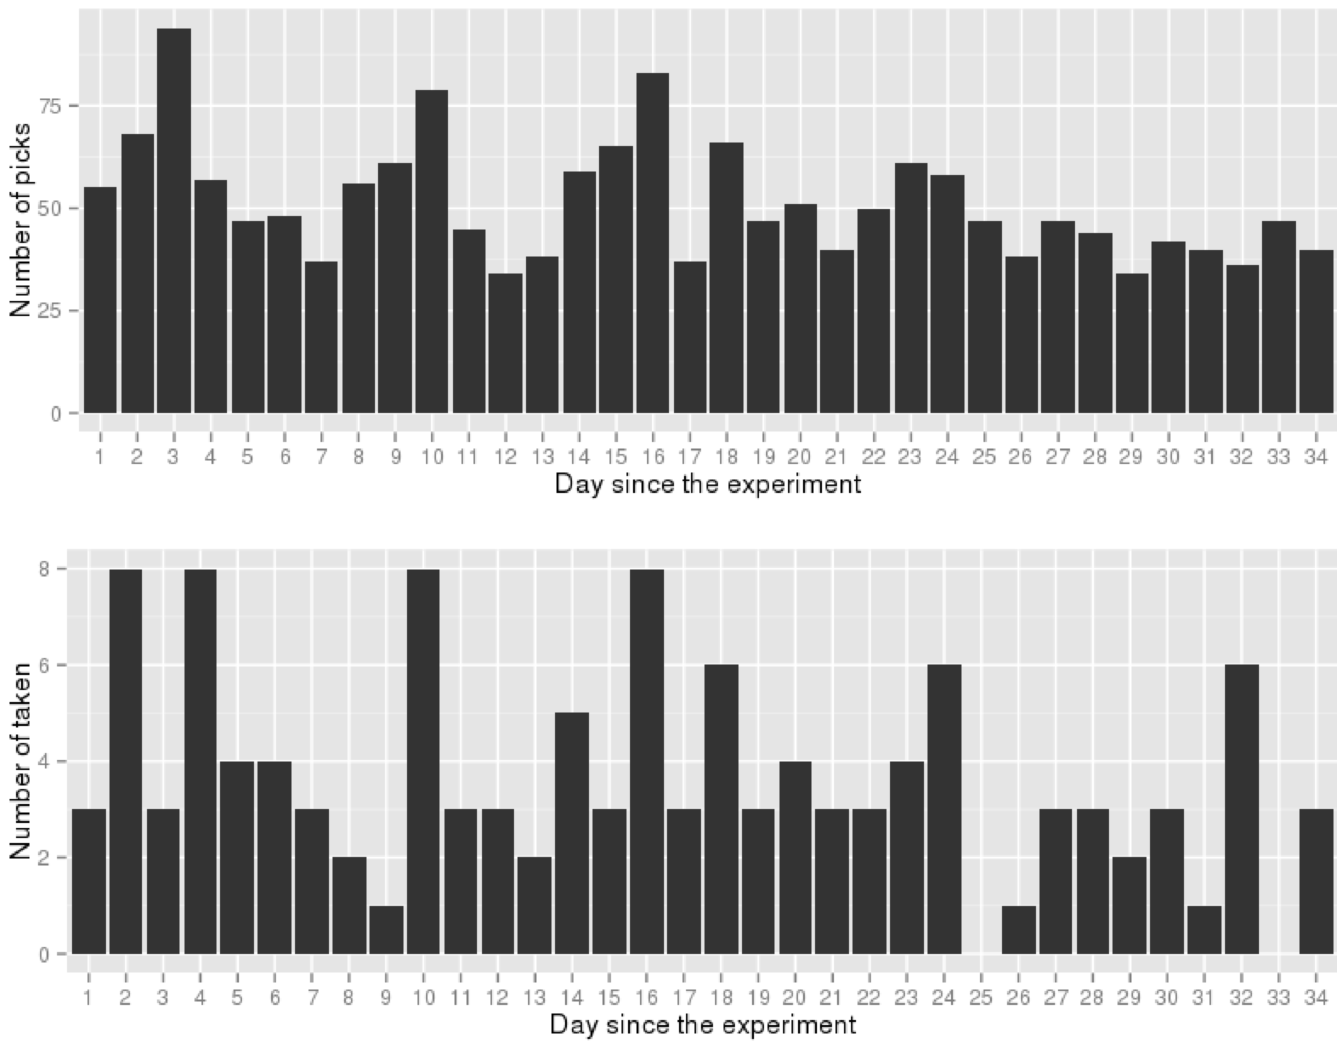

Figure 3.7: Daily number of picks and taken since the experiment

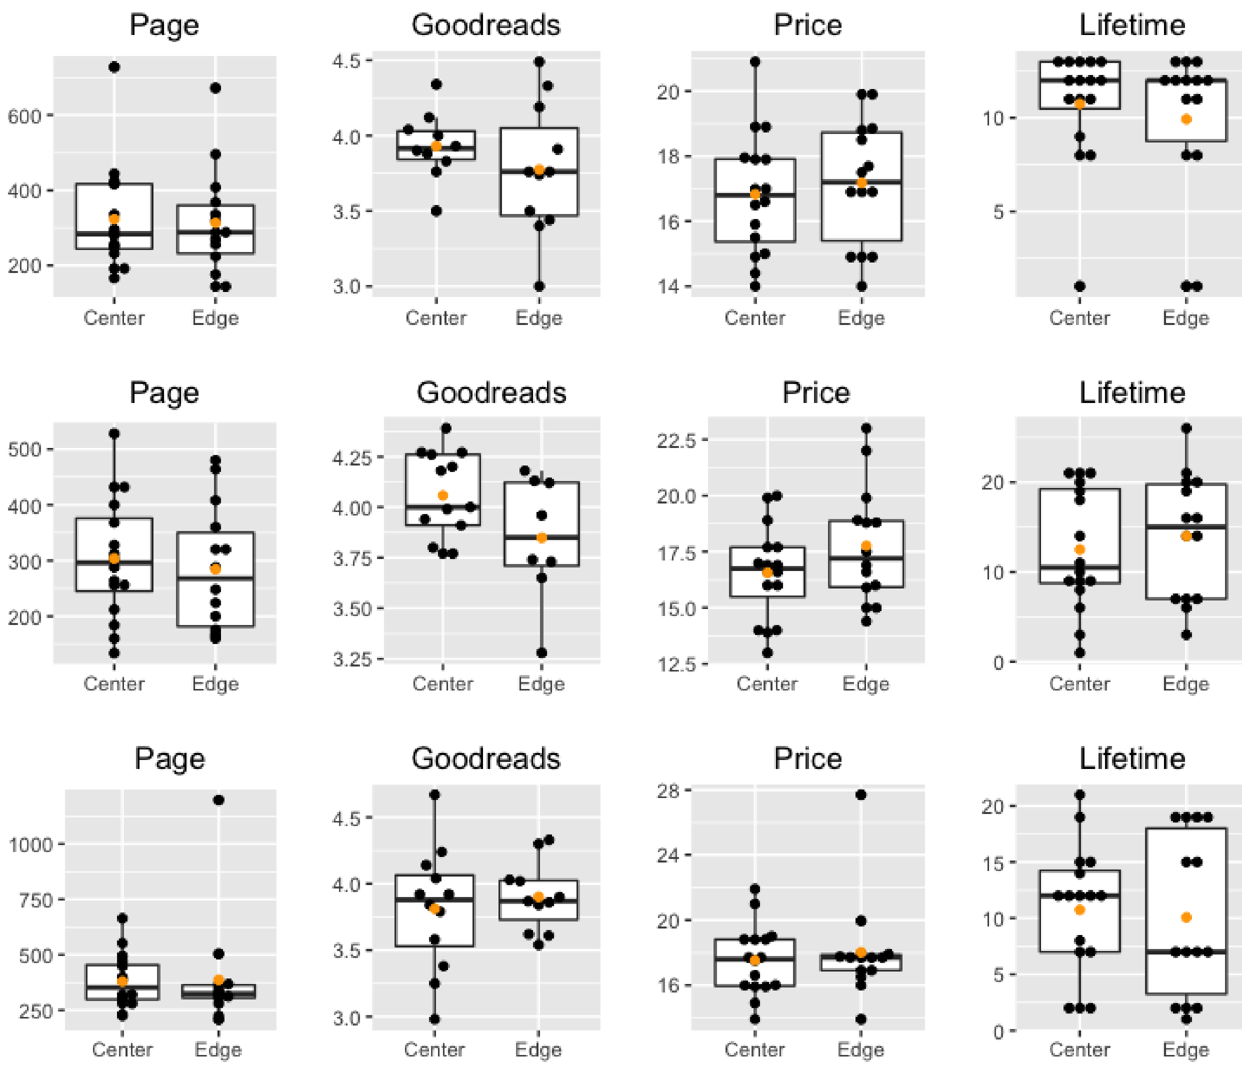

Figure 3.8: Comparing observed book characteristics between treated and control group in each cycle. The boxplot and beeswarm plot show the empirical distribution and density of each covariate. Yellow dots indicate the sample means of covariates.

Order	Row	Column	Edge	ISBN	Placement Date	Pages	Goodreads	Price
1	4	6	1	9789898839350	4/21/2016	256	3.44	17.69
2	3	5	0	9789896167103	4/21/2016	232		14.4
3	3	2	0	9789897242809	4/21/2016	424	4.04	17
4	5	5	0	9789896682989	4/21/2016	256	4.12	16.99
5	5	4	0	5601078420025	4/21/2016	254		15
6	3	4	0	9789899930407	4/21/2016	166		14
7	2	5	0	9789722059701	4/21/2016	192	4	18.9
8	1	3	1	9789722358262	4/21/2016	270	3.4	17.5
9	2	6	1	9789896379377	4/21/2016	288		16.9
10	5	1	1	9789897541575	4/21/2016	368	4.19	18.5
11	4	4	0	9789897222962	4/21/2016	280	3.5	16.6
12	2	4	0	9789897542374	4/21/2016	336	3.88	17.95
13	3	1	1	9789722060073	4/21/2016	176	3.5	14.9
14	4	1	1	9789897021961	4/21/2016	144	3	14
15	5	3	0	9789722059909	4/21/2016	192		14.9
16	2	1	1	9789896650872	4/21/2016	144	4.49	14.9
17	4	5	0	9789897021930	4/21/2016	288	3.9	15.5
18	2	3	0	9789722059923	4/21/2016	296		17.9
19	1	1	1	9789897222672	4/21/2016	408		18.8
20	3	3	0	9789896650728	4/21/2016	444	4.34	17.9
21	3	6	1	9789892335179	4/21/2016	336	3.76	16.9
22	2	2	0	9789892335032	4/21/2016	248	3.76	16.5
23	1	5	1	9789722059497	4/21/2016	224	3.76	14.9
24	4	2	0	9789897414619	4/21/2016	419	4	18.9
25	1	4	1	9789722059749	4/21/2016	288		16.9
26	5	6	1	9789896577643	4/21/2016	320	3.91	18.85
27	4	3	0	9789725305690	4/21/2016	416	3.93	15.9
28	1	2	1	9788416502530	4/21/2016	672	3.74	19.9
29	1	6	1	9789722059725	4/21/2016	496	4.33	19.9
30	5	2	0	9789896577711	4/21/2016	728	3.83	20.9

Table 3.5: The sequence of book placement at the beginning of the experiment (cycle 0). Row and column are the row number (front to back) and column number (left to right) of the slot on the table. ISBN is the unique numeric commercial book identifier. Placement Date is the date that book is placed on the table. At cycle 0, we shuffle the books that are already on the table, thus the placement date are the same for all books. Pages are page counts of the book. Goodreads is the average rating at the social book cataloging site Goodreads.com. Price is the list price (€) of the book.

Order	Row	Column	Edge	ISBN	Placement Date	Pages	Goodreads	Price
1	5	1	1	9789896232207	4/22/2016	176	3.65	15
2	1	1	1	9789722531733	4/22/2016	480		18.8
3	4	4	0	9789722531702	4/22/2016	264	4.26	16.6
4	2	1	1	9789896267490	4/29/2016	224		22
5	1	2	1	9789898839237	4/29/2016	320	3.74	18.79
6	4	2	0	9789898839558	4/29/2016	432	3.77	19.99
7	2	3	0	9789898839503	4/29/2016	256	4.27	16.99
8	2	2	0	9789898580405	4/30/2016	212		14
9	1	3	1	9789896761516	5/2/2016	168	3.28	14.4
10	5	2	0	9788416502646	5/2/2016	432		17.7
11	2	5	0	9789897414657	5/2/2016	328	3.77	16.6
12	1	4	1	9789722059688	5/2/2016	360	4.18	18.9
13	3	2	0	9781523371433	5/2/2016	312	4	16
14	5	3	0	9789897414305	5/3/2016	256		16.9
15	5	6	1	9789722059886	5/3/2016	248	3.96	16.9
16	4	6	1	9789724750415	5/3/2016	408	3.73	19.9
17	4	3	0	9789899947030	5/3/2016	160	4.39	14
18	5	4	0	9789899947047	5/3/2016	134	3.91	13
19	4	5	0	9789722356206	5/3/2016	184	4.2	13.9
20	3	6	1	9789896415891	5/3/2016	464	4.13	23
21	3	1	1	9789720048202	5/3/2016	320		16.6
22	4	1	1	9789896416003	5/3/2016	160		16
23	3	4	0	9789896416133	5/4/2016	400	3.99	16
24	2	4	0	9789896650650	5/4/2016	528	4.18	19.9
25	3	3	0	9789896577490	5/4/2016	288	4.27	17.7
26	2	6	1	9789896650865	5/4/2016	164		15.9
27	1	6	1	9789897022005	5/4/2016	200	4.12	15
28	1	5	1	9789897541858	5/4/2016	288		17.5
29	5	5	0	9789898827500	5/4/2016	304	3.8	16.9
30	3	5	0	9789722357944	5/4/2016	368	3.94	18.9

Table 3.6: The sequence of book placement at the beginning of the experiment (cycle 1). Row and column are the row number (front to back) and column number (left to right) of the slot on the table. ISBN is the unique numeric commercial book identifier. Placement Date is the date that book is placed on the table. Pages are page counts of the book. Goodreads is the average rating at the social book cataloging site Goodreads.com. Price is the list price (€) of the book.

Order	Row	Column	Edge	ISBN	Placement Date	Pages	Goodreads	Price
1	4	5	0	9789722357722	5/4/2016	224	3.92	13.9
2	5	1	1	9789897223013	5/6/2016	280	3.9	17.7
3	5	6	1	9789897220630	5/6/2016	1198	4.33	27.7
4	2	1	1	9789896379513	5/6/2016	368		17.76
5	5	4	0	9789897223006	5/6/2016	280	2.98	17.7
6	1	2	1	9789896379438	5/6/2016	224	3.61	16.9
7	5	5	0	9789897414664	5/10/2016	392	3.84	15.9
8	2	6	1	9789892334981	5/10/2016	312	3.62	16.5
9	3	6	1	9788416502547	5/10/2016	352	3.54	17.7
10	2	2	0	9789897414916	5/10/2016	280		15.9
11	5	3	0	9789720048349	5/11/2016	400		18.8
12	2	5	0	9789722530521	5/13/2016	320	3.79	16.6
13	2	4	0	9789896228187	5/13/2016	310		16
14	3	4	0	9789897102479	5/13/2016	496	3.92	15.98
15	3	3	0	9789722531146	5/13/2016	664	3.38	18.8
16	4	4	0	9789896443962	5/13/2016	472		17.7
17	2	3	0	9789892335223	5/17/2016	304	4.14	14.9
18	1	4	1	9789722358224	5/18/2016	328	3.84	17.9
19	1	6	1	9789896577810	5/18/2016	504	4.02	19.95
20	4	2	0	9789896650674	5/18/2016	552	4.04	21.9
21	3	5	0	9789897542336	5/18/2016	232	4.67	17.5
22	1	1	1	9789896577209	5/18/2016	336	4.03	17.95
23	1	3	1	9789722358255	5/18/2016	312	4.3	16.9
24	3	1	1	9789897061431	5/23/2016	304		15.99
25	4	3	0	9789896232214	5/23/2016	384	4.24	19
26	1	5	1	9789898839138	5/23/2016	368	3.86	17.69
27	5	2	0	9789898839893	5/23/2016	448	3.25	20.99
28	4	1	1	9789898839664	5/23/2016	320	3.87	17.69
29	3	2	0	9789898839879	5/23/2016	320	3.58	18.79
30	4	6	1	9789897414763	5/24/2016	208		13.9

Table 3.7: The sequence of book placement at the beginning of the experiment (cycle 2). Row and column are the row number (front to back) and column number (left to right) of the slot on the table. ISBN is the unique numeric commercial book identifier. Placement Date is the date that book is placed on the table. Pages are page counts of the book. Goodreads is the average rating at the social book cataloging site Goodreads.com. Price is the list price (€) of the book.

Chapter 4

Recommender Systems in Physical Retailing: Practices, Prospects and Policy Perspectives

4.1 Introduction

Today the business performance of physical retail industry is challenged by falling profit margins, rising cost pressures and intensified competition with online retailers (Brynjolfsson and Smith, 2000; Brynjolfsson et al., 2009; Lieber and Syverson, 2012). IT-related advances facilitate online retailers to maintain a nearly unlimited yet personalized "virtual inventory" that affect the market outcomes divergent from the traditional retailing (Brynjolfsson et al., 2009, 2011). For example, Brynjolfsson et al. (2003) show that the number of book titles available at Amazon.com is more than 23 times larger than the number of books of a Barnes & Noble superstore and 57 times greater than in a large independent bookstore. Meanwhile, the enhanced search features, such as price comparison tools and recommendation engines allow consumers to locate and acquire information about product and price at much lower costs (e.g. Bakos, 1997; Smith and Brynjolfsson, 2003; Oestreicher-Singer and Sundararajan, 2012b). These benefits lead to the proliferation of E-commerce and pose significant challenges for brick and mortar retailers to compete for their sustainable growth. According to the latest reported statistics, E-commerce sales for U.S. retailers increase over 50 times from 1998 to 2013, reaching 261 billion dollars, whereas the total sales merely double during the same period (U.S. Census Bureau 2015).

Despite setbacks, the future of retail will continue to be solidly anchored in conventional channel. Grieder et al. (2014) predict that more than 80 percent of U.S. retail sales will still remain within the walls in 2020. It is just not surprising that consumers prefer shopping non-digital products in physical retail stores: they can inspect products with sensory experience, engage with salespeople for assistance, consume immediately after purchase and return unwanted items without delivery lags (Burke, 2002; Forman et al., 2009; Lieber and Syverson, 2012). However, to survive and prosper in the current competitive environment, brick and mortar retailers have to make necessary adaptations to seamlessly integrate digital innovations in their most valuable channel - physical stores (Rigby, 2011). For example, German retailing giant Metro Group deployed several technologies, such as interactive kiosk, digital signage, electronic price tags and RFID chips in its newly opened "Extra Future Store" and reported high level of appraisal from consumers that has led to higher sales (Kalyanam et al., 2010). Likewise, British luxury brand Burberry's flagship store on Regent Street in London provides customers with an exciting shopping environment empowered by augmented reality: when consumers place products on a RFID embedded platform, they can see how those products fit on them through a projected image (Brown et al., 2014).

Online retailers increasingly rely on information systems such as recommender systems to ostensibly ease the search efforts for consumers when browsing through from the large variety of products (e.g. Hinz and Eckert, 2010; Senecal and Nantel, 2004). These systems provide substantial value for both consumers and retailers. On one hand, consumers are now living in an information-rich shopping environment and often faced with the "paradox of choice" when they are overwhelmed by the growing product assortment sizes (Schwartz, 2015). Typically, they initially incur search costs in discovering and processing product information from a myriad of choices and then only focus on a subset of alternatives before making purchase decisions (Haubl and Trifts, 2000; Mehta et al., 2003). Recommender systems can select relevant information to help consumers better match their preferences through reducing search costs and uncertainty associated with the information search (Ariely, 2000). On the other hand, retailers benefit from recommender systems to not only increase sales through enhancing both up-selling and cross-selling opportunities, but also build up consumer's perceived usefulness and loyalty (Schafer et al., 2001; Pathak et al., 2010). Therefore, the use of recommender systems has become a ubiquitous feature for most of major companies within their services (Jannach et al., 2016). For example, Netflix heavily leverage recommender systems to present movie suggestions to its users and influence about 80% of all streaming hours (Gomez-Uribe and Hunt, 2016).

Although recommender systems become so prevalent in online channel in recent years, systematic investigation on how physical retailing can benefit from its use to increase business values still lacks (Walter et al., 2012). In particular, as technological advances enable physical retailers to understand consumer preferences through tracking consumer behavior in a multitude of dimensions, there is a growing interest in learning how recommender systems built on evolving technical infrastructure in physical retailing would affect consumer choices and market outcomes. The emerging technologies bring about both opportunities and challenges for physical retailers to make use of recommender systems, and also raise public policy implications associated with the current and future practices in a similar fashion as in online world.

One recent notable example is Amazon Books, a physical extension of Amazon.com that integrate the offline and online shopping experience to help shoppers find desired books. More specifically, Amazon's physical bookstores select books based on a variety of metrics that reflect the consumer preferences, such as customer ratings and sales on its online channel and display these books under different recommendation schemes that resemble the online shopping interfaces. For example, Figure 4.1 shows a shelf display inside the Amazon's physical bookstore that recommends the book placed on the right based on its similarity to the book placed on the left (Wingfield, 2017). We show this display in an abstract form and show that such display essentially represent three different recommendation schemes: 1) product information in terms of product reviews and ratings; 2) item similarity between the two books; and 3) product placement recommendations. In most cases, as physical retailers provide recommendation implicitly through placing books at different slots, these practices essentially mark the potential for them to provide comparable recommendation services as in online setting.

In this chapter, we take an interdisciplinary approach to identify existing solutions and future opportunities that leverage emerging technologies to bring recommender systems to physical retailing. More specifically, we extend the conceptual framework for personalized recommen-

Figure 4.1: Example of book display inside the Amazon's physical bookstore that uses three recommendation schemes: 1) product information such as reviews and ratings; 2) item similarity between two neighboring books; 3) product placement recommendations.

dation and discuss about the existing solutions and potential opportunities in designing and implementing recommender systems in each stage of recommendation process. Then we discuss policy concerns that arise when applying recommender systems in physical retailing and provide policy recommendations mainly from three aspects. First, we examine whether physical recommender systems would promote or hinder consumer welfare, especially when retailers have incentives to manipulate the system design for profitability. Policy makers should be concerned about potential welfare loss when retailers mislead consumer preferences by manipulating their system design. Second, we examine whether physical recommender systems raise the same level of privacy concerns as in online setting, in particular when advanced retail tracking technologies to capture consumer behavior data have been extensively used in physical retailing industry. Third, as newly enacted data protection regulation potentially prohibits algorithmic profiling that lacks clear human interpretability in use, we note its impact on automated recommendations in physical recommender systems and discuss about how to provide explanation to accompany recommendations. Note that our policy recommendations are not in isolation from those in online setting, and instead we aim to bridge the the gap about policy discussions between online and offline context, while highlighting the distinctive features in retail stores.

4.2 Current Practices of Physical Recommender Systems

4.2.1 Non-Personalized *vs.* Personalized Recommendations

Physical retailers have long adopted certain recommendation schemes as part of marketing strategy even before the widespread use of recommender systems in online setting. For example, bookstores often recommend book collections that are derived from either manual selection (*e.g.,* editor choice) or simple statistical summaries (*e.g.,* best seller) in prominent spots (Schafer et al., 2001). This type of recommendation is regarded as non-personalized, *i.e.,* every customer gets the same recommendation independent of their tastes. Non-personalized recommendation is commonly used in retail store displays because it does not require input from particular customer and thus easy to implement (Schafer et al., 1999). By contrast, recommender systems in E-commerce also deliver personalized recommendations that are tailored to customer's own preferences. This type of recommendation requires extra information about individual customers and items, such as demographic characteristics, past browsing or purchase patterns, item characteristics, *etc.,* to establish one-to-one marketing relationship (Ansari et al., 2000). Several studies showed that personalized recommendation outperforms non-personalized recommendation in terms of better identifying customer interests and increasing customer satisfaction and loyalty (Liang et al., 2006; Zhang et al., 2011) and have greater influence on their product choices (Senecal and Nantel, 2004).

4.2.2 Conceptual Framework

Murthi and Sarkar (2003) conceptualized the personalized recommendation process into three stages, as shown in Figure 4.2: i) *learning* stage involves collection of consumer data and inference about their preferences; ii) *matching* stage provides relevant product recommendations based on the knowledge about the consumer; iii) *evaluation* stage develops appropriate metrics for assessing the effectiveness of learning and matching efforts. We extend this conceptual framework in physical retailing context with specific policy issues at each stage detailed as below.

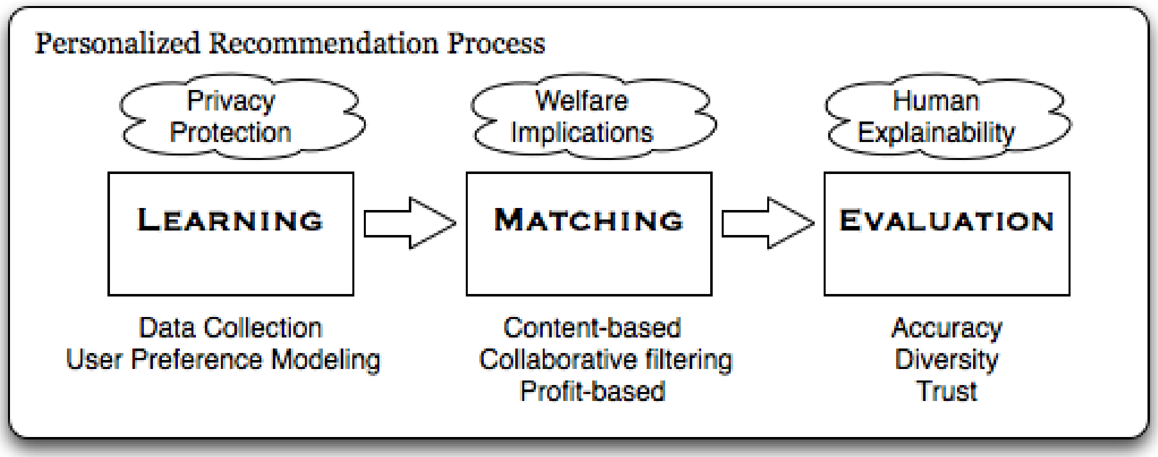

Figure 4.2: Conceptual Framework of Personalized Recommendation Process

First, physical retailers need to gather enough data to discern consumer preferences before offering them desired products. Traditionally, retailers directly interact with consumers, *e.g.,* through training salespeople or using surveys to receive individual feedbacks. However, consumers are generally unwilling to reveal much information unless they perceive clear benefits. Also, these approaches are rather expensive when taking into account of labor cost and thus may not be very effective (Lu et al., 2016). By contrast, online recommender systems enjoy many advantages from electronic data collection, as automated engine enables to instantaneously track and record consumer behavior through user registration or with the help of IP address or a cookie and can generate truthful data to avoid self-reported bias (Murthi and Sarkar, 2003). Therefore, it is crucial for physical retailers to equip certain digital technologies that allows real-time data collection and storage during the learning stage.

Second, once physical retailers obtain the information about the consumer behavior, they need to employ various algorithmic recommendation techniques to match a set of relevant products to the focal consumer accordingly. These techniques usually can be classified as content-based, collaborative filtering and hybrid (Adomavicius and Tuzhilin, 2005). Content-based recommenders utilize characteristics of items that the consumer liked previously and recommend additional items with similar characteristics. For example, a Sci-Fi movie fan may receive recommendations of more Sci-Fi movies, illustrated as "because you watched this movie" in Netflix. Collaborative filtering recommenders utilize characteristics of the consumer and recommend additional items that other consumers with similar preferences liked previously. Amazon's frequently used recommender showing "people who viewed/bought this item also viewed/bought" is a typical collaborative filtering example. However, each of the two types of recommenders has its own advantages and limitations: content-based recommenders may be effective at providing recommendations for new items but not new users that lacks the history of their past preferences; whereas collaborative filtering recommenders suffer from recommending new items that have never been rated. Hybrid approaches combine content-based and collaborative filtering methods in various ways to achieve some synergies between them. Physical retailers need to place priority on the design of recommender systems during the matching stage, particularly when taking into consideration of their profitability.

Third, evaluation of recommender system designs primarily focused on the *accuracy* metrics that empirically measure how well recommender systems infer the true preference of given users, *e.g.* to predict top N items that users are most likely to purchase (as exemplified by the well-known Netflix Prize competition) (Herlocker et al., 2004). More accurate recommendations can better capture consumer interests, especially when consumers do not have specific purpose in mind and thus result in higher decision-making satisfaction and quality (Liang et al., 2006; Zhang et al., 2011). However, accuracy-based recommender systems may yield suboptimal quality of recommendations, because a recommender can achieve high accuracy by safely giving users easy-to-predict items with little value, such as very popular or common ones that users are likely to purchase anyway (McNee et al., 2006). To increase the usefulness of recommender system, Herlocker et al. (2004) suggested to include measures of product *diversity* into the system design, *e.g.,* to recommend wider range of relatively unknown items to users, while keeping accuracy loss to a minimum (Adomavicius and Kwon, 2012). Empirical evidence showed that recommending diverse products to consumers may improve retention rates, which in turn create better sales performance (Park and Han, 2013).

Moreover, physical retailers should also center around increasing consumer's *trust* in the recommender systems. The consumer's trust plays an important role in influencing their intention to use the recommender systems (Wang and Benbasat, 2005). Panniello et al. (2016b) argue that when trust is low, physical retailers should aim at restoring trust even at a cost of profit reduction. On one hand, recommender systems may appear less trustworthy compared with other recommendation sources, such as from other consumers (Senecal and Nantel, 2004). However, development of recommender systems evaluated in a combination of accuracy and diversity can enhance cognitive consumer trust (Komiak and Benbasat, 2006). On the other hand, given that physical retailers have full control of recommender systems designs, they have incentives to manipulate what items to recommend and how to present to consumers (Pathak et al., 2010). The consumer trust in recommender system is susceptive to the consumer's perception of recommendations' manipulation. Therefore, evaluating performance of recommender systems involves many factors beyond accuracy. Retailers need to find a good balance between these evaluation metrics for recommender system design and build long-term trustworthy relationship with consumers that also benefits to their business objectives.

4.2.3 Recommender Systems Applications in Retail Stores

As we discussed in Chapter 3, physical retailers increasingly leverage marketing intelligence tools to capture and analyze multiple aspects of individual shopping behavior (Wedel and Kannan, 2016). For example, it is not uncommon nowadays for retailers to deliver personalized newsletters with coupon to consumers based on their purchase history recorded via loyalty cards. However, purchase history data does not reflect consumer's in-store decision-making process. Thus physical retailers need to deploy appropriate pervasive infrastructure *in situ* to learn about consumer preferences and influence the choice prior to purchase, thanks to the recent development in ubiquitous computing field.

Several important differences between E-commerce websites (which track consumer's path to purchase via event-driven clickstreams) and retail stores make the recent development of ubiquitous computing systems particularly useful to construct accurate user models in the latter (Walter et al., 2012). The physical shopping environment tends to be more heterogenous and complex compared to the online setting: shoppers need to visit many store areas and interact with items on the shelves, whereas they can have access to online product information from the same venue. The physical movement of shoppers requires sensing technologies to continuously monitor user activities across both spatial and temporal dimensions. Also, similar to the scenario of "who clicks what items" in clickstreams, interactions between shoppers and items need to be identified on both sides. The concept of "ubiquitous recommender systems" built upon autonomous networked sensing infrastructure may help retailers to address technical challenges to the construction of user preference model in physical world (McDonald, 2003).

There have been numerous attempts to the design and implementation of recommender systems in physical retailing context. We do not aim to perform an exhaustive survey in this chapter which is beyond the scope (see *e.g.* Anacleto et al. (2011) for an overview), but only highlight several key applications from technical perspectives. Lawrence et al. (2001) equipped supermarket shoppers with PDAs that allow them to submit their order for subsequent pickup at the store. The product recommendations are made offline based on the purchase history and only available to consumers for their next purchase. Miller et al. (2003) also used PDAs connected to network to help shoppers to select movies from online recommender systems. However, these pioneering works tend to ask consumer directly about their preferences, while ignoring valuable information about in-store consumer behavior.

In order to provide personalized recommendations for consumers at the point of purchase, physical retailers employ various technologies to track consumer behavior and learn about their preferences. A growing number of applications rely on RFID technology to sense the shopper-product interaction, as nowadays products and shopping carts are usually tagged with RFID chips that can be read and uniquely identified wirelessly. Decker et al. (2003) augmented shelves with RFID antennas to monitor the accurate location of RFID-tagged products as well as electronic price labels to show customized prices and texts. This "smart shelf" technology can detect shopper's action, such as picking a product and direct shopper's attention to read the recommendation message at that moment. Kowtsch and Maass (2010) attached RFID reader to PDA that can scan products to obtain detailed information together with related recommendations shown on the PDA screen.

The widely used smartphones open the new door for physical retailers to know more about consumers. There are many sensors embedded within the smartphones that may be very useful to infer shopping path inside the store, such as GPS, Wi-Fi, Bluetooth, *etc.* On one hand, retailers can deploy networked sensors to cover the area of store and compute shopper's location using the signal strength of the smartphone. On the other hand, each smartphone has a unique MAC address that can be used by retailers to uniquely identify the shopper. For example, Bajo et al. (2009) proposed a shopping multi-agent system that detects signals from Wi-Fi and Bluetooth to infer shopper's location in a shopping center and provide guidance on travel plan accordingly.

However, almost all these applications are merely isolated prototypes and still far from providing comparable personalized online shopping experience. Moreover, little is known about the impact of physical recommender systems on the individual choices and market outcomes. We will cover some potential opportunities in the next section.

4.3 Prospects of Physical Recommender Systems

Although the use of online recommender systems has exploded in recent times, it is imperative for physical retailers to leverage evolving new technologies to build comparable recommender systems to guide the consumers to the desired products. Given the recent technological advances in Internet of Things (IoT), AI and Cloud Computing resources, we believe that physical retailers have great potential in designing and implementing recommender systems in retail stores. We discuss about prospects of physical recommender systems from the perspective of three stages of recommender system design in conceptual framework.

First, as retailer's IT capabilities rise rapidly in order to gain competitive advantages, intelligent applications built upon pervasive IoT infrastructure that can enhance in-store shopping experience also proliferate. In particular, retail tracking and analytics leverage interconnected smart sensors to monitor consumer's movements and actions and provide rich in-store patterns. For example, Ganesan et al. (2016) showed that video feeds collected from surveillance cameras can be used to identify shoppers' physical characteristics, such as their demographics, age range as well as how they navigate, consider and determine the item choice. Costa (2014) described a set of applications that track shopper's location in retail stores by probing the signal strength from smartphones. Also, as physical loyalty cards have been quickly replaced by mobile apps, retailers can apply video and location analytics and further link the results with the loyalty program to gain business insights about individual consumer's profiles and preference signals in a situational context, such as items picked, dwell time before the shelf, *etc*. Thus it is natural to combine consumer behavior data from different sources through sensor fusion techniques to greatly improve the performance (Clifford and Hardy, 2013).

Adomavicius et al. (2005) proposed to incorporate the contextual information into recommender system design as *context-aware* recommender systems (CARS). Essentially, CARS factor in a set of contextual factors to characterize the specific physical or environmental information, such as time and location (Adomavicius et al., 2011). Prior works demonstrated that CARS outperform other types of recommender system not only in terms of relevancy (Panniello et al., 2009), but also in terms of trust and other critical business performance measures (Panniello et al., 2016a). Buser (2007) argued that the CARS design may also bring values to both consumers and sellers in physical retailing. Recommenders that lack input of individual profiles but take the current context into account could still offer satisfactory personalized services (Walter et al., 2012). Therefore, the combination of advanced tracking solutions enables unobtrusive and accurate data collection and preference identification in retail stores, which is of paramount importance in the learning stage of recommendation process.

Second, as retailers have incentives to manipulate recommender system to serve their economic goals, the recommender system design has drastically evolved to integrate profitability factors (Chen et al., 2008). Typically, the profit-based recommender systems adjust the item recommendations toward ones with higher margins. For example, several E-commerce vendors (including Amazon and Netflix) admitted to tweak its product recommendations to boost profitability with human intervention (Pathak et al., 2010). However, manipulating the outcomes of recommender systems may erode consumer trust, if consumers become aware of and perceive negatively about such manipulations (Simonson, 2005). Thus previous studies suggested recommender system design to balance several variables, such as recommendation's relevance, retailers' expected revenues and consumers' trust in the recommendation (Panniello et al., 2016b). For example, both Azaria et al. (2013) and Chen et al. (2008) demonstrated that modifying recommendations for profit maximization may be preferred by retailers without significant loss in recommendation accuracy and user satisfaction. Therefore, profit-based physical recommender systems lead to new design direction in the matching and evaluation stages of recommendation process different than the ones that best match consumer's preferences.

Third, complex machine learning models have been increasingly implemented in the matching stage of recommendation process to predict the consumer's preferences at a massive scale (Jannach et al., 2016). However, the "black-box" nature of many of these models tend to provide obscure recommendations that do not disclose the inner-workings of recommender systems (Friedrich and Zanker, 2011). Consumers may have difficulty in comprehending the underlying reasoning about item recommendations, which in turn decrease their trust and satisfaction in using the recommender systems (Wang and Benbasat, 2007). The provision of explanations is particularly favorable for physical recommender systems design, because contextual factors (*e.g.* physical or social context) may serve as high-quality supporting information to gain consumers' trust.

To sum up, we anticipate recommender system to be well received in physical retailing in the near future, due to the advances in terms of both technology and CARS design. More specifically, ubiquitous shopping environment enables physical retailers to capture in-store consumer behavior and infer preferences in a situational context through sensor fusion and large-scale analytics. Then, contextual factors may be incorporated to complement the traditional matching models and help to provide consumers with more relevant recommendations as well as associated explanations. For example, Lu et al. (2016) proposed an application of recommender system in a retailing chain for garment shoppers. The systems uses in-store cameras to capture the try-on video data and makes inferences about shoppers' preferences based on facial expression recognition. The systems then identifies shoppers with similar preferences and recommendations are made through collaborative-based matching model, which improve the experiences of garment shoppers and increase product sales. Moreover, similar to the online setting, physical retailers may also consider to place the profitability factors into the recommender system design. All these issues may potentially affect consumers' responses and market outcomes and therefore raise policy concerns to be addressed for consumer protection purposes.

4.4 Policy Perspective for Physical Recommender Systems

As policy makers aim to protect consumers in retail sector and is responsible for the enforcement of consumer law, it is important for them to understand whether using recommender systems in physical retailing would promote or hinder the consumers' decision-making while preserving their rights effectively. We cover three important policy implications that arise when physical retailers incorporate recommender systems into their business services across three stages of recommendation process: 1) welfare implications when physical retailers choose the recommender systems design to maximize the profits ; 2) privacy protection when consumer data is collected, stored and processed using emerging technologies; and 3) human interpretability in algorithmic recommendation techniques.

4.4.1 Welfare Implications for Physical Recommender Systems

As we explained in previous sections, welfare implications from using physical recommender systems remain ambiguous, due to the limited empirical studies on its impact on the individual consumer choice and market outcomes. We next discuss about potential welfare changes for using recommender systems in physical retailing through a comparative analysis of recommender systems in online setting.

In general, consumer welfare may increase as recommender system are introduced in retailing industry for consumers to browse and compare products, relative to when such practices are unavailable for the following reasons. First, recommender systems enable the reduction in search costs that facilitates consumers to build short and ordered lists with high utility among many alternatives, which prevent the negative effect of cognitive loads on consumer welfare (Botti and Iyengar, 2006). This is consistent with the effects brought by similar digitized innovations, such as shopbots and search engines (Baye et al., 2013). Second, recommender systems help consumers to ease the uncertainty associated with the purchase of new products from expansive product assortment (Brynjolfsson et al., 2003). In physical retailing context with even higher search costs, recommender systems may also work to promote consumer welfare to a greater extent. In particular, physical retailers may choose to incorporate contextual information to improve the accuracy of the recommender system, which in turn persuade consumers to take the recommended products without searching for more alternatives (Choudhary and Zhang, 2016). Therefore, consumers would also benefit from the provision of recommender systems in retail stores to aid them to locate desired products with less search efforts.

Numerous studies showed that recommender system designs play a significant role in shaping both individual and aggregate consumer choices, and thus may have differential impact on the consumer welfare, particularly with respect to different user and product segments. On one hand, lower search costs resulted from personalized recommendations may decrease the sales concentration (Anderson, 2006, p.52-57). Brynjolfsson et al. (2011) analyze a multichannel retailer that offers exactly the same product assortment online and offline, and find that consumer's usage of online recommender systems lead to higher demand for niche products that would otherwise be undiscovered. Likewise, Oestreicher-Singer and Sundararajan (2012a) show empirical evidences of flatter distribution of demand at Amazon's online bookstore. As such, consumers who prefer niche products are better off, because they can shift away from mainstream products to better match their preferences, which would not be the case without recommender systems due to high search cost (Hinz and Eckert, 2010). On the other hand, conflicting findings are reported when retailers choose different recommender system designs. Fleder and Hosanagar (2009) find that collaborative filtering based recommender systems tend to reduce sales diversity by favoring popular products with sufficient historical information, whereas individual diversity can increase. Hosanagar et al. (2014) further demonstrate that recommender systems that are assumed to cause fragmentation among the consumers and diversify demand, instead lead them to consume a more similar mix of products after recommendations. Along these lines, concerns about welfare loss at the societal level due to the fragmentation are mitigated.

As explained in the previous chapter, physical retailers with high inventory cost can only accommodate limited product variety and tend to place mainstream products at the prominent spots, *i.e.* niche products typically have higher search costs relative to the online channel. Personalized recommendations would similarly allow consumers to incur much lower search costs for niche products and result in more dispersed demand, while mainstream products still maintains the lion's share (Tan et al., 2016). Such improvement in consumer search can also drive physical retailers to include niche products that would otherwise remain unavailable (Yang, 2013). Accordingly, introduction of physical recommender systems may lead welfare gains particularly for the niche consumer segment.

However, retailers do not always benefit from the introduction of recommender systems in terms of their profitability without carefully tuning the designs (Li et al., 2014). For example, Hinz and Eckert (2010) illustrate the negative consequences of recommender systems on profits in terms of unfavorable shift in demand to niche products with lower margins. In response to this, Hervas-Drane (2015) suggest to over-represent mainstream products relative to their potential customer base, at the expense of a larger search cost increase for niche products. This indicates that retailers have incentives to manipulate the recommender systems for profit-maximization purposes (Pathak et al., 2010). Moreover, several studies demonstrated that bias in recommender systems may deliberately mislead consumer preferences (Adomavicius et al., 2013; Prawesh and Padmanabhan, 2014), *e.g.,* toward more profitable items (Panniello et al., 2016b). Such profit-based recommender system designs raise welfare concerns that have not been addressed until recently. Theoretical work by Choudhary and Zhang (2016) showed that welfare changes is contingent on the accuracy of the recommender systems. In particular, retailer's profit increases on the expense of consumer welfare, when targeting consumers with accurate product recommendations that have higher profit margins. Ferreira et al. (2016) empirically show the evidence of welfare loss generated by profit-based recommender systems. However, they also show that profit-based design may yield higher consumer welfare than traditional non-personalized design, implying that personalized recommender systems can still be useful, even when retailers consider to balance the recommendations' accuracy and product profitability for their own good.

4.4.2 Privacy Concerns associated with Retail Tracking Technologies

Physical retailers are facing the privacy-personalization tradeoff when applying physical recommender systems, because typically they need to extract more features such as the contextual in-store shopping behavior or cross-link multiple sources to augment the user models and in return provide consumers with more personalized recommendations (Awad and Krishnan, 2006). Friedman et al. (2015) performed extensive reviews on privacy risks imposed by the recommendation process and put privacy-enhancing solutions into three categories: 1) architectural approaches that minimize the data leakage threat through system design; 2) algorithmic techniques that apply cryptographic tools to protect the data; and 3) policy solutions to mandate the privacy protection through privacy laws or industry-wide self-regulations. We can further consider the first two approaches as technical solution and the last approach regarding the policy activities as non-technical solution that will be main focus in this section.

Major policymakers always aim to protect consumer privacy, not only through informing legislators to develop laws and policies governing privacy, but also urging retail sector to propose more effective self-regulatory guidelines. Although the self-regulatory practices to responsibly manage online data have been addressed, how such efforts may apply in physical contexts is largely under-explored (Sheehan and Hoy, 2000). In 2010 and 2012, FTC proposed and further amended a framework (referred hereafter as the Privacy Framework) for consumer privacy protection in the era of rapid technological change to reflect its longstanding objective to ensure Fair Information Practice Principle of Privacy Act (FTC, 2012). In particular, the Privacy Framework is issued as the guideline for all "commercial entities" that collect and use consumer data as they develop the best practices to operationalize privacy protection within their businesses. This ensures equal applicability to both online and offline retailers, implying that physical recommender systems should be subject to the privacy regulation to the same extent as in online context when retail tracking technologies are used.

FTC mandates the "commonly accepted practices" that companies can engage in without offering consumer choice. As such, online retailers need not provide choice when making product recommendations based on prior purchases that are believe to be acceptable. However, the use of recommender systems potentially fails to fit in full scope of "first-party marketing" at large, in the following scenarios when physical retailers may 1) deliver contextual recommendations through tracking shoppers through the use of sensors embedded in the third-party sources (e.g. the geo-location data from user's smartphone beyond the physical retail store) or across online third-party channels; 2) share the data with third-party firms to perform analytics.

Moreover, the Privacy Framework applies to the data that is "reasonably linkable to a specific consumer, computer or device", i.e., the data that contains "personally identifiable information". As the technological advances and the combined data from sources increasingly lead to a re-identification of traditional anonymous data, the scope of the privacy protection may essentially extend to a wide range of retail tracking practices. For example, cross-device tracking that enables retailers to combine information about consumer from both online and offline channel may fall into the scope of the Privacy Framework (FTC, 2017). Therefore, following the principles proposed in the Privacy Framework, we discuss about the implications of applying recommender systems in physical retailing as below.

Privacy by Design

Recommender systems used in physical retailing should place obligations to treat consumer data responsibly, by incorporating substantive privacy protections into the design of recommender systems. For example, physical retailers may consider introducing privacy-preserving system architecture that can compute recommendations without explicitly knowing the consumer's data in each of the three filtering methods (Friedman et al., 2015). In particular, this mechanism aims to take care of consumers who are highly sensitive to disclose their information. Moreover, the Privacy Framework recommends the collection and retention of consumer data to purposes consistent with the context in which consumers originally disclose their information. Recommender used in physical retailing specifically address this issue by designing the system to collect the contextual information for the purpose of providing personalized recommendation. Physical retailers should provide prominent notice to consumers when their data collection and retention is not for the use of physical recommender systems.

Simplified Consumer Choice

As physical retailers are increasingly creating linkage between multiple data sources to create accurate consumer profiles, they should make mechanism of robust choice meaningful to consumers. Tang et al. (2008) showed that clear disclosure can lead to enhanced consumer trust and improve the social welfare. To this end, recommender systems used in physical retailing should establish consumer choices as a baseline requirement to allows consumer to control over the collection and use of their data. In particular, physical retailers should provide clear disclosure about the data collection and use practices and offer consumer choice "in a context in which the consumer is making a decision about her data".

In general, two options are available in online data collection: 1) "opt-in" option that requires the consumer's affirmative express consent before data collection; 2) "opt-out" option that allows consumers to choose whether or not to be tracked. However, the practicality of providing choice in physical retailing needs further discussion, as it may incur extra costs to communicate with consumers and highly depend on the context of technology use. For example, the logistic problems in physical retailing prevent effective communication of privacy notices with consumers at the point of sale. The choice mechanism in physical retailing depends on the context and technology and sensitive nature of collected data that recommender system is based upon (Soltani, 2015). More specifically, we can categorize technologies used in collecting in-store shopping process into active monitoring technologies, which probe the shopping behavior through communicating with the consumer or device, such as mobile apps or persistent identifier from WiFi hotspot or cellular provider; or passive monitoring technologies, which intercepts signals from the consumer or device, such as cameras, passive cellular, WiFi and Bluetooth communications. On one hand, the former draws consumer attention when users choose to join the WiFi hotspot or uses the recommendation function on their app, thus should enable opt-in option to allow consumers to receive benefits by providing contextual information. On the other hand, the latter does not technically intrude shopping process such that physical retailers should make consumer aware, e.g. though the signage that inform consumers about the collection and use of their data and provide them with opt-out option. When the sensitive data such as facial information of the consumer or specific consumer segment such as children is collected, the retailers should enable the opt-in option regardless of the tracking technologies being used.

Tracking technology	Identifier	Active/Passive	Opt-in/Opt-out
WiFi probing	MAC address	Active	Opt-in during sign-up to hotspot
WiFi broadcast	MAC address	Passive	Opt-out via do-not-tracking
Cellular tracking	IMEI number	Active	Opt-in during sign-up to cellular service
Bluetooth sensing	iBeacon identifier	Active	Opt-in via enabling service
Mobile shopping apps	Consumer profiles	Active	Opt-in during sign-up to apps
Camera	Video	Passive	Opt-out via signage and do-not-tracking Opt-in for facial recognition

Table 4.1: Examples of retail tracking technologies for collection of in-store consumer data and consumer choice option for privacy protection.

The Privacy Framework encourages the enforceable self-regulatory initiatives that simplify disclosure and improve consumer choice mechanisms. For example, mobile location analytics companies now establish industry-wide code of conduct and offer consumers with choices to opt-out of mobile tracking by registering the MAC address of smartphones at smart-places.org. However, FTC still note the slow pace of self-regulation efforts and urge the industry to accelerate to implement enforceable self-regulatory codes under the FTC Privacy Framework.

Moreover, the Privacy Framework continues to enforce the section 5 of FTC act to take action against companies that fail to abide by self-regulatory initiatives in "unfair or deceptive" practices. (See settlement with retail tracking firm Nomi that fails to commit its opt-out option and mobile advertising firm InMobi that tracks locations of children without the consent).

Transparency of Tracking Practices

FTC recommends that provision of choice mechanism should be provided to consumers in a "prominent, relevant, and accessible place at a time and in a context when it matters to them". We list specific recommendations as below and discuss how it applies to the physical recommender systems. First, FTC encourages the development of standardized privacy statement for better comprehension and comparison. Physical retailing industry should propose standard formats and terminology for privacy statement applicable to the particular industry. Second, FTC mandates companies to provide reasonable access to the consumer data they maintain. However, this may impose extra costs for physical retailing industry to provide consumers with access to the data. Instead, physical retailers may choose to provide consumers with control over the data used for recommendation with certain granularity that does not outweigh the benefits. Third, FTC encourages companies to educate consumers about privacy choices. Physical retailers should consider to expand their efforts to educate consumers about their commercial recommendation practices, for example, make consumers aware of education material published by FTC.

4.4.3 Explainability of Recommender Systems

Recently, EU sets an important milestone for consumer data protection in the digital age, by adopting new General Data Protection Regulation (GDPR) (Regulation (EU) 2016/679) that will take effect in 2018 (European Commission, 2016). This regulation aims to define a unified data protection framework for the collection, storage and processing of personal information across the EU. In particular, the Article 22 of the regulation mandates consumers' right to obtain meaningful information about the logic behind the decision being made using algorithmic profiling[1], which potentially prohibits a wide range of machine learning algorithms with very little interpretability in use, including many being heavily implemented in each of the stages in recommendation process (Goodman and Flaxman, 2016).

The GDPR poses challenges for retailers to adapt the recommender systems to the new regulatory framework. As retailers characterize consumer preferences from many data sources in the learning stage, the resulting high-dimensional consumer data that are transformed to low-dimensional representation may be inherently difficult to interpret in the matching stage (Aggarwal, 2016, p.86). Moreover, many algorithms essentially function without transparency into the inner-working of recommender systems and thus offering minimal explanations associated with the recommendations (Tintarev and Masthoff, 2015). For example, content-based recommender systems generally provide more explainable recommendations such as the presence of features already shown in consumer's activity, whereas it would be challenging for hybrid recommender

[1] Profiling is defined as "any form of automated processing of personal data consisting of the use of personal data to evaluate certain personal aspects relating to a natural person".

systems to do the same, which use ensemble methods to represent recommendations from an aggregation measure (Friedrich and Zanker, 2011).

Generating explanations to accompany recommendation is one of the key issues that matter to the business performance of recommender systems during the evaluation stage. Tintarev and Masthoff (2007) review the issue and show the importance of improving the quality of explanation to increase usability of the recommendations. With the advent of GDPR, retailers should not only focus on improving the accuracy, but also ensure the explainability of the recommendation to satisfy consumers' needs and gain their trust (Wang and Benbasat, 2007).

4.5 Discussion

Recommender systems become increasingly prevalent in online retailing, however, it is still challenging for physical retailers to provide comparative shopping experience in their retail stores. As recent technological advances enable physical retailers to learn consumer preferences through tracking consumer behavior in a multitude of dimensions and then making relevant recommendations accordingly, there is a growing interest in understanding how recommender systems built on evolving technical infrastructure in physical retailing would affect consumer choices and market outcomes. In particular, the emerging technologies bring about both opportunities and challenges for physical retailers to make use of recommender systems, and also raise public policy implications associated with the current and future practices in a similar fashion as in online world.

We perform an interdisciplinary approach to systematically investigate the current practices, future prospects and policy implications when applying recommender systems in physical retailing. We extend the conceptual framework for personalization and discuss about the existing solutions and potential opportunities in designing and implementing recommender systems in each stage of recommendation process. More specifically, we find that physical recommender system applications are mostly isolated prototypes that focus on solving technology challenges but fail to reflect the impact on consumer behavior. However, as physical retailers may have great potential in building ubiquitous shopping environment to capture consumer preferences through many different sensors in retail stores, they can leverage the rich consumer data to incorporate contextual information to improve the recommender system design.

We discuss about policy implications for the use of physical recommender system from three perspectives. First, we examine whether physical recommender systems would promote or hinder consumer welfare, especially when retailers have incentives to manipulate the system design for profitability. We find that introduction of recommender systems generally eases the consumers' search efforts to match the consumers' preferences and increase the consumer welfare, particularly to the greater extent for consumers in the niche market. However, the profit-based recommender systems may help retailers to extract more consumer surplus, while still maintaining their trust. Policy makers should be concerned about potential welfare loss when retailers mislead consumer preferences by manipulating their system design. Second, we examine the privacy concerns propose brought by the practices of using advanced retail tracking technologies to capture consumer behavior data. According to the privacy framework proposed by FTC, physical retailers are subject to the privacy regulation at the same level as in online setting. Thus retailers should 1) promote consumer privacy at every stage of the development and design of recommender system; 2) offer simplified choice contingent on the tracking technology in use and obtain affirmative express consent before using consumer data; 3) ensure the transparency of tracking and recommendation practices compliance with the privacy rules. Third, the automated recommendation may be subject to the newly enacted data protection regulation, such that retailers should place explanations to accompany the recommendations to present to the consumers. All these issues highlight opportunities for retailers to design, implement and evaluate their recommender systems that offer convenience benefits and appropriate protection to consumers.

Chapter 5

Conclusion

5.1 Summary of Findings and Policy Implications

As recommender systems have increasingly become prevalent to guide consumers to find their desired products in many industries, understanding the impact of recommender systems on consumer behavior is critical to business performance and raises important policy implications. In this thesis, we examine the role of different recommendation schemes on consumers' switching and search behavior in two distinct case studies and followed with a comparative analysis between offline and online applications of recommender systems. Below we summarize the main findings and policy implications from these studies.

In the first study, we look at the effect of peer recommendations on subscriber churn in a large mobile network. We use the mobile phone dataset to analyze the relational dynamics between millions of individuals with fine granularity. We account for peer influence as the effect of number of churned friend on ego churn and perform generalized propensity score method to disentangle peer influence from confounding factors, such as homophily. We find that the ego's propensity to churn increases with the number of friends that churn. More specifically, the cumulative effect of up to 5 friend churning is at best 10% and peer influence that arise from strong friends churn are higher compared with weak friends churn, which is very indicative of role of tie strength in moderating recommendations from peers.

As churn rate is frequently used as proxy measure of competitive dynamics in the wireless industry, our results show that effective interpersonal communication may work as peer recommendations to influence consumer to make informed switching decisions, which may facilitate the competition among wireless carriers.

In the second study, we implement an in-vivo randomized field experiment to measure the effect of product display recommendations on shopper behavior at the point of purchase in a physical bookstore. We leverage video tracking technologies to monitor how shoppers respond to random book placement, which induces random search costs. We show that books placed at the edge of the table are more likely to be picked and taken than those placed at the center of the table. This is unsurprising as shoppers are conspicuous by books placed at prominent spots when browsing over the table due to the saliency effects. More interestingly, we also show that conditional on being picked, shoppers are equally likely to take books placed at the edge and at the center of the table. This suggests that display recommendations positively affect consumer choice mainly through its effect on the search process and not through its effect on the consideration process.

In our empirical setting, bookstore manager has incentives to place books with higher margin at the edge of the table for longer time, as she knows that this may incur higher sales. In particular, bookstore manager may strategically favor some publishers that she can collect more slotting allowance. To this end, consumer welfare will decrease because consumers are likely to purchase books that do not give them the highest utility, as a result of search cost obfuscation. Meanwhile, as our display recommendation is non-personalized, retailers may also consider favoring products that match most consumers' preferences. This would benefits mainstream consumers for offering them wanted products with lower search costs but hurts niche consumers. As such, retailers may enjoy market power over manufacturers through the design of recommendation scheme.

In the third study, we perform an interdisciplinary approach to systematically investigate the current practices, future prospects and policy implications when applying recommender systems in physical retailing. We extend the conceptual framework for personalization and discuss about the existing solutions and potential opportunities in designing and implementing recommender systems in each stage of recommendation process. However, we find that physical recommender system applications are mostly isolated prototypes that focus on solving technology challenges but very limited in reflecting the impact on consumer behavior. As retailer's IT capabilities rise rapidly in order to gain competitive advantages, intelligent applications built upon ubiquitous shopping environment show the great potential to capture in-store consumer behavior and infer preferences in a situational context through sensor fusion and large-scale analytics. Such contextual information may be incorporated to complement the traditional recommender systems design and help to provide consumers with more relevant recommendations.

The potential use of context-aware physical recommender system in physical retailing raises several policy implications for the from three perspectives. First, we find that introduction of recommender systems generally eases the consumers' search efforts to match the consumers' preferences and increase the consumer welfare, particularly to the greater extent for consumers in the niche market. However, the profit-based recommender systems may help retailers to extract more consumer surplus, while still maintaining their trust. Policy makers should be concerned about potential welfare loss when retailers mislead consumer preferences by manipulating their system design. Second, we examine the privacy concerns propose brought by the practices of using advanced retail tracking technologies to capture consumer behavior data. According to the privacy framework proposed by FTC, physical retailers are subject to the privacy regulation at the same level as in online setting. Thus retailers should 1) promote consumer privacy at every stage of the development and design of recommender system; 2) offer simplified choice contingent on the tracking technology in use and obtain affirmative express consent before using consumer data; 3) ensure the transparency of tracking and recommendation practices compliance with the privacy rules. Third, the automated recommendation may be subject to the newly enacted data protection regulation, such that retailers should place explanations to accompany the recommendations to present to the consumers.

5.2 Future Work

Our studies call for more empirical works to examine the impact of recommender systems on consumer choices, especially in industries with limited applications, such as physical retailing. We can further develop more sophisticated retail tracking technologies to detect and follow shopper's trace with more granularity. For example, we can capture the time that consumers spend in the store to account for different levels of cognitive loads, which enables us to explore the heterogenous effects of recommendations on each individual. Moreover, we can further combine the sales data to measure the aggregated market outcomes. This would allow us to uncover the linkage between the recommendations and business performance, while understanding the underlying mechanisms recommender systems in reducing costs of consumer search.

In this thesis, as we have put together several methodologies that can be very useful to study the causal impact of recommendations on consumer choices, we can also extend our findings in recommender systems to the fields that have social impact, such as healthcare, unemployment, *etc*. The public sectors typically hold rich data about the citizens across many dimensions and aim to provide personalized services to them. Thus it would be helpful for us to design and implement different recommendation schemes to ease the search efforts of citizens for better healthcare service or employment opportunities, which would have significant impact on the social welfare.

Bibliography

Adair, J. G. 1984. "The Hawthorne effect: A reconsideration of the methodological artifact," *Journal of Applied Psychology* (69:2), pp. 334–355.

Adomavicius, G., Bockstedt, J. C., Curley, S. P., and Zhang, J. 2013. "Do Recommender Systems Manipulate Consumer Preferences? A Study of Anchoring Effects," *Information Systems Research* (24:4), pp. 956–975.

Adomavicius, G., and Kwon, Y. 2012. "Improving Aggregate Recommendation Diversity Using Ranking-Based Techniques," *IEEE Transactions on Knowledge and Data Engineering* (24:5), pp. 896–911.

Adomavicius, G., Mobasher, B., Ricci, F., and Tuzhilin, A. 2011. "Context-Aware Recommender Systems," *AI Magazine* (32:3), pp. 191–226.

Adomavicius, G., Sankaranarayanan, R., Sen, S., and Tuzhilin, A. 2005. "Incorporating contextual information in recommender systems using a multidimensional approach," *ACM Transactions on Information Systems* (23:1), pp. 103–145.

Adomavicius, G., and Tuzhilin, A. 2005. "Toward the next generation of recommender systems: a survey of the state-of-the-art and possible extensions," *IEEE Transactions on Knowledge and Data Engineering* (17:6), pp. 734–749.

Agarwal, A., Hosanagar, K., and Smith, M. 2011. "Location, Location, Location: An Analysis of Profitability of Position in Online Advertising Markets," *Journal of Marketing Research* (48:6), pp. 1057–1073.

Aggarwal, C. C. 2016. *Recommender Systems: The Textbook*, Switzerland: Springer International Publishing.

Ahn, J.-H., Han, S.-P., and Lee, Y.-S. 2006. "Customer churn analysis: Churn determinants and mediation effects of partial defection in the Korean mobile telecommunications service industry," *Telecommunication Policy* (30:10-11), pp. 552–568.

Anacleto, R., Luz, N., Almeida, A., Figueiredo, L., and Novais, P. 2011. "Shopping Center Tracking and Recommendation Systems," *Advances in Intelligent and Soft Computing* (87), pp. 299–308.

Anderson, C. 2006. *The Long Tail*, New York NY: Hachette Book.

Ansari, A., Essegaier, S., and Kohli, R. 2000. "Internet Recommendation Systems," *Journal of Marketing Research* (37:3), pp. 363–375.

Aral, S., Muchnik, L., and Sundarajan, A. 2009. "Distinguishing influence-based contagion from homophily-driven diffusion in dynamic networks," *Proceedings of National Academy of Science* (106:51), pp. 21,544–21,549.

Ariely, D. 2000. "Controlling the Information Flow: Effects on Consumer's Decision Making and Preferences," *Journal of Consumer Research* (27:2), pp. 233–248.

Asai, S. 2016. "Determinants of demand and price for best-selling novels in paperback in Japan," *Journal of Cultural Economics* (40:3), pp. 375–392.

Atalay, A. S., Bodur, H. O., and Rasolofoarison, D. 2012. "Shining in the Center: Central Gaze Cascade Effect on Product Choice," *Journal of Consumer Research* (39:4), pp. 848–866.

Athey, S., and Imbens, G. 2017. "The Econometrics of Randomized Experiments," in *Handbook of Economic Field Experiments*, E. Duflo and A. Banerjee (eds.), Elsevier, chap. 5.

Awad, N. F., and Krishnan, M. S. 2006. "The Personalization Paradox: An Empirical Evaluation of Information Transparency and the Willingness to be Profiled and the Willing to be Profiled Online for Personalization," *MIS Quarterly* (30:1), pp. 13–28.

Azaria, A., Hassidim, A., Kraus, S., Eshkol, A., Weintrab, O., and Netanely, I. 2013. "Movie recommender system for profit maximization," in *Proceedings of the 7th ACM Conference on Recommender Systems*, , RecSys '13, ACM.

Bajo, J., Corchado, J. M., Paz, Y. D., Paz, J. F. D., Rodriguez, S., Martin, Q., and Abraham,

A. 2009. “SHOMAS: Intelligent guidance and suggestions in shopping centres,” *Applied Soft Computing* (9:2), pp. 851–862.

Bakos, J. Y. 1997. “Reducing Buyer Search Costs: Implications for Electronic Marketplaces,” *Management Science* (43:2), pp. 1676–1692.

Baye, M. R., los Santos, B. D., and Wildenbeest, M. R. 2013. “The Evolution of Product Search,” *Journal of Law, Economics & Policy* (9:2), pp. 201–221.

Berson, A., Smith, S., and Thearling, K. 2000. *Building Data Mining Applications for CRM*, Enterprise Computing Series, McGraw-Hill Osborne.

Bezawada, R., Balachander, S., Kannan, P., and Shankar, V. 2009. “Cross-Category Effects of Aisle and Display Placements: A Spatial Modeling Approach and Insights,” *Journal of Marketing* (73:3), pp. 99–117.

Blondel, V., Decuyper, A., and Krings, G. 2015. “A survey of results on mobile phone datasets analysis,” *EPJ Data Science* (4:10).

Botti, S., and Iyengar, S. S. 2006. “The Dark Side of Choice: When Choice Impairs Social Welfare,” *Journal of Public Policy and Marketing* (25:1), pp. 24–38.

Bramoullé, Y., Djebbari, H., and Fortin, B. 2009. “Identification of peer effects through social networks,” *Journal of Econometrics* (150:1), pp. 41–55.

Branco, F., Sun, M., and Villas-Boas, J. M. 2012. “Optimal Search for Production Information,” *Management Science* (58:11), pp. 2037–2056.

Braun, M., and Schweidel, D. A. 2011. “Modeling Customer Lifetimes with Multiple Causes of Churn,” *Marketing Science* (30:5), pp. 881–902.

Brown, M., Moriarty, M., and Mendoza-Pena, A. 2014. “On Solid Ground: Brick and Mortar is the Foundation of Omnichannel Retailing,” Analysis report, A.T. Kearney.

Brynjolfsson, E., Dick, A., and Smith, M. 2010. “A nearly perfect market? Differentiation vs. price in consumer choice,” *Quantitative Marketing and Economics* (8:1), pp. 1–33.

Brynjolfsson, E., Hu, Y. J., and Rahman, M. S. 2009. “Battle of the Retail Channels: How Product Selection and Geography Drive Cross-Channel Competition,” *Management Science*

(55:11), pp. 1755–1765.

Brynjolfsson, E., Hu, Y. J., and Simester, D. 2011. "Goodby Pareto Principle, Hello Long Tail: The Effect of Search Costs on the Concentration of Product Sales," *Management Science* (57:8), pp. 1373–1386.

Brynjolfsson, E., Hu, Y. J., and Smith, M. 2003. "Consumer Surplus in the Digital Economy: Estimating the Value of Increased Product Variety at Online Booksellers," *Management Science* (49:11), pp. 1580–1596.

Brynjolfsson, E., and Smith, M. 2000. "Frictionless Commerce? A Comparison of Internet and Conventional Retailers," *Management Science* (46:4), pp. 563–585.

Bucklin, R., and Gupt, S. 1999. "Commercial Use of UPC Scanner Data: Industry and Academic Perspectives," *Marketing Science* (18:3), pp. 247–273.

Burke, R. 2002. "Technology and the Customer Interface: What Consumers Want in the Physical and Virtual Store," *Journal of the Academy of Marketing Science* (30:4), pp. 411–432.

Burke, R. R. 2006. "The Third Wave of Marketing Intelligence," in *Retailing in the 21st Century: Current and Future Trends*, M. Krafft and M. K. Mantrala (eds.), Springer Berlin Heidelberg, pp. 113–125.

Buser, D. 2007. "Context-Based Recommender Systems in Conventional Grocery - an Economic Analysis," in *Proceedings of 40th Annual Hawaii International Conference on System Sciences*, , HICSS '07, IEEE.

Carvalho, J., Marques, M., Costeira, J. P., and Jorge, P. 2016. "Detecting People in Large Crowded Spaces using 3D Data from Multiple Cameras," in *Proceedings of the* 11^{th} *Joint International Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications*, , Porto, Portugal.

Chandon, P., Hutchinson, J. W., Bradlow, E., and Young, S. H. 2007. "Measuring the Value of Point-of-Purchase Marketing with Commercial Eye-Tracking Data," in *Visual Marketing: From Attention to Action*, M. Wedel and R. Pieters (eds.), Lawrence Erlbaum Associates, pp. 225–258.

Chandon, P., Hutchinson, J. W., Bradlow, E. T., and Young, S. H. 2009. “Does In-Store Marketing Work? Effects of the Number and Position of Shelf Facings on Brand Attention and Evaluation at the Point of Purchase,” *Journal of Marketing* (73:6), pp. 1–17.

Chen, L.-S., Hsu, F.-H., Chen, M.-C., and Hsu, Y.-C. 2008. “Developing recommender systems with the consideration of product profitability for sellers,” *Information Sciences* (178:4), pp. 1032–1048.

Chen, Y., and Yang, S. 2007. “Estimating Disaggregate Models Using Aggregate Data Through Augmentation of Individual Choice,” *Journal of Marketing Research* (44:4), pp. 613–621.

Chevalier, M. 1975. “Increases in Sales Due to In-Store Display,” *Journal of Marketing Research* (12:4), pp. 426–431.

Cho, D., Ferreira, P., and Telang, R. 2012. “The Impact of Mobile Number Portability on Price, Competition and Consumer Welfare,” in *Proceedings of the 40th Telecommunications Policy Research Conference*, , TPRC ’12.
URL `https://ssrn.com/abstract=2265104`

Cho, D., Kumar, A., and Telang, R. 2016. “IExclusivity: An Effect of iPhone Exclusivity Arrangement on Demand for Smartphones,” Working Paper No. 2016-004, KAIST College of Business.
URL `http://ssrn.com/abstract=2771403`

Choudhary, V., and Zhang, Z. J. 2016. “Recommender Systems and Consumer Product Search,” Tech. rep., Working Paper.

Christenfeld, N. 1995. “Choices from Identical Options,” *Psycological Science* (6:1), pp. 50–55.

Clerides, S. K. 2002. “Book value: intertemporal pricing and quality discrimination in the US market for books,” *International Journal of Industrial Organization* (20), pp. 1385–1408.

Clifford, S., and Hardy, Q. 2013. “Attention, Shoppers: Store is Tracking Your Cell,” *New York Times* (July 14, 2013).

Costa, T. 2014. “How Location Analytics Will Transform Retail,” *Harvard Business Review* (March 2014).

Court, D., Elzinga, D., Mulder, S., and Vetvik, O. J. 2009. "The consumer decision journey," *Mckinsey Quarterly* (June).

Cox, K. 1964. "The Responsiveness of Food Sales to Shelf Space Changes in Supermarkets," *Journal of Marketing Research* (1:2), pp. 63–67.

Curhan, R. H. 1972. "The Relationship between Shelf Space and Unit Sales in Supermarkets," *Journal of Marketing Research* (9:4), pp. 406–412.

Dasgupta, K., Singh, R., Viswanathan, B., Chakraborty, D., Mukherjea, S., Nanavati, A. A., and Joshi, A. 2008. "Social ties and their relevance to churn in mobile telecom networks," in *Proceedings of the 11th International Conference on Extending database technology: Advances in database technology*, , EDBT '08, ACM.

d'Astous, A., Colbert, F., and Mbarek, I. 2006. "Factors influencing readers' interest in new book releases: An experimental study," *Poetics* (34:2), pp. 134–147.

Decker, C., Kubach, U., and Beigl, M. 2003. "Revealing the Retail Black Box by Interaction Sensing," in *Proceedings of the 23rd International Conference on Distributed Computing Systems Workshops*, , ICDCSW '03, IEEE.

Dierkes, T., Bichler, M., and Krishnan, R. 2011. "Estimating the effect of word of mouth on churn and cross-buying in the mobile phone market with Markov logic networks," *Decision Support Systems* (51:3), pp. 361–371.

Dreze, X., Hoch, S., and Purk, M. 1994. "Shelf Management and Space Elasticity," *Journal of Retailing* (70:4), pp. 301–326.

Eshghi, A., Haughton, D., and Topi, H. 2007. "Determinants of customer loyalty in the wireless telecommunications industry," *Telecommunication Policy* (31:2), pp. 93–106.

European Commission 2016. "General Data Protection Regulation," .
URL `http://ec.europa.eu/justice/data-protection/reform/files/regulation_oj`

Farley, J. U., and Ring, L. W. 1966. "A Stochastic Model of Supermarket Traffic Flow," *Operations Research* (14:4), pp. 555–567.

FCC 2009. "Annual Report and Analysis of Competitive Market Conditions With Respect to

Commercial Mobile Services," WT Docket 08-27, Federal Communication Commission.

FCC 2011. "Annual Report and Analysis of Competitive Market Conditions With Respect to Mobile Wireless, Including Commercial Mobile Services," WT Docket 10-133, Federal Communication Commission.

Ferreira, P., Zhang, X., Belo, R., and Matos, M. 2016. "Welfare Properties of Recommender Systems: Theory and Results from a Randomized Experiment," Tech. rep., SSRN. URL http://ssrn.com/abstract=2856794

Fleder, D., and Hosanagar, K. 2009. "Blockbuster Culture's Next Rise or Fall: The Impact of Recommender Systems on Sales Diversity," *Management Science* (55:5), pp. 697–712.

Forman, C., Ghose, A., and Goldfarb, A. 2009. "Competition Between Local and Electronic Markets: How the Benefits of Buying Online Depends on Where You Live," *Management Science* (55:1), pp. 47–57.

Frank, R. E., and Massy, W. F. 1970. "Shelf Position and Space Effects on Sales," *Journal of Marketing Research* (7:1), pp. 59–66.

Friedman, A., Knijnenburg, B. P., Vanhecke, K., Martens, L., and Berkovsky, S. 2015. "Privacy Aspects of Recommender Systems," in *Recommender Systems Handbook*, F. Ricci, L. Rokach, B. Shapira, and P. B. Kantor (eds.), Springer US, chap. 19, pp. 649–688.

Friedrich, G., and Zanker, M. 2011. "A Taxonomy for Generating Explanations in Recommender Systems," *AI Magazine* (32:3), pp. 90–98.

FTC 2012. "Protecting Consumer Privacy in an Era of Rapid Change: Recommendations for Businesses and Policymakers," Commission report, Federal Trade Commission.

FTC 2017. "Cross-Device Tracking," Staff report, Federal Trade Commission.

Ganesan, V., Ji, Y., and Patel, M. 2016. "Video meets the Internet of Things," *McKinesy Quarterly* (December).

Gaur, V., and Fisher, M. 2005. "In-Store Experiments to Determine the Impact of Price on Sales," *Production and Operations Management* (14:4), pp. 377–387.

Gerpott, T. J., Ahmadi, N., and Weimar, D. 2015. "Who is (not) convinced to withdraw a contract termination announcement? ? A discriminant analysis of mobile communications customers in Germany," *Telecommunication Policy* (39:1), pp. 38–52.

Ghose, A., Goldfarb, A., and Han, S. P. 2013. "How Is the Mobile Internet Different? Search Costs and Local Activities," *Information Systems Research* (24:3), pp. 613–631.

Ghose, A., Ipeirotis, P., and Li, B. 2014. "Examining the Impact of Ranking on Consumer Behavior and Search Engine Revenue," *Management Science* (60:7), pp. 1632–1654.

Gilbride, T. J., Inman, J. J., and Stilley, K. M. 2015. "Examining the Impact of Ranking on Consumer Behavior and Search Engine Revenue," *Journal of Marketing* (79:3), pp. 57–73.

Glady, N., Baesens, B., and Croux, C. 2009. "Modeling Churn using Customer Lifetime Value," *European Journal of Operational Research* (197:1), pp. 402–411.

Gomez-Uribe, C. A., and Hunt, N. 2016. "The Netflix Recommender System: Algorithms, Business Value, and Innovation," *ACM Transactions on Management Information Systems* (6:4), pp. 1–19.

Goodman, B., and Flaxman, S. 2016. "European Union regulations on algorithmic decision-making and a "right to explanation"," in *Proceedings of the ICML Workshop on Human Interpretability in Machine Learning*, , WHI '16, ACM.

Granbois, D. H. 1968. "Examining the Impact of Ranking on Consumer Behavior and Search Engine Revenue," *Journal of Marketing* (32:4), pp. 28–33.

Granovetter, M. 1978. "Threshold Models of Collective Behavior," *American Journal of Sociology* (83:6), pp. 1420–1443.

Grieder, P., Buck, R., Banfi, F., Kment, V., and Fitzner, J. 2014. "The future of retail: How to make your bricks click," *McKinsey & Company Consumer and Shopper Insights* (September).

Guadagni, P. M., and Little, J. D. C. 1983. "A Logit Model of Brand Choice Calibrated on Scanner Data," *Marketing Science* (2:3), pp. 203–238.

Guardabascio, B., and Ventura, M. 2014. "Estimating the dose-response function through a generalized linear model approach," *The Stata Journal* (14:1), pp. 141–158.

Gupta, S., Lehmann, D., and Stuart, J. 2004. "Valuing Customer," *Journal of Marketing Research*

(41:1), pp. 7–18.

Haenlein, M. 2013. "Social interactions in customer churn decisions: The impact of relationship directionality," *International Journal of Research in Marketing* (30:3), pp. 236–248.

Han, J., Shao, L., Xu, D., and Shotton, J. 2013. "Enhanced Computer Vision with Microsoft Kinect Sensor: A Review," *IEEE Transactions on Cybernetics* (43:5), pp. 1318–1334.

Han, Q., and Cho, D. 2016. "Characterizing the technological evolution of smartphones: insights from performance benchmarks," in *Proceedings of the 18th International Conference on Electronic Commerce*, , ICEC '16, ACM.

Han, Q., and Ferreira, P. 2014. "The Role of Peer Influence in Churn in Wireless Networks," in *Proceedings of the 2014 International Conference on Social Computing*, , SocialCom '14, ACM.

Han, Q., Ferreira, P., and Costeira, J. P. 2015. "Asymmetric Peer Influence in Smartphone Adoption in a Large Mobile Network," in *Proceedings of the 14th International Conference on Mobile Business*, , ICMB '15, AIS.

Haubl, G., and Trifts, V. 2000. "Consumer Decision Making in Online Shopping Environments: The Effects of Interactive Decision Aids," *Marketing Science* (19:1), pp. 4–21.

Herlocker, J. L., Konstan, J. A., Terveen, L. G., and Riedl, J. T. 2004. "Evaluating Collaborative Filtering Recommender Systems," *ACM Transactions on Information Systems* (22:1), pp. 5–53.

Hervas-Drane, A. 2015. "Recommended for you: The effect of word of mouth on sales concentration," *International Journal of Research in Marketing* pp. 207–218.

Hinz, O., and Eckert, J. 2010. "The Impact of Search and Recommendation Systems on Sales in Electronic Commerce," *Business & Information Systems Engineering* (2:2), pp. 67–77.

Hirano, K., and Imbens, G. W. 2005. "The Propensity Score with Continuous Treatments," in *Applied Bayesian Modeling and Causal Inference from Incomplete-Data Perspectives*, A. Gelman and X.-L. Meng (eds.), John Wiley & Sons, Ltd, pp. 73–84.

Hoernig, S. 2007. "On-net and off-net pricing on asymmetric telecommunications networks," *Information Economics and Policy* (19:2), pp. 171–188.

Hosanagar, K., Fleder, D., Lee, D., and Buja, A. 2014. “Will the Global Village Fracture Into Tribes? Recommender Systems and Their Effects on Consumer Fragmentation,” *Management Science* (60:4), pp. 805–823.

Huang, Y., Zhu, F., Yuan, M., Deng, K., Li, Y., Ni, B., Dai, W., Yang, Q., and Zeng, J. 2015. “Telco Churn Prediction with Big Data,” in *Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data*, , SIGMOD ’15, ACM.

Hui, S., and Bradlow, E. 2012. “Bayesian Multi-Resolution Spatial Analysis with Applications to Marketing,” *Quantitative Marketing and Economics* (10:4), pp. 419–452.

Hui, S., Bradlow, E., and Fader, P. 2009. “The Traveling Salesman Goes Shopping: The Systematic Deviations of Grocery Paths from TSP-Optimality,” *Marketing Science* (28:3), pp. 566–572.

Hui, S. K., Huang, Y., Suher, J., and Inman, J. 2013. “Deconstructing the “First Moment of Truth”: Understanding Unplanned Consideration and Purchase Conversion Using In-Store Video Tracking,” *Journal of Marketing Resarch* (50:4), pp. 445–462.

Hung, S.-Y., Yen, D. C., and Wang, H.-Y. 2006. “Applying data mining to telecom churn management,” *Expert Systems with Applications* (31:3), pp. 515 – 524.

Hwang, H., Choi, B., and Lee, M.-J. 2005. “A model for shelf space allocation and inventory control considering location and inventory level effects on demand,” *International Journal of Production Economics* (97:2), pp. 185–195.

Hwang, H., Jung, T., and Euiho, S. 2004. “An LTV model and customer segmentation based on customer value: a case study on the wireless telecommunication industry,” *Expert Systems with Applications* (26:2), pp. 181–188.

Imai, K., King, G., and Stuart, E. A. 2008. “Misunderstandings between experimentalists and observationalists about casual inference,” *Journal of the Royal Statistical Society* (171:2), pp. 481–502.

Imai, K., and Van Dyk, D. 2004. “Causal Inference With General Treatment Regimes: Generalizing the Propensity Score," *Journal of American Statistical Association* (99:467), pp. 854–866.

Inman, J., Winer, R. S., and Ferraro, R. 2009. "The Interplay Between Category Characteristics, Customer Characteristics and Customer Activities on In-Store Decision Making," *Journal of Marketing* (73:1), pp. 19–29.

ITU 2015. "The World in 2015: ICT Facts and Figures," .
URL `http://www.itu.int/en/ITU-D/Statistics/Documents/facts/ICTFactsFigures`

Jannach, D., Resnick, P., Tuzhilin, A., and Zanker, M. 2016. "Recommender Systems - Beyond Matrix Completion," *Communications of the ACM* (59:11), pp. 94–102.

Kalyanam, K., Lal, R., and Wolfram, G. 2010. "Future Store Technologies and Their Impact on Grocery Retailing," in *Retailing in the 21st Century: Current and Future Trends*, M. Krafft and M. K. Mantrala (eds.), Springer Berlin Heidelberg, pp. 141–158.

Kempe, D., Kleinberg, J., and Tardos, E. 2003. "Maximizing the Spread of Influence through a Social Network," in *Proceedings of the 9th ACM SIGKDD International Conference on Knowledge discovery and data mining*, , KDD '03, ACM.

Kim, H.-S., and Yoon, C.-H. 2004. "Determinants of subscriber churn and customer loyalty in the Korean mobile telephony market," *Telecommunication Policies* (28:9), pp. 751–765.

Klein, B., and Wright, J. D. 2007. "The Economics of Slotting Contracts," *The Journal of Law & Economics* (50:3), pp. 421–454.

Komiak, S. Y. X., and Benbasat, I. 2006. "The Effects of Personalization and Familarity on Trust and Adoption of Recommendation Agents," *MIS Quarterly* (30:4), pp. 941–960.

Kowtsch, T., and Maass, W. 2010. "In-store consumer behavior: How mobile recommendation agents influence usage intentions, product purchases, and store preferences," *Computers in Human Behavior* (26:4), pp. 697–704.

Kumar, V., and Shah, D. 2004. "Building and sustaining profitable customer loyalty for the 21st century," *Journal of Retailing* (80:4), pp. 317–329.

La Fond, T., and Neville, J. 2010. "Randomization Tests for Distinguishing Social Influence and Homophily Effects," in *Proceedings of the 19th International Conference on World Wide Web*,

, WWW '10, New York, NY, USA: ACM.

Larson, J., Bradlow, E., and Fader, P. 2005. "An exploratory look at supermarket shopping paths," *International Journal of Research in Marketing* (22:4), pp. 395–414.

Lawrence, R., Almasi, G., Kotlyar, V., Viveros, M., and Duri, S. 2001. "Personalization of Supermarket Product Recommendations," *Data Mining and Knowledge Discovery* (5:1), pp. 11–32.

Lemmens, A., and Croux, C. 2006. "Bagging and Boosting Classification Trees to Predict Churn," *Journal of Marketing Research* (43:2), pp. 276–286.

Lemmens, A., and Gupta, S. 2013. "Managing Churn to Maximize Profits," Working Paper 14-020, Havard Business School.

Lewis, K., Gonzalez, M., and Kaufman, J. 2012. "Social Selection and peer influence in an online social network," *Proceedings of National Academy of Science* (109:1), pp. 68–72.

Li, L., Chen, J., and Raghunathan, S. 2014. "Recommender System Rethink: Implications for an Electronic Marketplace with Competing Manufacturers," Tech. rep., Working Paper.

Liang, D., Ma, Z., and Qi, L. 2011. "Service quality and customer switching behavior in China's mobile phone service sector," *Journal of Business Research* (66:8), pp. 1161–1167.

Liang, T.-P., Lai, H.-J., and Ku, Y.-C. 2006. "Personalized Content Recommendation and User Satisfaction: Theoretical Synthesis and Empirical Findings," *Journal of Management Information Systems* (23:3), pp. 45–70.

Lieber, E., and Syverson, C. 2012. "Online versus Offline Competition," in *The Oxford Handbook of the Digital Economy*, M. Peitz and J. Waldfogel (eds.), Oxford University Press, pp. 189–223.

Lim, A., Rodrigues, B., and Zhang, X. 2004. "Metaheuristics with Local Search Techniques for Retail Shelf-Space Optimization," *Management Science* (50:1), pp. 117–131.

Linoff, G. S., and Berry, M. J. A. 2011. *Data Mining Techniques: For Marketing, Sales, and Customer Relationship Management*, Indianapolis, IN: Wiley Publishing.

Lofgren, M. 2005. "Winning at the First and Second Moments of Truth: An Explorative Study," *Management Service Quality* (15:1), pp. 102–115.

Lu, S., Xiao, L., and Ding, M. 2016. "A Video-Based Automated Recommender (VAR) System for Garments," *Marketing Science* (35:3), pp. 484–510.

Matos, M., Ferreira, P., and Belo, R. 2015. "Target the ego or target the group: Evidence from a randomized experiment in proactive churn management," Tech. rep., SSRN.
URL `http://ssrn.com/abstract=2591671`

Matos, M., Ferreira, P., and Krackhardt, D. 2014. "Peer Influence in the Diffusion of the iPhone 3G over a Large Social Network," *MIS Quarterly* (38:4), pp. 1103–1133.

Matos, M., Ferreira, P., Smith, M. D., and Telang, R. 2016. "Culling the Herd: Using Real-World Randomized Experiments to Measure Social Bias with Known Costly Goods," *Management Science* (62:9), pp. 2563–2580.

McCarthy, E. J. 1964. *Basic marketing, a managerial approach*, Homewood, IL: Richard D. Irwin, Inc.

McDonald, D. W. 2003. "Ubiquitous Recommendation Systems," *IEEE Computer* (36:10), pp. 111–112.

McNee, S. M., Riedl, J., and Konstan, J. A. 2006. "Being Accurate is Not Enough: How Accuracy Metrics have hurt Recommender Systems," in *Proceedings of the ACM CHI Conference on Human Factors in Computing Systems*, , CHI '06, ACM.

McPherson, M., Smith-Lovin, L., and Cook, J. 2001. "Birds of a Feather: Homophily in Social Networks," *Annual Review of Sociology* (27), pp. 415–444.

Mehta, N., Rajiv, S., and Srinivasan, K. 2003. "Price Uncertainty and Consumer Search: A Structural Model of Consideration Set Formation," *Marketing Science* (22:1), pp. 58–84.

Miller, B. N., Albert, I., Lam, S. K., Konstan, J. A., and Riedl, J. 2003. "MovieLens Unplugged: Experiences with an Occasionally Connected Recommender Systems," in *Proceedings of the International Conference on Intelligent User Interfaces*, , IUI '03, ACM.

Mogan, S., and Winship, C. 2014. *Counterfactuals and Causal Inference: Methods and Principles for Social Research (Analyitcal Methods for Social Research), 2nd edition*, Cambridge University Press, New York NY.

Montgomery, A., Li, S., Srinivasan, K., and Liechty, J. C. 2004. "Modeling Online Browsing and Path Analysis Using Clickstream Data," *Marketing Science* (23:4), pp. 579–596.

Mozer, M., Wolniewicz, R., Grimes, D., Johnson, E., and Kaushansky, H. 2000. "Predicting Subscriber Dissatisfaction and Improving Retention in the Wireless Telecommunications Industry," *IEEE Transactions on Neural Networks* (11:3), pp. 690–696.

Murray, C. C., Talukdar, D., and Gosavi, A. 2010. "Joint Optimization of Product Price, Display Orientation and Shelf-Space Allocation in Retail Category Management," *Journal of Retailing* (86:2), pp. 125–136.

Murthi, B., and Sarkar, S. 2003. "The Role of the Management Sciences in Research on Personalization," *Management Science* (49:10), pp. 1344–1362.

Neslin, S. A., Gupta, S., Kamakura, W., Lu, J., and Mason, C. H. 2006. "Defection Detection: Measuring and Understanding the Predictive Accuracy of Customer Churn Models," *Journal of Marketing Research* (43:2), pp. 204–211.

Ngai, E., Xiu, L., and Chau, D. 2009. "Application of data mining techniques in customer relationship management: A literature review and classification," *Expert Systems with Applications* (36:2, Part 2), pp. 2592 – 2602.

Nielsen, R., and Sheffield, J. 2009. "Matching with Time-Series Cross-Sectional Data," in *The 26th Annual Society for Political Methodology Summer Conference*, , Polmeth.

Nitzan, I., and Libai, B. 2011. "Social Effects on Customer Retention," *Journal of Marketing* (75:6), pp. 24–38.

Nordfalt, J., Grewal, D., Roggeveen, A. L., and Hill, K. M. 2014. "Insights from In-Store Marketing Experiments," in *Review of Marketing Research*, vol. 11 of *Shopper Marketing and the Role of In-Store Marketing*, , Emerald, vol. 11 of *Shopper Marketing and the Role of In-Store Marketing*, chap. 6, pp. 127–146.

Oestreicher-Singer, G., and Sundararajan, A. 2012a. "Recommendation Networks and the Long Tail of Electronic Commerce," *MIS Quarterly* (36:1), pp. 65–83.

Oestreicher-Singer, G., and Sundararajan, A. 2012b. "The Visible Hand? Demand Effects of

Recommendation Networks in Electronic Markets," *Management Science* (58:1), pp. 1963–1981.

Panniello, U., Gorgoglione, M., and Palmisano, C. 2009. "Comparing Pre-filtering and Post-filtering Approach in a Collaborative Contextual Recommender System: An Application to E-Commerce," in *Proceedings of the 10th Conference on Electronic Commerce and Web Technologies*, , EC-Web '09, Springer.

Panniello, U., Gorgoglione, M., and Tuzhilin, A. 2016a. "In CARSs We Trust: How Context-Aware Recommendations Affect Customer's Trust and Other Business Performance Measures of Recommender Systems," *Information Systems Research* (27:1), pp. 182–196.

Panniello, U., Hill, S., and Gorgoglione, M. 2016b. "The impact of profit incentives on the relevance of online recommendations," *Electronic Commerce Research and Applications* (20), pp. 87–104.

Papke, L., and Wooldridge, J. M. 1996. "Econometric Methods for Fractional Response Variables with An Application to 401(k) Plan Participation Rates," *Journal of Applied Econometrics* (11), pp. 619–632.

Park, S.-H., and Han, S. P. 2013. "From Accuracy to Diversity in Product Recommendations: Relationships Between Diversity and Customer Retention," *International Journal of Electronic Commerce* (18:2), pp. 51–72.

Pathak, B., Garfinkel, R., Gopal, R. D., Venkatesan, R., and Yin, F. 2010. "Empirical Analysis of the Impact of Recommender Systems on Sales," *Journal of Management Information Systems* (27:2), pp. 159–188.

Phadke, C., Uzunalioglu, H., Mendiratta, V., Kushnir, D., and Doran, D. 2013. "Prediction of Subscriber Churn Using Social Network Analysis," *Bell Labs Technical Journal* (17:4), pp. 63–75.

Philips, H., and Bradshaw, R. 1993. "How customers actually shop: Customer interaction with the point of sale," *Journal of the Market Research Society Market Research Society* (35:1), pp. 51–62.

Prawesh, S., and Padmanabhan, B. 2014. "The Most Popular News Recommender: Count Amplification and Manipulation Resistance," *Information Systems Research* (25:3), pp. 569–589.

Radhakrishnan, M., Eswaran, S., Misra, A., Chander, D., and Dasgupta, K. 2016. "IRIS: Tapping Wearable Sensing to Capture In-Store Retail Insights on Shoppers," in *Proceedings of the 2016 IEEE International Conference on Pervasive Computing and Communications*, , PerCom '16, IEEE.

Resnick, P., and Varian, H. R. 1997. "Recommender Systems," *Communications of the ACM* (40:3), pp. 56–58.

Rigby, D. 2011. "The Future of Shopping," *Harvard Business Review* (89:12), pp. 65–76.

Rosenbaum, P., and Rubin, R. 1983. "The central role of the propensity score in observational studies for causal effects," *Biometrika* (70:1), pp. 41–55.

Rosenblum, M., and van der Laan, M. J. 2009. "Using Regression Models to Analyze Randomized Trails: Asymptotically Valid Hypothesis Tests Despite Incorrectly Specified Models," *Biometrics* (65:3), pp. 937–945.

Rosset, S., Neumann, E., Eick, U., Vatnik, N., and Idan, Y. 2002. "Customer Lifetime Value Modeling and Its Use for Customer Retention Planning," in *Proceedings of the 8th ACM SIGKDD international Conference on Knowledge discovery and data mining*, , KDD '02, ACM.

Russell, R. A., and Urban, T. L. 2010. "The location and allocation of products and product families on retail shelves," *Annals of Operations Research* (179:1), pp. 131–147.

Sam Hui, Y. H., J. Jeffrey Inman, and Suher, J. 2013. "Estimating the Effect of Travel Distance on Unplanned Spending: Applications to Mobile Promotion Strategies," *Journal of Marketing* (77:1), pp. 1–16.

Schafer, J. B., Konstan, J., and Riedl, J. 1999. "Recommender Systems in E-Commerce," in *Proceedings of the 1st ACM Conference on Electronic Commerce*, , EC '99, ACM.

Schafer, J. B., Konstan, J. A., and Riedl, J. 2001. "E-Commerce Recommendation Applications," *Data Mining and Knowledge Discovery* (5:1), pp. 115–153.

Schwartz, B. 2015. *The Paradox of Choice: Why More is Less*, New York, NY: Harper Perennial.

Seiler, S. 2013. "The impact of search costs on consumer behavior: A dynamic approach," *Quantitative Marketing and Economics* (11:2), pp. 155–203.

Senecal, S., and Nantel, J. 2004. "The influence of online product recommendations on consumers' online choices," *Journal of Retailing* (80:2), pp. 159–169.

Shaffer, G., and Zhang, Z. J. 2002. "Competitive One-to-One Promotions," *Management Science* (48:9), pp. 1143 – 1160.

Shalizi, C., and Thomas, A. 2011. "Homophily and Contagion Are Generically Confounded in Observational Social Network Studies," *Sociological Methods and Research* (40:2), pp. 211–239.

Shankar, V., Inman, J., Mantrala, M., Kelley, E., and Rizley, R. 2011. "Innovations in Shopper Marketing: Current Insights and Future Research Issues," *Journal of Retailing* (87:1), pp. 29–42.

Sheehan, K. B., and Hoy, M. G. 2000. "Dimensions of Privacy Concern Among Online Consumers," *Journal of Public Policy & Marketing* (19:1), pp. 62–73.

Shehu, E., Prostka, T., Schmidt-Stolting, C., Clement, M., and Blomeke, E. 2014. "The influence of book advertising on sales in the German fiction book market," *Journal of Cultural Economics* (38:2), pp. 109–130.

Simonson, I. 2005. "Determinants of Customer's Responses to Customized Offers: Conceptual Framework and Research Opportunities," *Journal of Marketing* (69:1), pp. 32–45.

Smith, M., and Brynjolfsson, E. 2003. "Consumer Decision-Making at an Internet Shopbot: Brand Still Matters," *Journal of Industrial Economics* (49:4), pp. 541–558.

Soltani, A. 2015. "Privacy trade-offs in retail tracking," Tech@FTC.
URL `https://www.ftc.gov/news-events/blogs/techftc/2015/04/privacy-trade-o`

Steglich, C., Snijders, T., and Pearson, M. 2010. "Dynamic Networks and Behavior: Separating Selection from Influence," *Sociological Methodology* (40:1), pp. 329–393.

Stigler, G. J. 1961. "The Economics of Information," *The Journal of Political Economy* (69:3), pp. 213–225.

Stilley, K., Inman, J. J., and Wakefield, K. 2010. “Spending on the Fly: Mental Budgets, Promotions, and Spending Behavior,” *Journal of Marketing* (74:3), pp. 34–47.

Stokmans, M., and Hendrickx, M. 1994. “The attention paid to new book releases on a display table,” *Poetics* (22:3), pp. 185–197.

Tan, T., Netessine, S., and Hitt, L. M. 2016. “Is Tom Cruise Threatened? An Empirical Study of the Impact of Product Variety on Demand Concentration,” Working Paper No. 2016/81/TOM/ACGRE, INSEAD.
URL `http://ssrn.com/abstract=2188354`

Tang, Z., Hu, Y. J., and Smith, M. D. 2008. “Gaining Trust Through Online Privacy Protection: Self-Regulation, Mandatory Standards, or Caveat Emptor,” *Journal of Management Information Systems* (24:4), pp. 153–173.

Tintarev, N., and Masthoff, J. 2007. “A Survey of Explanation in Recommender Systems,” in *Proceedings of the 23rd International Conference on Data Engineering Workshop*, , ICDEW ’07, IEEE.

Tintarev, N., and Masthoff, J. 2015. “Designing and Evaluating Explanations for Recommender Systems,” in *Recommender Systems Handbook*, F. Ricci, L. Rokach, B. Shapira, and P. B. Kantor (eds.), Springer US, chap. 8, pp. 353–382.

Uncles, M. D., Dowling, G. R., and Hammond, K. 2003. “Customer Loyalty and Customer Loyalty Programs,” *Journal of Consumer Marketing* (20:4), pp. 294–316.

Underhill, P. 2009. *Why We Buy: The Science of Shopping (3rd edition)*, New York, NY: Simon & Schuster.

Ursu, R. 2016. “The Power of Rankings: Quantifying the Effect of Rankings on Online Consumer Search and Purchase Decisions,” Working paper, University of Chicago.

Valdimar Sigurdsson, H. S., and Foxall, G. 2009. “Brand Placement and Consumer Choice: An In-Store Experiment,” *Journal of Applies Behavior Analysis* (42:3), pp. 741–745.

Valenzuela, A., and Raghubir, P. 2009. “Position-based beliefs: The center-stage effect,” *Journal of Consumer Psychology* (19:2), pp. 185–196.

Valenzuela, A., and Raghubir, P. 2015. "Are Top-Bottom Inferences Conscious and Left Right Inferences Automatic? Implications for Shelf Space Positions," *Journal of Experimental Psychology: Applied* (21:3), pp. 224–241.

Valenzuela, A., Raghubir, P., and Mitakakis, C. 2013. "Shelf space schemas: Myth or reality?" *Journal of Business Research* (66:7), pp. 881–888.

van Nierop, E., Fok, D., and Franses, P. H. 2008. "Interaction Between Shelf Layout and Marketing Effectiveness and Its Impact on Optimizing Shelf Arrangements," *Marketing Science* (27:6), pp. 1065–1082.

Verbeke, W., Dejaeger, K., Martens, D., Hur, J., and Baesens, B. 2012. "New insights into churn prediction in the telecommunication sector: A profit driven data mining approach," *European Journal of Operational Research* (218:1), pp. 211 – 229.

Verbeke, W., Martens, D., Mues, C., and Baesens, B. 2011. "Building comprehensible customer churn prediction models with advanced rule induction techniques," *Expert Systems with Applications* (38:3), pp. 2354–2364.

Walter, F. E., Battiston, S., Yildirim, M., and Schweitzer, F. 2012. "Moving recommender systems from on-line commerce to retail stores," *Information Systems and e-Business Management* (10:2), pp. 367–393.

Wang, W., and Benbasat, I. 2005. "Trust in and Adoption of Online Recommendation Agents," *Journal of Association for Information Systems* (6:3), pp. 72–101.

Wang, W., and Benbasat, I. 2007. "Recommendation Agents for Electronic Commerce: Effects of Explanation Facilities on Trusting Beliefs," *Journal of Management Information Systems* (23:4), pp. 217–246.

Wedel, M., and Kannan, P. 2016. "Marketing Analytics for Data-Rich Environments," *Journal of Marketing* (80:6), pp. 97–121.

Wedel, M., and Pieters, R. 2008. "A Review of Eye-Tracking Research in Marketing," in *Review of Marketing Research*, vol. 4, N. K. Malhotra (ed.), Emerald Publishing, vol. 4, chap. 5, pp. 123–147.

Wei, C.-P., and Chiu, I.-T. 2002. “Turning telecommunications call details to churn prediction: a data mining approach,” *Expert Systems with Applications* (23:2), pp. 103–112.

Wingfield, N. 2017. “Amazon?s Living Lab: Reimagining Retail on Seattle Streets,” *New York Times* (February 12, 2017).

Yang, H. 2013. “Targeted Search and the Long Tail Effect,” *RAND Journal of Economics* (44:4), pp. 733–756.

Yang, M.-H. 2001. “An efficient algorithm to allocate shelf space,” *European Journal of Operational Research* (131:1), pp. 107–118.

Yang, M.-H., and Chen, W.-C. 1999. “A Study on Shelf Space Allocation and Management,” *International Journal of Production Economics* (60-61), pp. 309–317.

Zhang, T., Agarwal, R., and Lucas, H. C. 2011. “The Value of IT-Enabled Retailer Learning: Personalized Product Recommendations and Customer Store Loyalty in Electronic Markets,” *MIS Quarterly* (35:4), pp. 859–881.

Zhang, X., Li, S., Burke, R. R., and Leykin, A. 2014. “An Examination of Social Influence on Shopper Behavior Using Video Tracking Data,” *Journal of Marketing* (78:5), pp. 24–41.