

 Open access • Journal Article • DOI:10.1177/0022243719833631

P2V-MAP: Mapping Market Structures for Large Retail Assortments:

— [Source link](#) 

Sebastian Gabel, Daniel Guhl, Daniel Klapper

Published on: 06 Jun 2019 - Journal of Marketing Research (SAGE Journals)

Topics: Big data

Related papers:

- [Identifying Customer Needs from User-Generated Content](#)
- [Marketing Analytics for High-Dimensional Assortments](#)
- [Product Choice with Large Assortments: A Scalable Deep-Learning Model](#)
- [Classification of retail products: from probabilistic ranking to neural networks](#)
- [B2B analytics in the airline market: Harnessing the power of consumer big data](#)

Share this paper:    

View more about this paper here: <https://typeset.io/papers/p2v-map-mapping-market-structures-for-large-retail-2wpk9h0qam>

One-to-One Marketing in Grocery Retailing

DOCTORAL THESIS

to acquire the academic degree of
doctor rerum politicarum
(Doctor of Economics and Management Science)

submitted to the

School of Business and Economics of
Humboldt-Universität zu Berlin

by

Dipl.–Phys. Dipl.–Wirtsch.–Phys. Sebastian Gabel

President of Humboldt-Universität zu Berlin:

Prof. Dr.–Ing. Dr. Sabine Kunst

Dean of the School of Business and Economics:

Prof. Dr. Daniel Klapper

Reviewers:

1. Prof. Dr. Daniel Klapper
2. Prof. Dr. Stefan Lessmann

Date of Colloquium: June 26, 2019

Abstract

Target marketing has evolved from targeting large customer segments to one-to-one marketing. Retailers personalize promotions based on customer-level transaction data, search engines optimize results based on users' past queries, and online advertisers take into account users' online behavior. Personalizing their marketing mix to individual customers increases firm sales and profits, and improves customer satisfaction. Customers benefit from better services, more relevant offers, and tailored communication.

The increasing volume, variety, and velocity of data that firms collect open up promising opportunities for better target marketing. Nonetheless, research on one-to-one marketing with a focus on retailing is scarce in academic literature. The two main reasons are that the target marketing approaches proposed by researchers do not scale to the size of typical retail applications and that data regarding one-to-one marketing remain locked within retailers and marketing solution providers.

This dissertation (1) develops new descriptive, predictive, and prescriptive marketing models for automated target marketing that are based on representation learning and deep learning and (2) studies the models' impact in real-life applications.

First, this thesis shows that representation learning is capable of analyzing market structures at scale without requiring any human interaction. The proposed approach to visualizing market structures is fully automated and superior to existing mapping methods that are based on the same input data, such as multidimensional scaling and principal component analysis. Understanding product relationships and competition is the basis for any target marketing application, so this study is a necessary first step toward new deep learning models for predictive and prescriptive marketing analytics.

Based on these results, the thesis then proposes a scalable, nonparametric model that predicts product choice for the entire assortment of a large retailer. The model is based on a custom deep neural network architecture, that is specifically designed for the application to time series purchase data from retailer loyalty programs. The end-to-end neural network outperforms benchmark methods for predicting customer purchases and generalizes out-of-sample. Coupon policies based on the

proposed model lead to substantially higher revenue lifts than policies based on the benchmark models.

The remainder of the thesis then studies a real-time offer engine that is based on the proposed models. A close collaboration with a leading German grocery retailer and its target marketing solution provider makes it possible to evaluate the business impact of one-to-one marketing in a real-life application. The comparison of personalized promotions to non-targeted promotions shows that sophisticated machine learning systems for automated one-to-one marketing increase redemption rates, revenues, and profits. A study of customer responses to personalized price promotions within the retailer's loyalty program reveals that personalized marketing also increases loyalty program usage. This illustrates how targeted price promotions can be integrated smoothly into loyalty programs.

In summary, this thesis is highly relevant for both researchers and practitioners. The new deep learning models outperform existing approaches to market structure analysis and predicting customer decisions. This facilitates more scalable and efficient one-to-one marketing. The models' flexibility makes them well suited to deal with large-scale data sets from heterogeneous data sources. In addition to the methodological contribution, this research offers several pertinent implications for promotion management and one-to-one marketing.

Zusammenfassung

Target Marketing hat sich von der Targetierung großer Kundensegmente zum One-to-One-Marketing weiterentwickelt. Einzelhändler personalisieren Werbeaktionen auf der Grundlage ihrer Transaktionsdaten, Suchmaschinen optimieren Suchergebnisse basierend auf vergangenen Nutzeranfragen und Firmen verwenden das beobachtete Nutzerverhalten, um Online Werbung zu personalisieren. Durch die Personalisierung des Marketing-Mixes auf Kundenebene können Unternehmen ihren Umsatz und Gewinn steigern und gleichzeitig die Kundenzufriedenheit verbessern. Kunden ihrerseits profitieren von nützlicheren Dienstleistungen, relevanteren Angeboten und maßgeschneiderter Kommunikation.

Das zunehmende Volumen und die Vielfalt gesammelter Daten sowie die hohe Beobachtungsfrequenz eröffnen vielversprechende Möglichkeiten für besseres Target Marketing. Dennoch existieren in der akademischen Fachliteratur kaum Forschungsergebnisse zu One-to-One-Marketing, die auf Anwendungen im Einzelhandel ausgerichtet sind. Zu den Hauptgründen zählen, dass die von Forschern vorgeschlagenen Ansätze für Target Marketing nicht auf die Größe typischer Einzelhandelsanwendungen skalieren und dass die Verfügbarkeit relevanter Daten auf Händler und Marketing-Systemanbieter beschränkt ist.

Die vorliegende Dissertation (1) entwickelt neue deskriptive, prädiktive und präskriptive Marketingmodelle für automatisiertes Target Marketing, die auf Representation Learning und Deep Learning basieren und (2) untersucht die Auswirkungen dieser Marketingansätze in Praxisanwendungen.

Im ersten Schritt zeigt die Arbeit, dass Representation Learning in der Lage ist, skalierbar Marktstrukturen zu analysieren, ohne dass menschliches Eingreifen erforderlich ist. Der vorgeschlagene Ansatz zur Visualisierung von Marktstrukturen ist vollständig automatisiert und vorhandenen Methoden wie multidimensionaler Skalierung und Hauptkomponentenanalysen, die auf denselben Eingabedaten basieren, überlegen. Produktbeziehungen und Wettbewerb abzubilden, ist die Grundlage für jede Target Marketing-Anwendung. Diese Studie ist somit ein notwendiger erster Schritt in Richtung neuer Deep Learning-Ansätze für prädiktive und präskriptive Marketingmodelle.

Auf Basis dieser Erkenntnisse entwickelt die Arbeit anschließend ein skalierbares, nichtparametrisches Modell, das Produktwahl auf Konsumentenebene für alle Produkte im Sortiment großer Einzelhändler vorhersagt. Das Modell basiert auf einer neuartigen Deep Learning-Architektur, die im Rahmen dieser Arbeit gezielt für die Anwendung auf Zeitreihen-Transaktionsdaten aus Kundenbindungsprogrammen entwickelt wurde. Das vorgeschlagene neuronale Netzwerk generalisiert über die Stichprobe hinaus und übertrifft die Vorhersagekraft existierender Benchmarks. Die unter Nutzung des Modells abgeleiteten Coupons führen im Vergleich zu Coupons aus Benchmark-Modellen zu signifikanten Umsatzsteigerungen.

Die Dissertation untersucht anschließend eine Coupon-Engine, die auf den entwickelten Modellen basiert. Eine Zusammenarbeit mit einem führenden deutschen Lebensmitteleinzelhändler und einem Anbieter von Target Marketing-Anwendungen ermöglicht es, die wirtschaftlichen Konsequenzen von Target Marketing in der Praxis zu untersuchen. Der Vergleich personalisierter Werbeaktionen mit Massenmarketing belegt, dass der Einsatz moderner Machine Learning-Verfahren Coupon-Einlösungsraten, Umsätze und Gewinne steigern kann. Eine Analyse der Kundenreaktionen auf personalisierte Coupons im Rahmen des Kundenbindungsprogrammes des Einzelhändlers zeigt außerdem, dass personalisiertes Marketing Systemnutzung erhöht. Diese Erkenntnisse illustrieren, wie Händler Target Marketing und Kundenbindungsprogrammen effizient und nahtlos kombinieren können.

Zusammenfassend ist die vorliegende Dissertation sowohl für Forscher als auch für Praktiker relevant. Die entwickelten Deep Learning-Modelle übertreffen die Leistungsfähigkeit existierender Ansätze zur Marktstrukturanalyse und zur Vorhersage von Konsumentenverhalten und bilden die Grundlage für skalierbarere und effizientere Marketingpersonalisierung. Die Universalität der Modelle erlaubt zudem die Nutzung heterogener Datenquellen. Neben methodischer Beiträge bietet diese Arbeit relevante Implikationen für effizientes Promotion-Management und One-to-One-Marketing im Einzelhandel.

Acknowledgments

Most of all, I thank my advisor Prof. Dr. Daniel Klapper for the opportunity to work with him. His guidance, feedback, and support were indispensable during the last three years. I am grateful to Prof. Dr. Stefan Lessmann and Prof. Dr. Lutz Hildebrandt for their valuable feedback on my research.

My gratitude belongs to the whole team of SO1 GmbH, especially Raimund Bau and Dieter Ebbers, who made my research possible. I thank Dr. Daniel Guhl with whom I worked on several research projects. I enjoyed our discussions and learned much during our joint work. I am grateful to my other co-authors, the PhD students at the School of Business and Economics, and the research assistants for their support.

Lastly, I thank my family. My parents who supported me in all my pursuits. And especially my loving wife Karoliina. Without her support my research would not have been possible.

Contents

1	Introduction	1
2	P2V-MAP: Mapping Market Structures for Large Retail Assortments	11
3	Cross-Category Product Choice: A Scalable Deep-Learning Model	13
3.1	Introduction	14
3.2	Related Literature	16
3.2.1	Product Choice Modeling	16
3.2.2	Coupon Personalization and Targeting	18
3.2.3	Deep Learning Applications in Marketing	19
3.3	The Proposed Cross-Category Product Choice Model	20
3.3.1	Overview	20
3.3.2	Model Architecture	21
3.3.3	Model Calibration	24
3.3.4	Discussion	25
3.4	Simulation Setup	26
3.4.1	Stage 1: Category Purchase Incidence	26
3.4.2	Stage 2: Product Choice	27
3.4.3	Simulation Calibration	27
3.5	Model Evaluation Based on Simulated Data	29
3.5.1	Aggregate Prediction Performance	30
3.5.2	Prediction Performance Decomposition	31
3.5.2.1	Product Choice Dynamics	31
3.5.2.2	Coupon Effects	32
3.5.2.3	Time Dynamics and Inventory Effects	34
3.5.2.4	Identifying Product Category Structure	35
3.5.3	Performance Gains for Coupon Optimization	37
3.6	Model Evaluation Based on Empirical Data	40
3.6.1	Data	40
3.6.2	Evaluation Results	41
3.7	Conclusion	44
3.8	Appendix	46

4	The Impact of Personalization on Coupon Performance	61
4.1	Introduction	62
4.2	Related Work and Contribution	64
4.2.1	Targeted Couponing–Methods and Effects	64
4.2.2	Promotional Effects and Their Drivers	66
4.3	Setup	66
4.3.1	Data Set	67
4.3.2	Targeting Policy	68
4.3.3	Descriptive Analysis of Redemption Rates	69
4.4	Approach	70
4.4.1	Model	70
4.4.2	Analysis Overview	73
4.5	Part 1: Redemption Rate Analysis	74
4.5.1	Redemption Rate Uplift Through Brand Personalization	74
4.5.2	Drivers of Redemption Rate Uplift	76
4.6	Part 2: Revenue and Profit Analysis	80
4.6.1	Revenue and Profit Uplift Through Personalization	80
4.6.2	Revenue and Profit Uplift for Varying Degrees of Personalization	83
4.7	Conclusion	85
4.8	Appendix	89
5	The Impact of Personalized Coupons on Loyalty Program Usage	95
5.1	Introduction	96
5.2	Related Work and Contribution	98
5.2.1	Customer Responses to LPs	98
5.2.2	LP Rewards	99
5.2.3	Effects of Personalized Coupons	100
5.3	Loyalty Program Setup and Data	101
5.4	Empirical Analysis	103
5.4.1	Descriptive Analysis	103
5.4.2	Duration Analysis	106
5.4.3	Discussion of Findings and Cost Evaluation	112
5.4.4	Supporting Insights from an Online Experiment	115
5.5	Conclusion	116
5.6	Appendix	119
	References	124

List of Tables

1.1	Four essays on one-to-one marketing in grocery retailing.	9
3.1	Aggregate prediction performance (simulation).	30
3.2	True and estimated discount elasticities (average across products).	33
3.3	Time series correlation scores for model predictions.	35
3.4	Coupon optimization results.	39
3.5	Summary statistics: data sets for empirical application.	41
3.6	Aggregate prediction performance (empirical application).	43
3.7	Binary cross-entropy loss regression analysis.	44
3.8	Adam optimizer parameters (PyTorch).	46
3.9	Simulation parameters.	46
3.10	Additional metrics for aggregate prediction performance (simulation).	53
3.11	Benchmarking results (test set) for four nested DNN models.	54
3.12	Coupon policy analysis.	57
3.13	Category characteristics for loyalty card data set.	58
4.1	Summary of data set statistics.	68
4.2	Results for binary logit models to explain targeting policy.	69
4.3	Descriptive analysis of redemption rates.	70
4.4	Estimation results for redemption models.	72
4.5	Meta-regression estimation results.	78
4.6	Comparison of RTO engine revenue with mass market promotions.	82
4.7	Comparison of RTO engine profit with mass market promotions.	82
4.8	Model results based on random data.	90
4.9	Model results based on targeted data.	90
4.10	References and descriptive statistics for explanatory variables.	91
4.11	Correlation between brand and category characteristics.	92
4.12	Items for impulse and ability-to-stockpile scales.	93
4.13	Ten highest and lowest ranked categories for the impulse and stockpiling scores.	93
5.1	Prior studies on LP rewards.	100
5.2	Effect of lagged rewards on $\log(IUT)$	105
5.3	Variable operationalization and descriptive statistics.	107
5.4	PHM model estimation results.	109

List of Figures

1.1	Four essays on one-to-one marketing in grocery retailing.	3
3.1	Neural network architecture for the proposed product choice model.	22
3.2	Simulated product market shares.	28
3.3	Category incidence probability histograms for different inventory levels.	29
3.4	Time-series prediction (hold-out set).	32
3.5	Difference between true simulated elasticities and predicted elasticities.	34
3.6	Heat-map of the product embedding W_H	36
3.7	t-SNE projection of the product embedding W_H	37
3.8	Category incidence probability histograms for different inventory levels for three product categories.	47
3.9	Data split for hold-out evaluation.	49
3.10	Loss curves for training, validation, and test data.	50
3.11	Probability scatter plots for six product categories.	51
3.12	Probability scatter plots for six products in category 1.	52
3.13	Test loss, correlation metric and product embedding W_H (products 1 to 40).	55
3.14	t-SNE Projection of the DNN Product Embedding W_H	59
3.15	Observed purchase rates as a function of predicted probabilities. . .	60
4.1	Analysis overview.	73
4.2	Approach for redemption rate comparison with stratified discount distributions.	75
4.3	Redemption probability comparison.	76
4.4	Revenue and profit per customer/coupon vs. circular.	83
4.5	Revenue and profit for varying degrees of personalization.	85
5.1	LP usage cycle.	102
5.2	Longitudinal and cross-sectional data on print and reward events. .	104
5.3	Median customer inter-usage time.	105
5.4	Estimated baseline hazard, survivor and density functions.	110
5.5	Effect of reward stock values on relative hazard.	111
5.6	Effect of reward values on relative hazards.	113
5.7	Costs of products to achieve a certain relative hazard value.	114

5.8	Perceived value of LP rewards.	117
-----	--	-----

Abbreviations

AIC	Akaike information criterion
ANOVA	Analysis of variance
AUC	Area under curve
BFGS	Broyden-Fletcher-Goldfarb-Shanno
CI	Confidence interval
CPG	Consumer packaged goods
DL	Deep learning
DNN	Deep neural network
Est.	Parameter estimate
FSI	Freestanding insert
GLM	Generalized linear model
GLMM	Generalized linear mixed model
IUT	Inter-usage time
KL	Kullback-Leibler
LL	Log-likelihood
LP	Loyalty program
MCMC	Markov chain Monte Carlo
ML	Machine learning
MTurk	Mechanical Turk
P2V-MAP	Product2Vec-Map
PCA	Principal component analysis
PHM	Proportional hazard model
ReLU	Rectified linear unit
RMSE	Root mean square error
RTO	Real-time offer
SD	Standard deviation
SE	Standard error
Sig.	Parameter significance
t-SNE	(Barnes-Hut) t-distributed stochastic neighbor embedding
US	United States

1 | Introduction

Target marketing has a long history in marketing research and practice. Marketers tailor marketing activities to their customers' characteristics and preferences with the goal to improve the firm's position in the target segment (Palmatier and Sridhar, 2017). Technological progress has made it feasible to continuously reduce the size of target segments. One-to-one marketing, the most granular form of target marketing, tailors the firm's marketing mix to each customer (Peppers and Rogers, 1997; Peppers et al., 1999; Shaffer and Zhang, 2002). Firms track the purchases of individual shoppers, observe service usage in real time, collect rich behavioral and attitudinal data to learn about customer preferences, and use digital channels to personalize marketing communications (Wedel and Kannan, 2016; Kannan et al., 2017). These developments open up new, exciting opportunities for one-to-one marketing. Customer-centric firms understand that customer heterogeneity necessitates tailoring marketing efforts to individuals, and that doing so leads to substantially higher profits (Rust and Verhoef, 2005; Fader, 2012). The personalization of marketing activities is based on the statistical analysis of customer data that yields predictions about customer responses to marketing activities such as promotions and advertising (Arora et al., 2008). This enables firms to deliver "the right content to the right person at the right time, to maximize immediate and future business opportunities" (Tam and Ho, 2006, p. 867).

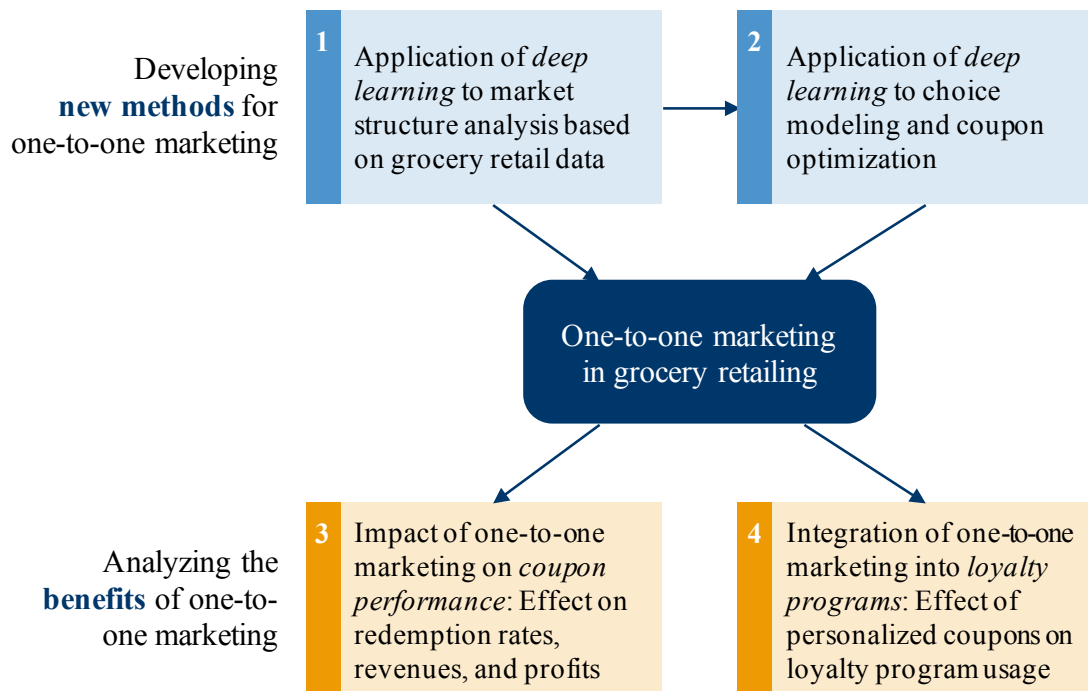
Pioneers of personalization and one-to-one marketing can be found across different industries, both offline and online (Aguirre et al., 2015). The search engines Google and Bing analyze past queries and contextual information to produce faster and better search results (Arora et al., 2008). Facebook targets online advertisements based on the users' online behavior (Goldfarb and Tucker, 2011), and publishers such as nytimes.com recommend articles based on the users' interests (Arora et al., 2008). Retailers collect vast amounts of customer-level data, which they use to analyze customer purchasing habits (Blattberg et al., 2008; Bradlow et al., 2017). Amazon.com and Barnes & Noble, for example, provide personalized product recommendations (Montgomery and Smith, 2009). In grocery retailing, the availability of customer data, especially those obtained through loyalty programs (LP), and the targeting engines offered by solution providers (e.g., dunnhumby or Catalina Marketing) promise to leverage the potential of promotion personalization

(Rowley, 2005; Guillot, 2016). The brick-and-mortar retailers Target and Safeway's, for example, tailor circulars and coupons to the customers' shopping histories (Bleier et al., 2018).

Academic literature confirms that personalization in marketing benefits both customers and firms (for an overview, see Vesanen, 2007). Research has shown that target marketing increases firm profits and revenues by offering more relevant products to customers and differentiating prices according to customers' preferences and willingness-to-pay (Arora et al., 2008; Rossi et al., 1996; Rust and Verhoef, 2005; Zhang and Wedel, 2009). Differentiation helps firms to gain a competitive advantage (Murthi and Sarkar, 2003) which, in turn, might make it possible to charge higher prices (Vesanen, 2007). Personalized offers yield higher recall and are more effective in influencing customer decisions (Tam and Ho, 2006; Tucker, 2014). From the customer perspective, marketing personalization can simplify decisions (Murthi and Sarkar, 2003) and contribute to increased satisfaction and loyalty (Ansari and Mela, 2003). Practitioners confirm these findings. They list higher response rates, more relevant customer interactions, higher conversion rates, better differentiation against competitors, and higher loyalty as key benefits of adopting personalized marketing strategies (eMarketer, 2016). Industry studies also report increased profitability, larger shopping baskets, higher purchase frequencies and improved customer retention (Lindsay, 2014; Hawkins, 2012).

Despite these benefits for firms and customers, research on one-to-one marketing in retailing is scarce in academic literature. The two main reasons are that the target marketing approaches proposed by researchers do not scale to the size of typical retail applications and that data regarding one-to-one marketing remain locked within retailers and marketing solutions providers. The **goal of this thesis** is to (1) develop and validate new descriptive, predictive, and prescriptive marketing models for automated one-to-one marketing that are explicitly designed for the application in retailing and (2) study the impact of these models in real-life applications. The close collaboration with a leading German grocery retailer and its target marketing solution provider sets the context for this research: The proposed models use large market basket data and loyalty card data sets as input for modeling market structure, predicting customer choices, and deriving policies for personalized coupons. Retailers can apply the models directly to raw transaction data. This eliminates the need for extensive data preparation and assumptions about category delineation and cross-product effects. The implementation is based on modern deep learning frameworks for automated inference, so the models can be easily modified and extended. The loyalty card data also contains the customers' responses to targeted coupons—the targeting is based on the models proposed in this thesis—so this opens up an exciting opportunity to study the impact of one-to-one marketing on customer behavior and coupon performance.

Figure 1.1. Four essays on one-to-one marketing in grocery retailing.



Note: Best viewed in color.

The thesis is a collection of four essays on one-to-one marketing (see Figure 1.1). The first two essays have a **methodological focus** in that they propose new, scalable approaches to one-to-one marketing that are particularly suitable for retailing. Models that have been used for target marketing in the past break down when the number of customers or promoted alternatives increase (Naik et al., 2008). Consider, for example, that Walmart collects data on billions of shopping baskets every year and stocks up to 150,000 distinct products in its brick-and-mortar stores, along with more than one million products on walmart.com (Walmart, 2005, 2016). As a consequence, most studies on one-to-one marketing adopt the perspective of a brand by focusing on a small number of products, product categories, and customers. In their seminal work on target marketing, Rossi et al. (1996) show how to model, measure, and optimize price discounts in a brand choice setting, highlighting that household purchase histories are valuable to manufacturers for optimizing coupon profitability, according to their study in a single product category. Zhang and Wedel (2009) study retailer-customized checkout coupons in online and offline stores across two product categories. Johnson et al. (2013) propose a model that adds the dimension of timing to target marketing and apply it to optimize coupons for four brands in a single product category. Dubé and Misra (2017) propose a machine learning approach for price personalization and apply it to (business-to-business) subscription pricing at an online recruiting company. In applying such approaches to full product assortments, retailers would have to implement hundreds of complex

models, one for each of their products or product categories. Given the level of sophistication, each model necessitates careful data preparation and calibration (e.g., data pruning, choice set definition, master data collection). Also, consider that the quality of the results can depend directly on the assumptions made during data preparation, for example, on how retailers delineate product categories. Even if retailers make the “right” assumptions, it remains unclear how they should combine the category-level results in a global one-to-one marketing policy across all categories, and whether modeling categories independently might adversely affect outcomes. Reducing the size of data on the other hand—in terms of customers and/or products (Zanutto and Bradlow, 2006)—is not viable in the context of one-to-one marketing (Jacobs et al., 2016).

A possible alternative for researchers is to borrow tools from machine learning (ML) to complement traditional econometric techniques (Einav and Levin, 2014). Big (retail) data offers the potential of understanding causal effects of marketing instruments to a greater extent (Sudhir, 2016), and ML methods are a promising approach to combine a variety of heterogeneous data sources, such as purchase histories, responses to past promotions, click-stream data, and browsing histories to inform one-to-one marketing, explicitly accounting for complex interaction effects (Bradlow et al., 2017). In online retailing, for instance, firms use collaborative filtering algorithms to predict customers’ next purchases by analyzing their purchase histories (Mild and Reutterer, 2003; Liu et al., 2009; Jannach et al., 2011). Although applications of ML have proven useful in practice, it is important to note their limitations: Simple response models that are frequently used in targeting engines to predict (binary) outcomes such as clicks or purchases (Chapelle et al., 2015) can be used to predict the purchase probability conditional on marketing interventions (e.g., targeted coupons), but they typically fail to account for relationships between alternatives, for example, competition and complementarity between products. A challenge for count-based approaches such as collaborative filtering algorithms is incorporating customer characteristics and marketing variables (e.g., coupons). Research that addresses these shortcomings is only beginning to emerge. Jacobs et al. (2016) extend latent Dirichlet allocation (Blei et al., 2003) to allow using customer characteristics in predicting the purchases of 11,783 customers for 394 products. The generative model proposed by Ruiz et al. (2018) jointly predicts the purchase probabilities for 5,590 products and 11,783 customers accounting for product prices and the sequential decision process of shoppers. Although this research is a promising first step toward more scalable approaches for modeling customer choices, neither model predicts individual responses to marketing actions. This is a prerequisite for one-to-one marketing applications. And consider too, that both applications are still small compared to the vast size of typical retail data sets, so more work is required that explicates how models from other disciplines such as

data science, ML, and natural language processing can be applied appropriately to marketing problems (Montgomery and Smith, 2009; Chintagunta et al., 2016; Sudhir, 2016).

An especially promising research direction to tackle these challenges is deep learning (DL). The universality of DL and its applicability to large-scale data sets is well established and deep learning pioneers, such as Google and Facebook, have illustrated its usefulness in marketing applications (Covington et al., 2016; Park et al., 2018). DL is a general-purpose learning procedure that is capable of processing data in their raw form. It can be applied to a variety of different big data sources (including images, video, audio, speech, and text), and features very good predictive performance (LeCun et al., 2015). In contrast to classic ML approaches that require the manual design of features as model input, DL models utilize increasing amounts of computational resources and data in automatically learning representations (or features) from raw input data. These representations capture intricate structures in large data sets without requiring manual effort or domain-specific expert knowledge (LeCun et al., 2015). DL has achieved performances close to the level of humans in image (Krizhevsky et al., 2012), face (Taigman et al., 2014), and speech recognition (Hinton et al., 2012), while producing promising results for natural language understanding (Collobert et al., 2011), question answering (Bordes et al., 2014), language translation (Sutskever et al., 2014), and automatic image annotation (Vinyals et al., 2015). It is therefore not surprising that researchers expect DL to play a central role in marketing applications in the future (Kannan et al., 2017). Yet, little work in academic marketing literature has addressed deep neural networks (Wedel and Kannan, 2016). This thesis proposes new DL models that can be applied to one-to-one marketing, specifically personalized coupons, so it is a first step toward target marketing based on deep learning. The contribution of this thesis includes novel deep learning architectures that are specifically designed to model market structures and cross-category product choice based on retail data.

The **first essay** shows that DL is capable of analyzing market structures at scale without requiring any human interaction and ex ante assumptions about product relationships (e.g., product categorization). Understanding product relationships and competition is the basis for any target marketing application, so this study is a necessary foundation toward new predictive and prescriptive marketing models based on DL. The study shows that the proposed approach for visualizing market structures is superior to existing mapping methods (e.g., multidimensional scaling, principal component analysis) that are based on the same input data. A comprehensive simulation study contributes to a better understanding of how DL models capture product attributes, product relationships and market structures. The in-depth comparisons with the results of state-of-the-art methods for analyzing product relationships such as multivariate probit models (Manchanda et al., 1999)

and mixed logit models (Train, 2009) provide evidence that the learned representations approximate the true market structure well. The application of this approach using data collected at a leading German grocery retailer underlines its usefulness and generates novel findings that are relevant to promotion management and assortment-related decisions.

The **second essay** directly builds on the results of the first essay in that it proposes a DL model for cross-category product choice. Accurately predicting what customers will likely buy on their next shopping trip is at the core of efficient one-to-one marketing. Prior research has had great success in modeling the choices of individual customers within a single or across a small number of selected product categories. This study deals with product choice across the entire assortment of grocery retailers. Such retailers typically operate hundreds of product categories and handle millions of transactions per day. The dimensionality and scale of the problem require new methods for efficient product choice modeling. The essay proposes a scalable, nonparametric model that predicts product choice for the entire assortment of a large retailer. The model is based on a custom deep neural network architecture, that is specifically designed for the application to time series purchase data from retail LPs. The model inputs customer-level purchase histories and coupon assignments to predict purchases of individual customers. The proposed neural network builds on the results of the first essay in that it first estimates latent product representations using market basket data. It then combines purchase histories, marketing mix variables, and additional meta data to predict product choice. Retailers can apply the model directly to raw loyalty card data without making assumptions about product relationships (e.g., category structure) and extensive data preparation (e.g., product attributes, choice sets). This paper provides an in-depth evaluation of the model's performance in a simulation study and verifies its prediction performance using empirical data. The simulation study explicates that the end-to-end neural network generalizes out-of-sample, achieves a higher prediction accuracy than state-of-the-art benchmark methods, and is more scalable than classic econometric approaches, both in the number of products and the volume of historical purchase data. The model captures own- and cross-product coupon effects, adjusts the predicted probabilities for consumption dynamics, and automatically learns market structure. The study illustrates the value of improved product choice prediction in the context of real-time offer (RTO) engines for grocery coupons. The deep neural network facilitates more effective and efficient coupon policies for one-to-one marketing. Coupon personalization based on our model achieves substantially higher revenues compared to the baseline prediction methods. The application of the deep neural network to loyalty card data from the same retailer studied in the first essay confirms the superiority of the proposed model over benchmark solutions. Retailers can easily extend the model input, so the

proposed product choice model offers practical value for other retail analytics tasks that require quantifying how marketing decisions impact business performance.

The last two essays are motivated by the practitioners' view of target marketing, in that they evaluate the **impact of one-to-one marketing** on coupon performance and LP usage. Research on this aspect of one-to-one marketing and coupon personalization in retailing is limited. In a quasi-experiment, Venkatesan and Farris (2012) find that coupon exposure and redemption have positive effects on trip incidence and revenues. Sahni et al. (2016) evaluate the revenue effect of personalized email promotions in a field experiment at an online ticket resale platform. Osuna et al. (2016) study the performance of checkout coupons, targeted such that eligibility to receive the coupons depend on the households' purchase histories. When it comes to LPs, researchers have demonstrated that personalization can act as a loyalty-building mechanism (Bijmolt et al., 2011; Meyer-Waarden, 2007; Verhoef, 2003), but the link between personalized promotions and LPs are understudied in academia. Research is needed that explicates how LPs can "be combined or even integrated with other marketing-mix instruments" (Bijmolt and Verhoef, 2017, p. 161) and what the effects of such combinations are. This thesis uses the loyalty card data provided by a German grocery retailer and its marketing solution provider to study how targeted coupons impact coupon performance and loyalty card usage. The RTO engine that personalizes the coupons uses the DL models proposed in the first two essays, so the close collaboration with the retailer and its target marketing solution provider opens up a unique opportunity to evaluate the impact of one-to-one marketing in a real-life application.

The **third essay** studies a rich retail data set that comprises market basket data, loyalty card data, and customer responses to 12 million personalized coupons across 1,116 brands in 115 product categories. For almost 1 million coupons, the brand and the discount were randomized, so the exogenous variation pertaining to both coupon dimensions facilitates an unbiased measurement of the effect of decision variables on customer responses. This makes it possible to study the impact of personalization on coupon performance in policy simulations. To this end, the essay quantifies the effect of targeting on redemptions, revenues, and profits. Personalized coupons achieve an average redemption rate of 4.2%. This equals an increase of 64.0% relative to non-targeted coupons. One-to-one marketing increases revenues by up to 182.2% and profits by up to 111.8% compared to non-targeted mass marketing policies (e.g., circulars). The impact of targeting on coupon effectiveness varies significantly across categories and brands, and much of the variance can be explained by brand and category characteristics, such as brand loyalty, price position, and purchase frequency. This research helps retailers to use targeting engines more efficiently. The results underline the benefits of sophisticated systems for automated one-to-one marketing that are based on DL and allow retailers

to compare the costs associated with implementing personalization engines to the financial benefits that such systems offer. Beyond the analysis of the impact of personalization on retail performance metrics, this essay provides an effective framework for measuring the success of personalized price promotions.

The **fourth essay** studies the effect of one-to-one marketing on user behavior within an LP, specifically customer responses to personalized coupons produced by an RTO engine. Prior research has extensively studied the revenue and profit implications of LPs, yet little is known about how the LP design and LP rewards affect LP usage. The link between the LP design and LP usage is becoming increasingly important for practitioners. Retailers send too many communications, it takes too long to earn points for rewards, and the rewards provided in LPs are often not relevant, so LP usage is at an all-time low (Fruend, 2017). A rich longitudinal data set makes it possible to use a latent-class proportional hazard model (PHM) to analyze how personalized coupons and classic LP rewards affect LP usage. The results indicate that the effect of personalized coupons is stronger and that the two reward types complement each other. An Amazon Mechanical Turk (MTurk) experiment confirms the main findings and conclusions derived from the hazard model and contributes to a better generalizability of the findings. This essay provides empirical evidence for how RTO engines and targeted price promotions can be integrated smoothly into LPs, to drive customer retention and LP usage. At the same time, it contributes to the understanding of the interaction between two of the most fundamental aspects of retail management: LP design and price promotions. For practitioners, the essay outlines practical insights that are useful in increasing LP usage.

In summary, the research presented in this thesis is relevant to both researchers and practitioners. Table 1.1 provides an overview of the thesis' key contributions, main findings, and research methodology. The proposed DL models outperform existing approaches to market structure analysis and predicting customer choices. Their flexibility and scalability make them well suited for the application to large-scale data sets and a variety of heterogeneous input data. This thesis outlines how to use these models in one-to-one marketing. In addition to the methodological contribution, this research offers several pertinent implications for promotion management and one-to-one marketing. The results indicate that retailers can use RTO engines to increase coupon redemption rates, revenues, and profits. The essays provide generalizable insights that can guide retailers in implementing and using RTO engines, and illustrate the value of tightly integrating personalized coupons with LPs.

Table 1.1. Four essays on one-to-one marketing in grocery retailing.

	Chapter 2	Chapter 3	Chapter 4	Chapter 5
Title	P2V-MAP: Mapping Market Structures for Large Retail Assortments	Cross-Category Product Choice: A Scalable Deep-Learning Model	The Impact of Personalization on Coupon Performance	The Impact of Personalized Coupons on Loyalty Program Usage
Contribution	Fully automated, scalable method for mapping market structures; foundation for deep learning applications in marketing; application to assortment decisions	Scalable method for predicting purchases and deriving coupon policies; fully automated cross-category choice model	Impact of one-to-one targeting on coupon performance (redemption rates, revenues, and profits); insights for design and usage of targeting engines	Extend research on LP design, LP rewards and integration of personalized promotions/RTO engines into loyalty programs
Key findings	Improved mapping accuracy (e.g., adjusted mutual information +69.3%); validation of approach in simulation study and empirical application	Improved prediction accuracy; more efficient coupon policies (e.g., revenue +74.0%); validation of approach in simulation study and empirical application	Personalization increases redemption rates (+64.0%), revenues (+182.2%), and profits (+111.8%); uplift varies for brands, categories, and degree of personalization	Positive effect of LP rewards on LP usage; considerable customer heterogeneity; effect of personalized coupons stronger than effect of classic LP rewards
Data	Market basket data; simulated data	Loyalty card data; market basket data; simulated data	Loyalty card data; data on coupon redemptions (targeted and random coupons); MTurk survey data	Longitudinal data on purchases, kiosk usage, and coupon redemptions; MTurk survey data
Approach	Deep learning/neural networks; dimensionality reduction; multivariate probit model; multinomial probit/logit model	Deep learning/neural networks; multivariate probit model; multinomial logit model	Binary logistic regression and (weighted) linear regression random effects models; policy simulations	Latent class PHM; linear regression random effects model; ANOVA
Comments	Published in <i>Journal of Marketing Research</i> Finalist EMAC 2017 Best Paper Award Based on a Doctoral Work	Research cooperation with <i>MIT Sloan School of Management</i>		Research cooperation with <i>ETH Zürich</i>

2 | P2V-MAP: Mapping Market Structures for Large Retail Assortments

Publication

Gabel, S., Guhl, D., and Klapper, D. (2019). P2V-MAP: Mapping Market Structure for Large Retail Assortments. *Journal of Marketing Research* (forthcoming). Available at <https://journals.sagepub.com/doi/10.1177/0022243719833631>.

Abstract

The authors propose a new, exploratory approach for analyzing market structures that leverages two recent methodological advances in natural language processing and machine learning. They customize a neural network language model to derive latent product attributes by analyzing the co-occurrences of products in shopping baskets. Applying dimensionality reduction to the latent attributes yields a two-dimensional product map. This method is well-suited to retailers because it relies on data that are readily available from their checkout systems and facilitates their analyses of cross-category product complementarity, in addition to within-category substitution. The approach has high usability because it is automated, scalable, and does not require a priori assumptions. Its results are easy to interpret and update as new market basket data are collected. The authors validate their approach both by conducting an extensive simulation study and by comparing their results with those of state-of-the-art, econometric methods for modeling product relationships. The application of this approach using data collected at a leading German grocery retailer underlines its usefulness and provides novel findings that are relevant to assortment-related decisions.

3 | Cross-Category Product Choice: A Scalable Deep-Learning Model

Sebastian Gabel, Artem Timoshenko

Abstract

Automated coupon personalization requires predictions of how coupons affect customer purchasing behavior. We propose a scalable product choice model that inputs individual purchase histories and coupon assignments to predict purchase decisions across the entire assortment of a retailer. The model is based on a custom deep neural network architecture. We rely on convolutional filters, bottleneck layers, and weight sharing to efficiently capture cross-product relationships and dynamic consumption patterns. Retailers can apply the model directly to loyalty card transaction data, without predefined categories or product attributes. We provide a detailed evaluation of the model in a simulation. Our model achieves a higher prediction accuracy than the baseline machine learning methods. We demonstrate that the model infers coupon effects and adjusts the predicted probabilities for recent purchases. Using the proposed model for coupon personalization leads to substantially higher revenue lifts. We verify the prediction performance by applying the model to transaction data with experimental coupon assignment variation provided by a large retailer.

Keywords

product choice model, neural networks, deep learning, cross-category choice, retail analytics, coupon optimization

3.1 Introduction

Retailers provide coupons to promote products and categories, stimulate incremental purchases, and improve customer retention (Blattberg and Neslin, 1990). In 2018, US retailers distributed 256.5 billion coupons for consumer packaged goods (CPG) alone, and consumers redeemed over 1.7 billion coupons with a combined face value of \$2.7 billion (NCH Marketing Services, 2019).

Providing coupons is costly for retailers. For example, freestanding inserts (FSI) account for about 90% of the CPG coupons and the estimated cost per redemption is \$.35 (Biafore, 2016). Moreover, customers often redeem coupons for products for which they would have been willing to pay the regular price (Forrester, 2017).

To increase redemption rates and coupon profitability, retailers adopt coupon personalization solutions (Peppers and Rogers, 1997; Fader, 2012). CVS offers personalized coupons at the store entrance through kiosk systems, Food Lion (Ahold Delhaize) provides coupons for the next visit at the checkout, and Whole Foods distributes coupons via its mobile application.

Automated coupon personalization requires a product choice model that predicts how marketing actions affect customer purchasing behavior (Arora et al., 2008). In our conversations with major retailers and solution providers in the US and Europe, practitioners emphasized that implementing such models can be challenging. Current product choice models used for coupon optimization adopt a brand perspective and focus on a single product category (e.g., Rossi et al., 1996; Johnson et al., 2013). Models need careful calibration and require the modeler to delineate categories, prune input data, define choice sets, and collect product attributes. Large retailers such as Walmart handle millions of transactions per day and stock products in over 500 product categories (Walmart, 2005, 2016). Sophisticated by-category product choice models achieve substantially higher prediction accuracies than models that predict responses by-product (e.g., binary response models), but implementing these models for hundreds of categories and maintaining them is hardly feasible. Even if retailers were able to implement hundreds of by-category models in parallel, ignoring cross-category product relationships leads to sub-optimal coupon targeting policies across the full assortment.

Retailers understand that complex choice models can achieve higher targeting efficiency, but the limited scalability and high implementation effort of existing approaches force them to revert to targeting heuristics that allocate coupons based on manually defined scoring rules. The scores aggregate redemption rates and purchase frequencies scaled by the products' prices. Customers then receive coupons for the highest-scoring products. Simple heuristics can improve coupon effectiveness but certainly do not leverage the full potential of data-based personalization.

In this paper, we develop a scalable product choice model that predicts customer-specific purchase likelihoods in response to personalized coupon discounts for the entire assortment. The model is based on a custom deep learning architecture which inputs purchase histories of individual customers and coupon assignments to predict purchase decisions.

The proposed model is highly practical. Retailers can apply the model directly to raw transaction data from loyalty programs. This eliminates the need for extensive data preparation and assumptions about category delineation and cross-product effects. Our customized implementation leverages an established deep learning framework for automated inference, so the model can be easily modified and extended.

To achieve scalability to large product assortments, we keep most of the neural network transformations product-specific and use weight sharing between the neurons (Alain and Bengio, 2014). The parsimonious model architecture has a regularization effect and simplifies model training. We rely on the bottleneck layers to encode relevant cross-product relationships in the hidden layers of the neural network, thereby adjusting the predicted probabilities. For example, the model automatically infers that coupons for Coke and Pepsi have a similar effect on the purchase likelihoods of other soft drinks.

We evaluate the proposed product choice model using simulated and empirical data. We first simulate a retailer with many products across multiple categories. Purchase decisions follow a two-stage process: customers first decide whether to purchase a product from a category (category choice), and then choose products within the selected categories (product choice). We assume customer heterogeneity and category-specific consumption dynamics. Customers receive coupons every time period. Each coupon affects the own-product purchase probability and purchase probabilities of other products in the category.

The simulation study validates that our model accurately predicts purchase probabilities for all products in the assortment and generalizes out-of-sample. We compare the proposed custom neural network to two binary response baselines and conclude that our model achieves superior prediction accuracy. The model successfully approximates own- and cross-product coupon effects and dynamically adjusts the predicted probabilities for customer-specific consumption patterns. It infers the underlying product category structure and accounts for cross-product relationships without a manual *ex ante* definition of categories.

We further use the simulated data to demonstrate the value of the proposed product choice model for coupon personalization. Coupon personalization requires a model to predict purchase probabilities as a function of coupon assignments and an optimization approach to allocate coupons given the predicted effects. We evaluate

coupon personalization approaches with one or five coupons per customer. In both cases, we keep the optimization algorithm constant and vary the underlying product choice models. The higher prediction accuracy of our product choice model leads to approximately 75% larger revenue lifts through coupon personalization. The coupon policy based on our model (1) targets more expensive but less frequently purchased products without sacrificing redemption rates and (2) generates more incremental category purchases.

We finally evaluate the prediction performance of the proposed product choice model using transaction data provided by a leading German grocery retailer. The retailer distributed random coupons to a small fraction of customers. Experimental data allows us to train and evaluate the model without endogeneity concerns. In line with the results obtained from simulated data, our model achieves higher out-of-sample prediction accuracy than the baseline models. The outperformance margins over the reference product choice models are particularly large for observations shortly after a category purchase and observations in categories characterized by smaller interpurchase times.

The proposed product choice model also offers high practical value for retailing problems other than coupon optimization. Potential applications include retail analytics tasks that require quantifying how marketing decisions impact business performance based on purchase data (Hanssens, 2014). For example, offline retailers forecast demand to optimize fulfillment and predict response lifts to improve targeted promotions. Online retailers can leverage our model to optimize product recommendations or personalized landing pages.

Section 3.2 proceeds with a review of related literature. Section 3.3 introduces the proposed product choice model. In Section 3.4, we describe the simulation setup. We use simulated data to evaluate the prediction performance of the proposed model and demonstrate its value for coupon personalization in Section 3.5. Section 3.6 validates the prediction performance using empirical data. We summarize our findings and suggest directions for future research in Section 3.7.

3.2 Related Literature

Our research relates to three streams of literature: product choice modeling, methods for targeting and coupon optimization, and deep learning applications in marketing. We next discuss each of these areas and highlight our respective contributions.

3.2.1 Product Choice Modeling

Product choice models quantify how marketing actions affect business outcomes such as market shares and profits. Predicting the effects of marketing activities

is the basis for efficient resource allocation (Hanssens, 2014). Winer and Neslin (2014) provide a comprehensive overview of the product choice modeling literature.

Traditionally, product choice models estimate purchase decisions for a single product/brand or a category. For example, Fader and Hardie (1996) propose a latent class multinomial logit model to predict customer choices for 56 products within the fabric softener category. The authors represent products as a combination of attributes (e.g., brand, package size) and demonstrate that their model significantly outperforms a model specification with 55 product-specific intercept terms, even though it uses less parameters.

Attribute-based choice models achieve better predictive performance, but require a retailer to maintain comprehensive product attribute data bases and to identify the relevant attributes for each category-level model. Doing this for all products in the retailer’s assortment is a very complex and laborious task. Our proposed product choice model infers product similarities directly from customer-level transaction data. The neural network represents products using low-dimensional vectors (embeddings), and a common product embedding space makes products comparable. This approach does not require manual definitions of product attributes.

Models that study multi-category product choice include the multivariate probit model (Manchanda et al., 1999) and the multivariate logit model (Russell and Petersen, 2000). Multivariate choice models infer product co-occurrence, complementarity and substitution by estimating the covariance structure of purchase decisions across categories from market basket data. The number of possible choice alternatives in the multivariate choice models increases exponentially with the number of product categories, which limits their scalability. For example, Manchanda et al. (1999) and Russell and Petersen (2000) each study four product categories. Our proposed model encodes product relationships within and across categories implicitly in the hidden layers of the deep neural network. This makes simultaneously modeling hundreds of product categories and scaling to the size of typical retail applications possible.

Recently, machine learning approaches for product choice modeling have been gaining more popularity in marketing. For example, Jacobs et al. (2016) propose LDA-X, an extension of latent Dirichlet allocation (Blei et al., 2003), to predict customer-specific purchase probabilities for products in the assortment of an online retailer. LDA-X first infers small-dimensional customer embeddings from the data through Markov chain Monte Carlo (MCMC) and then uses customer embeddings to inform predictions of future purchases. Ruiz et al. (2018) propose SHOPPER to sequentially predict the purchase probabilities for products from multiple product categories given the current content of the shopping cart. SHOPPER describes products through latent attributes (embeddings) that capture product characteris-

tics and product relationships. Both LDA-X and SHOPPER account for customer heterogeneity and are more scalable than classic discrete choice models.

We contribute to machine learning models in marketing in three ways. First, our model is specifically designed to predict individual responses to marketing actions. The model incorporates customer-specific marketing mix variables and customer-level purchase histories. Both effects are important for coupon personalization and other targeting applications. Second, our model scales both to the entire assortment and rich transaction data at a large retailer. The neural network architecture allows parallel implementation and inference via mini-batch gradient descent. For example, MCMC inference for LDA-X takes several days even for small product assortments (e.g., 2,500 products). Our model trains in a few hours with similar hardware specifications. Third, the proposed model has high usability. We provide an implementation of the custom neural network architecture in an established deep learning framework. Retailers can easily modify and extend our model with new data sources. For example, the retailer that provided data for the empirical application in our paper already uses neural networks in supply chain management and is likely to adopt our approach for marketing applications.

3.2.2 Coupon Personalization and Targeting

Our product choice model is motivated by the coupon personalization problem in retail. Coupon personalization and targeting are important topics in marketing research and practice (Bradlow et al., 2017; Grewal et al., 2017). Rossi et al. (1996) propose a model to derive profit maximizing coupon personalization policies and highlight the value of household purchase histories for optimizing coupon profitability. Zhang and Wedel (2009) jointly model purchase incidence, product choice, and quantity decisions in online and offline stores to maximize brand profit through promotion customization. Dubé and Misra (2017) propose a machine learning approach for price personalization and apply it to (business-to-business) subscription pricing at an online recruiting company. Simester et al. (2019b) evaluate the robustness of the machine learning models for targeting direct mail promotions for customer acquisition in retail.

Coupon personalization solutions require a product choice model and an optimization approach. The choice model predicts how different combinations of coupons affect individual purchasing behavior, and the coupon optimization approach allocates coupons given the predicted effects. Our research develops a product choice model that predicts the impact of coupons on purchase probabilities for the entire assortment of a large retailer. We evaluate the performance of the proposed product choice model for coupon personalization by comparing it to reference models and estimate the expected profits of the simulated retailer with the coupons allocated by the optimization approach with different underlying product choice models.

The basis for training and evaluating our proposed product choice model is experimental data. Our simulation and the empirical application assign coupons to customers at random. Random coupon assignment allows training the prediction model without endogeneity concerns. We validate the coupon optimization approaches in the simulation using a randomization-by-policy experimental design (Simester et al., 2019a). In particular, we evaluate coupon personalization by implementing different algorithms to assign coupons to different groups of customers (or equivalently using independent simulation runs).

3.2.3 Deep Learning Applications in Marketing

Our proposed product choice model is based on a neural network. Neural network models have achieved remarkable performance in computer vision and natural language processing applications (LeCun et al., 2015). Marketing researchers have recently started applying deep neural networks to marketing problems.

For example, Liu et al. (2017) develop an approach to automatically extracting content information from online product reviews and predict conversion. Timoshenko and Hauser (2019) propose a deep learning framework that enables firms to identify customer needs from online reviews more efficiently. Zhang and Luo (2018) use deep learning to extract sentiments from photos and reviews posted on Yelp and find that sentiments predict restaurant survival, even after controlling for other covariates. Liu et al. (2018) apply deep convolutional neural networks to social media images with the goal to measure consumers' perception of brands. Gabel et al. (2019) propose a machine learning method based on neural networks to map market structures in grocery retailing based on market basket data.

The properties of deep neural networks make them well-suited for applications to loyalty card data. First, deep learning methods can handle large volumes of training data (Goodfellow et al., 2016). Large retailers process millions of transactions daily, which creates an enormous amount of data for model calibration. Second, deep learning models can effectively operate with high-dimensional inputs. Our proposed model uses purchase histories as one of its inputs. With 2,500 products in the retail assortment and a 30-week history window, the purchase history of a single customer contains 75,000 values. This dimensionality is comparable to 256×256 images often used in computer vision applications (Krizhevsky et al., 2012). The sequential nature of the purchase histories also resembles the structure of words in texts in the natural language processing tasks (Collobert et al., 2011).

Our contribution is a novel deep learning architecture to model cross-category product choice in the context of large product assortments. We provide an in-depth evaluation of the model's performance in a simulation study, and verify the prediction performance using empirical data.

3.3 The Proposed Cross-Category Product Choice Model

3.3.1 Overview

Consider a retail store operating J products. The products may be related both in terms of cross-price elasticities and purchase co-incidence (Manchanda et al., 1999). The relationship between the products is unknown ex ante.

There are I customers who shop at the store. For ease of exposition, we assume that the customers visit the store at every time period (e.g., week, day), but may leave the store without making a purchase. We use a binary vector $\mathbf{b}_{it} = [b_{it0}, \dots, b_{itJ}] \in \{0, 1\}^{J \times 1}$ to denote the purchase decisions of customer i at time t . The binary variable $b_{itj} \in \{0, 1\}$ indicates whether customer i purchased product j at time t . We summarize information about past purchasing behavior of customer i by a purchase history of length T and product purchase frequencies over the entire available time horizon. We denote the purchase history of length T for customer i at time t by $B_{it}^T = [\mathbf{b}_{i,t}, \mathbf{b}_{i,t-1}, \dots, \mathbf{b}_{i,t-T+1}] \in \{0, 1\}^{J \times T}$ and the vector of product-specific purchase frequencies for customer i over the entire customer purchase history available at time t by $B_{it}^\infty = [\bar{b}_{it1}, \dots, \bar{b}_{itJ}] \in [0, 1]^{J \times 1}$.

Customers receive personalized, product-specific coupons before each shopping trip (e.g., by email, through a mobile app, at in-store kiosks). A coupon provides a percent discount on a product at the checkout. We denote personalized coupons by $D_{it} = [d_{it1}, \dots, d_{itJ}] \in [0, 1]^{J \times 1}$, where $d_{itj} \in [0, 1]$ indicates the size of the coupon (i.e., the discount) received by customer i in time t for product j .

We propose a product choice model that predicts probabilities

$$P_{i,t+1} = [p_{i,t+1,1}, \dots, p_{i,t+1,J}] \quad (3.1)$$

that customer i will purchase product j at time $t + 1$ for every product $j \in \{1, \dots, J\}$, given the coupon assignment $D_{i,t+1}$, the purchase history B_{it}^T , the purchase frequencies B_{it}^∞ , and the model parameters θ :

$$P_{i,t+1} = f(D_{i,t+1}, B_{it}^T, B_{it}^\infty; \theta). \quad (3.2)$$

The vector $P_{i,t+1}$ contains the probabilities for the (binary) purchase events for all products j :

$$p_{i,t+1,j} = \mathbb{P}(b_{i,t+1,j} = 1). \quad (3.3)$$

Including both B_{it}^∞ and B_{it}^T as an input to the model serves two purposes. First, the model uses B_{it}^∞ to learn the customer's base preferences, whereas it models

purchasing patterns over time based on B_{it}^T . Separating the information already at the model input simplifies the learning process and speeds up the training. Second, providing B_{it}^∞ in addition to B_{it}^T reduces dimensionality of the input data. Our model could learn B_{it}^∞ directly from B_{it}^T if the window length were set to infinity (i.e., $T = \infty$). However, only recent purchases are relevant to model purchase timing, so we reduce dimensionality by considering a smaller window T and including B_{it}^∞ as a summary of the older purchases.

Figure 3.1 summarizes the proposed model architecture. The model is non-parametric and based on a neural network. Each observation in our model is a customer-time pair (i, t) . For every training sample, the model transforms the inputs (i.e., $D_{i,t+1}$, B_{it}^T , B_{it}^∞) to create product-specific feature maps $\mathbf{z}_{i,t+1,j} \in \mathbb{R}^{K \times 1}$, which are then used to predict the purchase probabilities $p_{i,t+1,j}$ for every product in the assortment:

$$\begin{aligned}
 p_{i,t+1,j} &= p(\mathbf{z}_{i,t+1,j}; \theta_p), \\
 \mathbf{z}_{i,t+1} &= [\mathbf{z}_{i,t+1,1}, \dots, \mathbf{z}_{i,t+1,J}] \in \mathbb{R}^{J \times K}, \\
 \mathbf{z}_{i,t+1} &= Z(D_{i,t+1}, B_{it}^T, B_{it}^\infty; \theta_z).
 \end{aligned} \tag{3.4}$$

The feature maps, $\mathbf{z}_{i,t+1,j}$, summarize information about coupons and information about the customer purchasing behavior into customer-product-specific K -dimensional vectors. The feature maps infer cross-product relationships directly from the transaction data.

3.3.2 Model Architecture

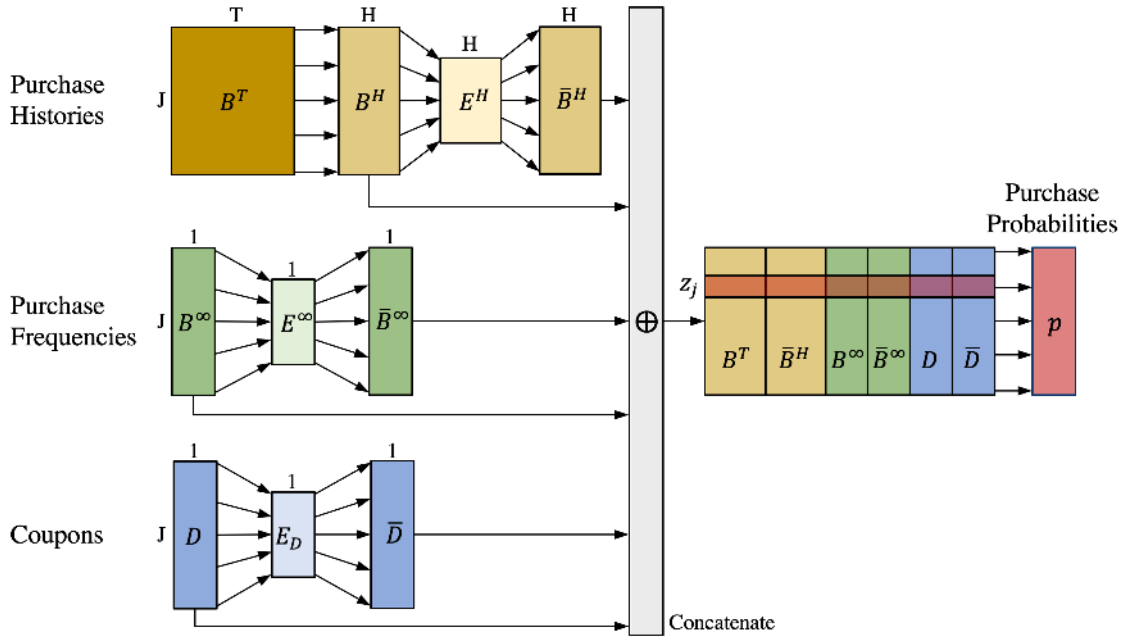
We next describe the details of the model architecture and the calibration of the model. The inputs to the model are a coupon assignment $D_{i,t+1}$, a purchase history B_{it}^T , and product purchase frequencies B_{it}^∞ .

The model first transforms the purchase histories B_{it}^T . In a retail setting, purchase histories B_{it}^T are sparse. We apply convolutional operations with H different real-valued filters $\mathbf{w}_h \in \mathbb{R}^{T \times 1}$:

$$B_{it}^H = [\sigma(B_{it}^T \mathbf{w}_1), \dots, \sigma(B_{it}^T \mathbf{w}_H)] \in \mathbb{R}^{J \times H}, \tag{3.5}$$

where $\sigma(\cdot)$ is a leaky ReLU activation function (Xu et al., 2015):

$$\sigma(x) = \begin{cases} x & \text{for } x \geq 0 \\ 0.2x & \text{for } x < 0. \end{cases} \tag{3.6}$$

Figure 3.1. Neural network architecture for the proposed product choice model.

Note: Best viewed in color.

The filters apply the same transformations to the purchase histories of every product and create H product-specific summary statistics. These summary statistics represent information about recent purchases in a dense form. We calibrate the weights of the time filters using the training data.

Our nonparametric approach for summarizing timing information is more flexible than manually defined transformations of the purchase histories (e.g., weighted averages). This flexibility is important. Retail products have a substantial variation in the interpurchase times. For example, customers typically purchase milk every few days, whereas detergent purchases happen once every few months. Observing a purchase of milk or detergent in period t thus requires different adjustments of the predicted probabilities in period $t + 1$. Our model inputs purchase histories for all products in the assortment, and defining product-specific transformations manually is not feasible. In contrast, the neural network's time filters automatically calibrate these transformations by observing purchasing patterns in the training data.

Purchase frequencies B_{it}^∞ , the aggregated purchase histories B_{it}^H , and coupon assignments $D_{i,t+1}$ are product-specific. We use linear bottleneck layers at the neural network to share information across products. In particular, we apply the following transformations:

$$\begin{aligned}
 B_{it}^\infty &= W_\infty^\top E_{i,t}^\infty & E_{i,t}^\infty &= W_\infty B_{it}^\infty \\
 B_{it}^H &= W_H^\top E_{i,t}^H & \text{with } E_{i,t}^H &= W_H B_{it}^H \\
 \bar{D}_{i,t+1} &= W_d^\top E_{i,t+1}^D & E_{i,t+1}^D &= W_d D_{i,t+1}^H.
 \end{aligned} \tag{3.7}$$

W_d , W_∞ , and W_H are $(L \times J)$ weight matrices with $L \ll J$, and W^\top refers to the transpose of matrix W . The bottleneck layer encodes the inputs into low-dimensional representations $E_{i,t}^\infty$, $E_{i,t}^H$, and $E_{i,t+1}^D$. For example, in Section 3.4 we simulate a retailer with $J = 250$ products, and we estimate the model with $L = 30$. The model infers the weight matrices W_d , W_∞ , and W_H during training.

The bottleneck layers are the basis for modeling cross-product relationships. Consider the following illustrative example. Customer i is indifferent between Coke and Pepsi, and purchases one of the two products when the combined stock of soft drinks at home is low. When the customer purchases Coke or Pepsi at time t , the retailer needs to adjust the estimates of the probabilities that the customer will purchase these soft drinks at time $t + 1$. The adjustment in probabilities is independent of which particular product was purchased in time t . The model recognizes this by creating similar L -dimensional representations of the purchase histories B_{it}^H and the purchase frequencies B_{it}^∞ for the two different scenarios (Coke or Pepsi). These L -dimensional representations are then expanded back to J dimensions to keep further operations at the by-product level.

Applying the bottleneck layer to the discounts $D_{i,t+1}$ captures a different type of relationship between products. Under the assumption of negative price elasticities, a coupon for Coke increases a purchase probability for Coke. Other soft drinks in the soft drinks category exhibit a combination of two effects. A substitution effect decreases their purchase probabilities. On the other hand, the coupon for Coke increases overall consideration of the soft drink category (own-category price effect) increasing the purchase probabilities of all soft drinks, even brands besides Coke. The bottleneck layer allows to capture these cross-product effects of discounts.

We combine the inputs and outputs of the bottleneck layers to create feature maps $\mathbf{z}_{i,t+1}$:

$$\mathbf{z}_{i,t+1} = [\mathbf{1}^{J \times 1}, D_{i,t+1}, \bar{D}_{i,t+1}, B_{it}^\infty, \bar{B}_{it}^\infty, B_{it}^H, \bar{B}_{it}^H] \in \mathbb{R}^{J \times K} \quad (3.8)$$

where $K = 2H + 5$. Combining the inputs and outputs of the layer is a standard method to improve the predictive performance of the neural networks (Orhan and Pitkow, 2017). We input the feature maps $\mathbf{z}_{i,t+1,j}$ to a softmax layer to predict purchase probabilities $P_{i,t+1} = [p_{i,t+1,1}, \dots, p_{i,t+1,J}]$ for every product in the assortment:

$$p_{i,t+1,j} = \frac{\exp\{\theta_p \mathbf{z}_{i,t+1,j}\}}{1 + \exp\{\theta_p \mathbf{z}_{i,t+1,j}\}}. \quad (3.9)$$

The feature maps $\mathbf{z}_{i,t+1,j}$ summarize relevant information about the customer purchasing behavior and the coupon assignment from the inputs, and the softmax

layer uses $\mathbf{z}_{i,t+1,j}$ as input to predict the purchase probability for customer i and product j at time t . The parameters θ_p are shared between the products.

The functional form of the softmax layer is similar to a binary logit model, but they are conceptually different. Traditional binary logit models assume category-specific weights and variation in the product attributes. The product attributes are defined by the researchers. Our model encodes product differences and cross-product effects in the feature maps \mathbf{z} and keeps the weights shared between all the products across categories. The feature maps \mathbf{z} are inferred by the model from the transaction and coupon assignment data.

3.3.3 Model Calibration

The parameters of the model are the time filters \mathbf{w}_h , bottleneck layer parameters W_d , W_∞ , and W_H , and the parameters of the softmax layer θ_P :

$$\theta = (\theta_z; \theta_p), \quad \theta_z = (\mathbf{w}_{h=1..H}; W_d; W_\infty; W_H). \quad (3.10)$$

We calibrate the parameters by minimizing the binary cross-entropy loss

$$\theta^* = \underset{\theta}{\operatorname{argmin}} \sum_{i=1}^I \sum_{j=1}^J \sum_{t=1}^T L(b_{i,t+1,j}, \hat{p}_{i,t+1,j}), \quad (3.11)$$

with

$$L(b_{i,t+1,j}, \hat{p}_{i,t+1,j}) = -b_{i,t+1,j} \log(\hat{p}_{i,t+1,j}) - (1 - b_{i,t+1,j}) \log(1 - \hat{p}_{i,t+1,j}). \quad (3.12)$$

We use the adaptive moment estimation (Adam; Kingma and Ba, 2014) algorithm with mini-batches to optimize the parameters. Adam is a gradient descent method that computes automatic, adaptive learning rates for each parameter of the model to improve learning stability and speed. We provide a complete specification of the optimization algorithm in Appendix 3.8.

The proposed neural network model architecture incorporates two constraints on the parameters to facilitate faster model convergence and prevent overfitting. We first assume the weights at the bottleneck layer decoder to be the transpose of the encoder parameters. For example, we estimate $\bar{D}_{i,t+1} = W_d^\top E_{i,t+1}^D = W_d^\top W_d D_{i,t+1}$, where W_d^\top is a transpose of the weight matrix W_d . The tied weights constraint helps to reduce the number of model parameters and serves as a regularization technique (Alain and Bengio, 2014).

Similarly, we assume tied weights θ_P . The softmax layer applies to product-specific feature maps $\mathbf{z}_{i,t+1,j}$, but the parameters θ_P are shared between the products. This weight sharing is possible because the feature maps \mathbf{z} encode purchase information, including cross-product effects.

3.3.4 Discussion

Our proposed neural network architecture provides a flexible functional form to closely approximate customers' purchasing behavior. The model incorporates information about the purchase histories and current discounts to make customer- and time-specific predictions for every product in the assortment. The standalone parameters of the model have no behavioral or economic interpretation, but the model effectively predicts purchasing behavior required for efficient coupon targeting. For example, in Section 3.5 we demonstrate that the model is capable of adjusting predicted probabilities to account for consumption dynamics and cross-product relationships.

The neural network architecture also makes the model computationally tractable and scalable. We optimize the parameters of the model using a gradient descent algorithm with mini-batches. Training the model in mini-batches allows parallel computing and not having all training data in memory. The proposed neural network architecture allows to efficiently compute gradients via back-propagation. Training deep neural networks is therefore feasible even with a large number of customers I and alternatives J (Covington et al., 2016).

One important characteristic of the deep neural network is that it can be easily extended to incorporate additional information relevant for targeting. For example, retailers can leverage information about the timing of the shopping trip, information about the location of the store, or customer demographic variables. Additional information can also include unstructured data such as product reviews (Archak et al., 2011) or images (Zhang and Luo, 2018). These data can be preprocessed by additional (or even pretrained) neural network layers and added to the feature maps $\mathbf{z}_{i,t+1,j}$ by concatenation:

$$\mathbf{z}_{i,t+1,j}^* = [\mathbf{z}_{i,t+1,j} I_{itj}]. \quad (3.13)$$

This extension increases the number of parameters θ_P but the optimization of the model stays tractable.

The neural network can also be trained in stages. Retailers often have rich market basket data with no customer identifiers. Lack of customer purchase histories limits the ability to target. However, our model can leverage these data to better identify cross-product relationships. In particular, the unlabeled market basket data can be used to train product embeddings (Gabel et al., 2019), and the model can initialize the bottleneck layer parameters with the embeddings. Initialization with pretrained parameters improves the prediction performance of the neural network models and helps to achieve faster convergence (Bengio et al., 2007).

3.4 Simulation Setup

The proposed deep neural network aims to approximate customer purchasing behavior and predict future purchases. We use a simulation study to evaluate the performance of the model in a controlled environment. We draw on previous research in marketing to design the simulation study (Manchanda et al., 1999; McFadden, 1974; Fader and Hardie, 1996). A key benefit of using a simulation is that the true purchase probabilities and the parameter of the data generating process are known. We can thus better evaluate the model’s performance and decompose performance gains.

We simulate a retailer with I customers and an assortment of J products. The products are grouped into C product categories of equal size. Customers visit the store every period, and make purchase decisions in two stages. The customers first decide whether to buy a product in a category and then choose one product in each of the selected categories (Neslin et al., 2009).

The purchase probability for customer i and product j (in category c) at time t is given by

$$p_{itjc} = p_{itc}^{(1)} \cdot p_{ijt}^{(2)}, \quad (3.14)$$

where $p_{itc}^{(1)}$ is the category purchase incidence probability (Section 3.4.1) and $p_{ijt}^{(2)}$ is the product choice probability conditional on the category incidence (Section 3.4.2).

3.4.1 Stage 1: Category Purchase Incidence

We model the category incidence as a multivariate probit model (Manchanda et al., 1999). Customer i ’s utility of a category c purchase incidence depends on the customer-specific base preference, the coupon assignment in the category, and the current inventory:

$$u_{itc} = \gamma_c + \gamma_{ic} + \gamma_{ic}^p \bar{d}_{itc} + \gamma_c^{Inv} Inv_{ic}^t + \varepsilon_{itc}. \quad (3.15)$$

Here, $\gamma_c + \gamma_{ic}$ is the (customer-specific) base utility, \bar{d}_{itc} is the average coupon discount in the category, and Inv_{ic}^t is customer i ’s inventory for category c at time t . Assuming that the random noise has a standard normal distribution, $\varepsilon_{itc} \sim N(0, 1)$, the purchase incidence probability becomes

$$p_{itc}^{(1)} = \mathbb{P}(y_{itc} = 1) = \Phi(u_{itc}), \quad (3.16)$$

where Φ is the cumulative density function of the standard normal distribution and y_{itc} indicates the category purchase incidence, that is the purchase of any product j in C :

$$y_{itc} = \mathbb{1} \left(\sum_{j \in C} b_{itj} > 0 \right). \quad (3.17)$$

Customers are characterized by the latent taste preferences Θ_i , and we model $\gamma_{ic} = \Gamma_c \Theta_i$. Parameters Γ_c define purchase coincidence between product categories (Manchanda et al., 1999). Customers tend to purchase categories c and c' together if Γ_c and $\Gamma_{c'}$ are similar.

Products within the categories have different purchase frequencies (see Section 3.4.2). In some product categories a small number of products account for most of the sales. We thus weight the coupon discounts by the customer's purchase share of each product, that is

$$\bar{d}_{itc} = \frac{1}{|C|} \sum_{j \in c} p_{ijt}^{(2)} d_{itj}. \quad (3.18)$$

Inventory dynamics are determined by the customer-specific consumption rates, $Cons_{ic}$. The inventory is aggregated to the category level and consumption rates are different between the categories:

$$Inv_{ic}^t = Inv_{ic}^{t-1} + \sum_{j \in C} b_{itj} - Cons_{ic}. \quad (3.19)$$

3.4.2 Stage 2: Product Choice

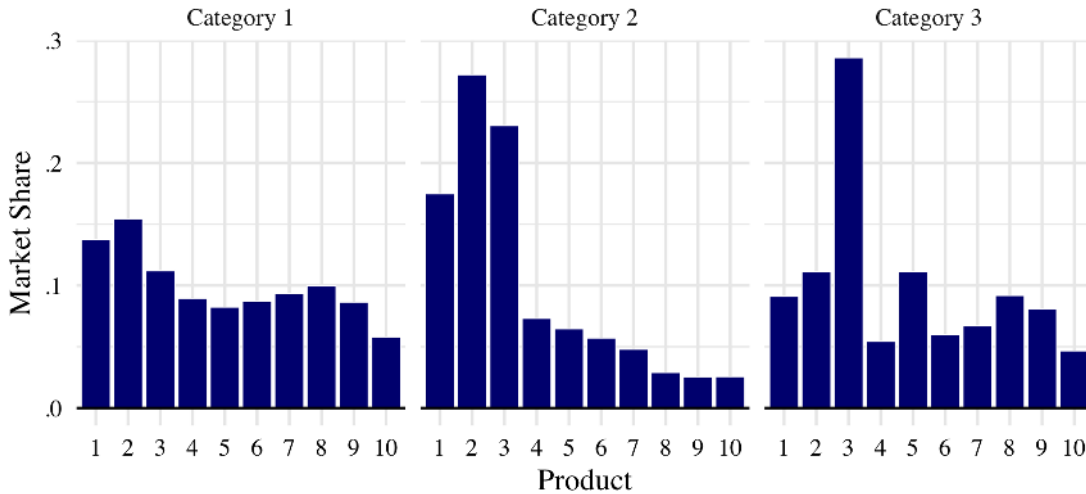
Product choice within a category follows a multinomial logit model (McFadden, 1974; Guadagni and Little, 1983). We assume the following form of customer i 's utility for product j at time t :

$$u_{itj} = \beta_{ij}^0 - \beta_i^p (1 - d_{itj}) price_j + \varepsilon_{itj}, \quad (3.20)$$

where β_{ij}^0 indicates customer i 's base utility for product j , β_i^p is a customer-specific price sensitivity, $price_j$ is a (regular) price of the product j , and d_{itj} is the size of the coupon provided to customer i for product j at time t . Assuming that the error term ε_{itj} follows a Gumbel extreme value distribution, the probability that the customer chooses product j in category c becomes

$$p_{itj}^{(2)} = \mathbb{P}(b_{itj} = 1 | y_{itc} = 1) = \frac{\exp\{u_{itj}\}}{\sum_{k \in C} \exp\{u_{itk}\}}. \quad (3.21)$$

The base utility, β_{ij}^0 , is customer- and product-specific. We define $\beta_{ij}^0 = B_j \Theta_i$, where Θ_i is the customer taste characteristic vector used in Stage 1. Customer i 's price sensitivity, β_i^p , is constant across categories. We assume that product prices, $price_j$, are constant over time, and coupons are the only source of price variation.

Figure 3.2. Simulated product market shares.

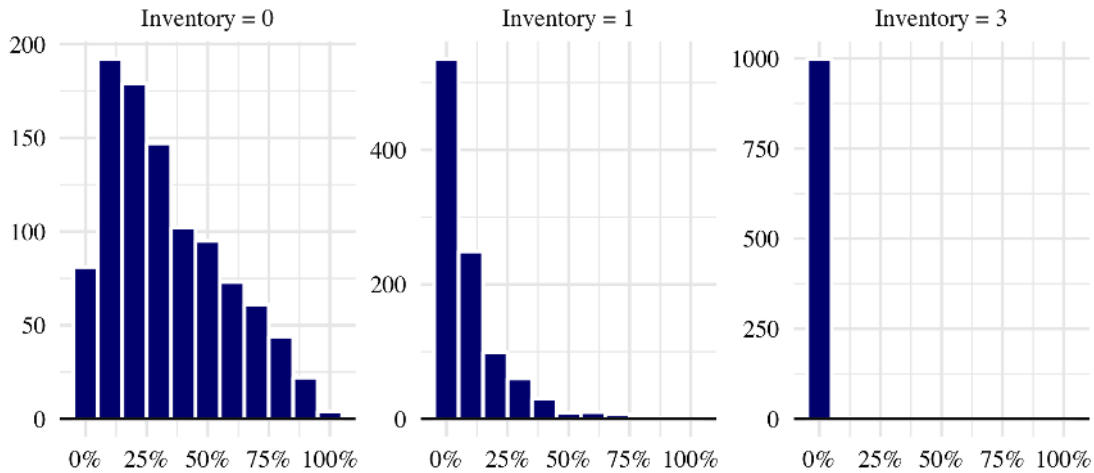
3.4.3 Simulation Calibration

We simulate a retailer with $J = 250$ products grouped into $C = 25$ categories and $I = 75,000$ customers. We also tested the proposed model with a larger number of products ($J > 1,000$) and categories ($C > 100$). The substantive findings reported in Sections 3.5 and 3.6 are robust, so for ease of exposition we opt for a smaller assortment size.

For every customer, we draw the taste characteristics from the multivariate normal distribution $\Theta_i \sim MVN(0^{h \times 1}, h^{-1}I^{h \times h})$, where h is the dimensionality of the latent tastes. We simulate 50 burn-in periods to allow the inventory to converge, and we simulate an additional 100 periods for model training and evaluation.

Customers receive coupons every time period. For model training and evaluation, we assume that the coupons are assigned randomly and that discounts range from 10% to 40%. We benchmark the predictive performance of the proposed product choice model in a simulation with five coupons per customer, and we evaluate coupon personalization based on product choice models in scenarios with one and five coupons per customer. Random coupon assignment in our simulation is consistent with the empirical application.

We define the parameters of the category purchase incidence model ($\gamma_c, \Gamma_c, \gamma_{ic}^p, \gamma_c^{Inv}, Cons_{ic}$) and the product choice model (B_j, β_i^p) to balance customer heterogeneity, inventory dynamics, and coupon and inventory effects on the product purchase rates. We calibrate the data generating process, such that the characteristics of the simulated data are similar to the empirical transaction in Section 3.6. For example, market concentration systematically varies between categories. Figure 3.2 shows product market shares in three different product categories. The products have similar market shares in Category 1, while a few products (or even one product) account for a large fraction of the revenue in Categories 2 and 3.

Figure 3.3. Category incidence probability histograms for different inventory levels.

We also demonstrate the sensitivity of the product purchase probabilities to the inventory in Figure 3.3. For a single category, we plot a histogram of the category incidence probabilities $p_{itc}^{(1)}$ across customers at three different levels of inventory, $Inv_{ic}^t \in \{0, 1, 3\}$. The distribution of the category incidence rates shrinks toward zero as we increase the inventory.

We provide all sampling distributions and parameter values in Appendix 3.8, and we include additional examples for the effect of the customers' inventory on category purchase incidence probabilities in Appendix 3.8.

3.5 Model Evaluation Based on Simulated Data

We compare the performance of the proposed product choice model to two baselines. The first baseline is a binary logit model (hereafter Binary Logit). We apply the Binary Logit model by-product. For each product, the independent variables are the customer-specific purchase frequency \bar{b}_{itj} , the purchase histories $[b_{itj}, \dots, b_{i,t-T+1,j}]$, and the current discount $d_{i,t+1,j}$. We use these independent variables to predict the purchase decision $b_{i,t+1,j}$.

The second baseline is a binary classifier based on LightGBM (Ke et al., 2017). LightGBM is an efficient implementation of the gradient boosting decision tree algorithm. We estimate LightGBM with an extended set of independent variables: the independent variables used in the Binary Logit model, the customer-product purchase histories, the current discounts of all J products, and customer embeddings based on the Product2Vec model (Gabel et al., 2019). We provide a complete description of the LightGBM independent variables in Appendix 3.8.

The proposed model comparison is nested in terms of the information used for prediction. For every customer, the Binary Logit model uses only the product-

Table 3.1. Aggregate prediction performance (simulation).

Model	Cross-Entropy Loss	Scaled Cross-Entropy Loss
True Probabilities	.0537	100.0%
Our Model	.0563	92.6%
LightGBM	.0589	85.2%
Binary Logit	.0662	64.5%

Note: All differences are significant at $p < .01$, based on standard errors (SE) computed using a nonparametric bootstrap with 100 replications.

specific information, that is the purchase history and the current discounts. The LightGBM model extends the Binary Logit model by incorporating cross-product effects (i.e., cross-product discounts and predefined summary statistics of the full customer purchase history across all products). Using the purchase histories for all products is not feasible in the LightGBM model due to high dimensionality and data sparseness. For completeness, we also evaluate the LightGBM model with the same independent variables as used by the Binary Logit. The performance of this model is similar to the Binary Logit across all comparisons.

Our proposed neural network model extends LightGBM by using all information about all products as an input to predict purchase incidence for a focal product. Leveraging rich high-dimensional information for all products is possible due to the proposed model architecture, including the bottleneck layers to encode cross-product relationships and the weight sharing to reduce the number of parameters and regularize the model.

3.5.1 Aggregate Prediction Performance

We evaluate the model’s prediction performance on holdout test data. We simulate 100 time periods. The first 90 time periods are the input for the model training. We use the trained models to make predictions for the last ten periods and compare the predicted purchase probabilities to the true simulated probabilities. The models never access the data from the last ten time periods during training and validation, so we can evaluate whether the models overfit the data. We provide details on the holdout test set construction in Appendix 3.8 and report the cross-entropy loss curves (as a function of training epochs) in Appendix 3.8.

Table 3.1 evaluates the prediction performance of the proposed neural network in a simulation with five random coupons per customer. We report the binary cross-entropy loss calculated using the holdout data. The binary cross-entropy measures how well the predicted probabilities approximate the binary purchase

decisions. We also present a scaled cross-entropy loss for interpretability. The scaled cross-entropy is based on a linear scale between the loss achieved by the true probabilities from the simulation and the loss achieved by the best uniform prediction.

Our model achieves significantly higher prediction performance than the reference models. The result is robust to the choice of the evaluation metric. In Appendix 3.8, we report the aggregate prediction performance of the models based on the area-under-curve metric (AUC) and Kullback-Leibler (KL) divergence. The cross-entropy loss and the AUC compare predicted and true purchase probabilities to the realized purchase decisions, while the KL divergence compares the predicted probabilities to the true simulated probabilities directly.

3.5.2 Prediction Performance Decomposition

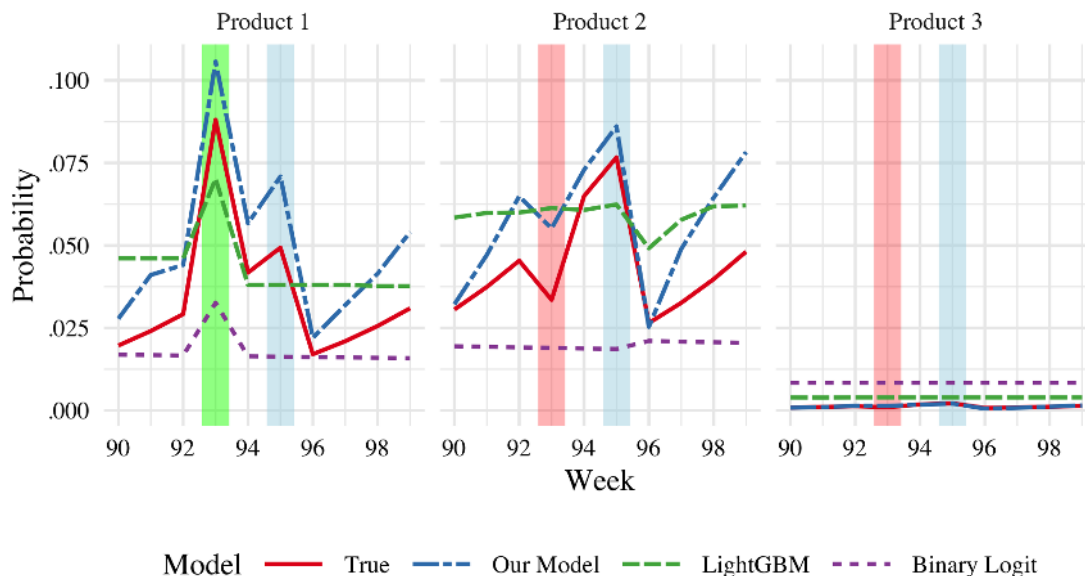
For effective coupon personalization two features are key: models need to (1) capture time dynamics in product choice (e.g., individual consumption patterns) and (2) predict the effects of coupons. We therefore provide a more detailed evaluation of the model’s predictive performance in the next subsections.

3.5.2.1 Product Choice Dynamics

Time dynamics of purchase probabilities in our simulation are determined by category inventory dynamics and coupon assignments. Figure 3.4 demonstrates the purchase probabilities of three products for one customer over ten holdout periods. The products belong to the same product category.

Note two interesting observations. First, the customer receives a coupon for Product 1 at time $t = 93$. The coupon affects purchase probabilities for all considered products. We observe a substantial positive effect on the purchase probability of Product 1, a negative effect on Product 2, and a small negative effect on Product 3. Our proposed model captures the first two changes, and underestimates the last effect. The Binary Logit model and LightGBM do not adjust the estimated probabilities for Products 2 and 3. The Binary Logit model only incorporates the coupon discount information of the focal product, so this behavior is expected. Although the LightGBM model has access to all product discounts, the model cannot capture the cross-product coupon effects either. High dimensionality and sparseness of $D_{i,t+1}$ are a reasonable explanation for this observation.

The second important observation in Figure 3.4 is that the customer purchases Product 2 at time $t = 95$. When the purchase happens, our simulation increases the category inventory for the customer and the increased inventory decreases purchase probabilities for all products in the category in the following time periods. We observe that the proposed neural network model captures this and adjusts

Figure 3.4. Time-series prediction (hold-out set).

Note: Best viewed in color.

the probabilities for all products accordingly. The LightGBM model adjusts probabilities only for Product 2. The Binary Logit model even increases the predicted purchase probability for Product 2 as a result of the increased purchase frequency for this product. The Binary Logit model fails to adjust estimated probabilities for Products 1 and 3.

The next two sections unfold this illustrative result by a deeper analysis of discount effects (Section 3.5.2.2) and time dynamics (Section 3.5.2.3).

3.5.2.2 Coupon Effects

The simulation setup implies that coupons affect the purchase probabilities of the promoted products and all other products in the products' categories. We can evaluate whether the model is able to recover coupon effects at the holdout data by comparing the true coupon discount elasticities in the simulation to the models' elasticity predictions.

To calculate the true discount elasticities, we save the simulation after period 90 (the last training period) and calculate purchase probabilities for each customer-product combination (i, j) in period 91 (the first test period) for two scenarios:

1. The retailer does not provide coupons to the customers.
2. All customers receive a 30% discount for product j_c .

We repeat this process for all products $j_c \in \{1, \dots, J\}$, average probabilities across the customers, and calculate product-specific discount elasticities

Table 3.2. True and estimated discount elasticities (average across products).


Model	Own-Elasticity	Cross-Elasticity Within Category	Cross-Elasticity Between Categories
True Probabilities	2.204	−.047	.000
Our Model	2.189	−.008	.001
LightGBM	1.472	.000	.000
Binary Logit	2.133	.000	.000

$$\varepsilon_{j,jc} = \frac{p_j^{30\%} - p_j^{0\%}}{.3 \cdot p_j^{0\%}}, \quad (3.22)$$

where $p_j^{30\%}$ is the average purchase probability for product j given a 30% discount on product j_c , and $p_j^{0\%}$ is the average purchase probability for product j assuming no coupons. This evaluation process yields a matrix of $J \times J$ own- and cross-product elasticities. We compute elasticities for the three models (our model, LightGBM, and Binary Logit) by replacing the true with the predicted purchase probabilities.

Table 3.2 reports the true and predicted discount elasticities. All models capture the positive impact of price discounts on purchase probabilities (“own-effect”). The average predicted elasticities for our model and the Binary Logit model are close to the true elasticity, whereas the LightGBM model underestimates the impact of discounts. Regarding cross effects, recall that our simulation assumes substitution and category incidence effects within categories, but no cross-price effects between categories (see the first row in Table 3.2). The Binary Logit model does not incorporate cross-price effects, so all cross-elasticities are zero. The LightGBM inputs the discounts of all products, but does not estimate significant cross effects either. The only model that finds significant cross-price effects is the deep neural network (although it underestimates the within-category cross-price elasticity). The estimated price effects across categories for our model are close to zero.

Figure 3.5 compares the (product-level) true own-price elasticities and the predicted elasticities for all three models. For our model, the mean error is .016, and 95% of the product-level errors fall into the interval $[-.60, .58]$. The Binary Logit achieves similar mean error with a much larger variance. In line with the results presented in Table 3.2, we observe that the LightGBM model systematically underestimates price elasticities. The root mean square error (RMSE) between the true and model-based elasticities is lowest for our model ($RMSE = .396$), followed by the Binary Logit model ($RMSE = .958$), and the LightGBM model ($RMSE = 1.164$).

Figure 3.5. Difference between true simulated elasticities and predicted elasticities.

Notes: The vertical lines on the ridges indicate the 95% CIs for each model. Best viewed in color.

We conclude that the proposed model predicts discount effects more accurately than the benchmark models. It is the only model that captures all types of discount effects (i.e., positive own-effects, negative cross effects within categories, no effects across categories). However, the neural network underestimates within-category cross effects of discounts.

3.5.2.3 Time Dynamics and Inventory Effects

We quantify how well the models capture the time dynamics of purchase probabilities by estimating the correlation of the predicted probabilities and the true probabilities over time (i.e., for the ten hold-out weeks) for every customer-product pair:

$$\rho^{time} = \frac{1}{IJ} \sum_{ij} corr_t^{ij} \quad (3.23)$$

with

$$corr_t^{ij} = \frac{cov_t(\hat{p}_{itj}, p_{itj}^{true})}{\sigma_{\hat{p}} \sigma_p}, \quad (3.24)$$

where $\sigma_{\hat{p}}$ and σ_p are the standard deviations of the predicted and true probabilities over time. To ensure numerical stability in the computation, we set $corr_t^{ij}$ to zero

Table 3.3. Time series correlation scores for model predictions.

	Data With Coupons		Data Without Coupons	
	Absolute	Scaled	Absolute	Scaled
True Probabilities	.8076	100.0%	.6455	100.0%
Our Model	.5791	71.7%	.5399	83.6%
LightGBM	.1033	12.8%	.0121	1.9%
Binary Logit	.1503	18.6%	.0572	8.9%

Note: All differences are significant at $p < .01$, based on SEs computed using a nonparametric bootstrap with 100 replications.


for time series with $\sigma_{\hat{p}} \cdot \sigma_p < 10^{-12}$. We compute the correlation metric for a simulated data set with coupons and a simulated data set without coupons. The first data set includes both sources of probability variation over time: the effect of the consumers' inventories (compare Figure 3.3) and the coupon effects. The second data set isolates the inventory effect.

Table 3.3 reports the average time correlation score ρ^{time} for our proposed model and the two baselines. The scaled correlation is based on a linear scale between zero and the time correlation achieved by the true simulation probabilities. The results confirm our analysis in Figure 3.4. The proposed neural network architecture achieves an average time correlation of $\rho^{time} = .58$. This value is considerably higher than the correlation scores for the baseline models and very close to the optimal score that we derive from the true simulation probabilities.

The correlation scores for the data without coupons are lower for all models. This is a result of the true probabilities exhibiting less variation over time (i.e., σ_p is smaller). The scaled correlation scores are lower for baseline models, but the scaled performance of our model is even higher for the data without coupons, suggesting that our model efficiently recovers the consumption patterns from the transaction data.

3.5.2.4 Identifying Product Category Structure

Our analysis of the cross-product coupon effects and the inventory time dynamics indicates that the proposed neural network model identifies cross-product relationships within categories. However, the model does not require specifying the product categories ex ante. The model learns cross-product relationships from the customer purchasing behavior at the training data.


Figure 3.6. Heat-map of the product embedding W_H .

Note: Best viewed in color.

The cross-product relationships are encoded in the parameters of the bottleneck layers. In Figure 3.6, we plot the heat-map of the bottleneck layer weight matrix W_H . The weight matrix W_H has 250 columns corresponding to $J = 250$ products in the simulated data. We order products by product categories, such that the first ten products correspond to the first product category, the next ten products correspond to the second category, etc. The heat-map reveals $C = 25$ groups of ten similar columns in the matrix W_H . The groups correspond to product categories. We refer to the columns of matrix W_H as *product embeddings*, as they incorporate information about product similarities. Products from the same categories have similar product embeddings.

Figure 3.7 depicts the two-dimensional t-SNE projections (Maaten and Hinton, 2008) of the product embeddings. Each dot represents one product, and we identify the true (simulated) categories by different colors. We observe that the products form clusters corresponding to different categories, and the clusters are perfectly separated, which confirms that the trained product embeddings encode information about products and product category structure.

Appendix 3.8 contains a deeper analysis of how the different components of the neural network architecture impact the predictive performance of our proposed model. We sequentially remove components of the full architecture and demonstrate that both the time filter and the bottleneck layers are critical to the model's predictive performance.

Figure 3.7. t-SNE projection of the product embedding W_H .

Notes: Colors indicate true product categories. Best viewed in color.

3.5.3 Performance Gains for Coupon Optimization

We conclude the evaluation of the proposed product choice model in the context of the simulation by evaluating how the improved prediction performance translates into the efficiency gains for the coupon personalization problem.

The performance of coupon personalization depends not only on the product choice model, but also on the coupon optimization algorithm. The coupon optimization algorithm allocates coupons to customers based on the estimated effects of the coupons on purchase probabilities. We evaluate the overall revenue gains with one coupon per customer or five coupons per customer. In both cases, we focus our analysis on the product choice model by keeping the optimization algorithm constant and changing the underlying choice models.

We first evaluate the performance of the coupon assignment for the case that every customer receives one coupon. We assume that customers act independently, so with a single coupon per customer we can enumerate and evaluate all possible

coupon allocations. For a customer with a purchase history B_{it}^T and purchase frequencies B_{it}^∞ , we select the coupon that maximizes the expected revenue

$$D_{it}^* = \underset{D=[d_1, \dots, d_J]}{\operatorname{argmax}} \sum_j \hat{p}_{itj} \left(D, B_{it}^T, B_{it}^\infty \right) (1 - d_j) \operatorname{price}_j \quad (3.25)$$

s.t. $D \in \{.1, .2, .3, .4\}^{J \times 1}$
and $\sum_j \mathbb{1}(d_j > 0) = 1.$

We also evaluate coupon policies that generate five coupons per customer. Allocation by a complete enumeration is no longer feasible in this case. The evaluation time for one combination of five coupons for all customers takes approximately .5s. There are over 8×10^{12} possible coupon combinations which results in over 100,000 years of computing time to solve the problem through a complete enumeration. Instead, we consider a greedy heuristic for coupon allocation. The greedy heuristic begins by selecting a single coupon that maximizes revenues. It then sequentially adds coupons, one coupon at a time, to maximize revenues given the previously chosen coupons. The method stops when the five coupons are selected. The greedy heuristic has previously been successfully applied in product line optimization (Green and Krieger, 1985; Belloni et al., 2008).

We demonstrate the coupon optimization results for different product choice models in Table 3.4. We report the expected revenue lift per customer and the percent improvement of revenue over the no-coupon baseline. The expected revenue lift measures the difference between the revenue with coupon (based on the respective coupon policies) and the revenue without coupons. We integrate over the error terms of the product and category choice models in the simulation by evaluating the responses to coupons 100 times with different seeds for the random number generator. The results in Table 3.4 are the average uplift over the 100 replications and we report the SEs of the sample means over the replications in parentheses.

A random coupon allocation defines the lower bound for coupon performance. If the products are too expensive, providing random coupons can improve the revenue without optimization. However, revenue uplifts are very small. Coupon policies that optimize revenues should outperform this lower bound. We thus compare the coupon optimization methods based on the product choice models with the random coupon assignment and find that all optimized methods outperform the lower bound for both one and five coupons per customer.

A second reference point is a mass marketing coupon policy that provides the same revenue-maximizing price promotion to all customers (Best Uniform). As

Table 3.4. Coupon optimization results.

	One Coupon per Customer		Five Coupons per Customer	
	\$ Revenue Lift	% Revenue Lift	\$ Revenue Lift	% Revenue Lift
Our Model	\$.72 (\$.01)	2.26%	\$2.61 (\$.02)	7.68%
LightGBM	\$.45 (\$.01)	1.41%	\$1.50 (\$.02)	4.55%
Binary Logit	\$.50 (\$.01)	1.57%	\$1.40 (\$.02)	4.26%
Best Uniform	\$.21 (\$.01)	.65%	\$.91 (\$.01)	2.83%
Random	\$.01 (\$.00)	.03%	\$.05 (\$.00)	.17%

Note: All differences are significant at $p < .01$, based on SEs computed using a nonparametric bootstrap with 100 replications.

expected, this policy leads to a larger revenue increase than random coupons, but is outperformed by all three models that personalize coupons.

The proposed neural network model significantly improves coupon optimization over the LightGBM and Binary Logit baselines. For a single coupon, our policy yields a $3.4\times$ higher revenue lift than the Best Uniform policy and a $1.6\times$ higher revenue increase than the LightGBM model. The advantage over the LightGBM model increases even further in the case of five coupons. Note that the revenue lift per coupon decreases with a larger number of coupons for all policies. We expect diminishing returns as the policies by construction select the most profitable coupons first, so this result offers face validity.

A deeper analysis of the simulated purchase probabilities with and without coupons allows us to evaluate why our model outperforms the baselines in the coupon optimization problem. Recall that the neural network is more successful in identifying purchase timing. We observe two patterns in the coupon policies that can explain the increased revenue lift. First, the policy based on our proposed deep neural network creates additional category purchases by providing coupons to the customer-category combinations with lower inventory at the time of the coupon distribution. This leads to higher coupon redemption rates. At the same time, we find that the neural network policy provides coupons to the customer-category combinations with a large discount sensitivity and a low base probability for a category purchase more often. A coupon yields the highest own-product revenue lift in these cases and our model finds such opportunities more reliably than the other baselines.

Second, the neural network policy provides coupons for more expensive products. In line with the empirical data, more expensive products in our simulation have lower interpurchase times. The baseline policies do not model purchase timing

as well as the neural network. They consequently focus on cheaper products that have higher purchase frequencies and fail to capitalize on the revenue potential that more expensive products offer (if coupons are timed well). We provide more details on the comparison of targeting policies in Appendix 3.8.

We conclude that the coupon optimization generates a significant revenue increase in our simulation. Our proposed product choice model improves prediction accuracy of customers' purchase rates and this translates into larger revenue gains at the coupon optimization.

3.6 Model Evaluation Based on Empirical Data

3.6.1 Data

We validate the prediction performance of our model using transaction data provided by a leading German grocery retailer. The data set comprises three data sources: loyalty card data, market basket data, and coupon data. Loyalty card data follows a panel structure and contains transactions by the loyalty card holders. Market basket data includes information about purchases without customer identifiers. Coupon data contains information about the coupons provided to the customers (including unredeemed coupons). Overall, the data set spans over 83 weeks (2015-2016) and includes 22,740,377 purchases by 489,438 loyalty card customers and 73,048,605 shopping baskets by customers without a loyalty card. We provide summary statistics in Table 3.5.

We focus our analysis on 50 product categories for which the retailer distributed coupons during the period of the analysis. The categories include food products such as milk, bread, chocolate bars, and coffee, and nonfood products such as shampoo, fabric softener, and toothbrushes. The categories vary in the interpurchase times and a competitive intensity. We provide a complete list of the product categories in Appendix 3.8.

Most customers visit the store no more than once a week, so we aggregate data to a weekly level. The median time between two shopping trips across all customers is two weeks ($SD = 4.02$). This value is typical for a supermarket in a German metropolitan area. The retailer used coupons to promote category-brand combinations that group stock keeping units of the same package size and price range. We follow the retailer's product grouping and use this level of aggregation for our analysis. For a small subset of customers, the retailer provided coupons randomly. Our empirical analysis only considers customers with randomly assigned coupons, which allows us to avoid endogeneity concerns in model training and validations.

Table 3.5. Summary statistics: data sets for empirical application.

Data Set	Variable	Value
Loyalty Card Data	# of users	489,438
	# of weeks (date range)	83 (2015/05 - 2016/12)
	# of brands (# of retailer categories)	758 (50)
	# of stores	155
Coupon Data	# of coupons	650,973
	Avg. # of coupons per customer (SD)	4.76 (6.33)
	Discount range	[5%, 50%]
	Avg. redemption rate (SE)	1.529% (.014%)
Market Basket Data	Avg. discount (SD)	23.7% (10.0%)
	# of baskets	73,048,605
	# of months	12 ^{a)}
	Avg. # of products per basket	4.91

Note: a) First year of loyalty card data.

3.6.2 Evaluation Results

For the model evaluation, we follow the approach used in the simulation study and create a hold-out test set by splitting the data in the time dimension. The first 73 weeks are used for model training, whereas the last ten weeks comprise the test data. We predict the purchase probabilities for all products and 1,000 customers.

We train the deep neural network in two stages. In the first stage we apply P2V-MAP (Gabel et al., 2019) to the market basket data to derive product embeddings. We then use the pretrained embeddings to initialize W_H , W_d , and W_∞ . Pretraining embeddings is a common approach to training deep neural network architectures in computer vision and natural language processing. Pretraining helps to avoid local minima in supervised learning and facilitates a better generalizability of results by having a regularizing effect on the neural network (Erhan et al., 2010). Initializing the product embeddings in our neural network with the output of P2V-MAP reduces the number of training iterations that is required to achieve model convergence by approximately 25%. The training time for mini-batch gradient descent scales linearly with the number of training iterations, so pretraining the product embedding reduces training time significantly.

In the second step of the model training, we initialize the parameters of the bottleneck layers with the product embedding and train the full neural network by minimizing the binary cross-entropy loss (see Section 3.3.3). This step fine-tunes the pretrained bottleneck layer weights.

We calibrate the hyperparameters of the model on a validation set by comparing the binary cross-entropy loss over a small number of randomly sampled hyperparameter sets. Random search is a common approach to configuring neural networks and typically finds solutions that are as good as grid searched results within a small fraction of the computation time of grid search (Bergstra and Bengio, 2012). We initialize the hyperparameter search with the values used in the simulation. We find that a larger embedding size ($L = 50$) improves the test loss and that the model converges in less epochs ($n_{epoch} = 10$). For the other hyperparameters, the random search did not yield a significant loss improvement, so we use the same values as in the simulation. We calibrate the simulation to mimic the behavior of the empirical data, so the similarities between the simulation study and the empirical application are not surprising. In line with the results of the simulation study we find that the parsimonious architecture makes our neural network robust to overfitting (see Appendix 3.8). We present a two-dimensional t-SNE projection of the product embedding W_H trained on empirical data in Appendix 3.8.

An important difference to the simulation study is that true purchase probabilities are unknown in the context of the empirical application. We therefore focus on the comparison of binary cross-entropy loss that evaluates the predictions based on the observed (binary) purchase indicator and the predicted purchase probabilities. We evaluate the prediction performance of our proposed model and compare it to the baselines used in Section 3.5.

Table 3.6 reports the evaluation results. We find that the ranking of the models based on the predictive performance is in line with the results obtained from the simulated data, and our proposed model achieves a lower cross-entropy loss than the reference methods.

We conduct an additional regression analysis to understand the performance differences between our deep neural network (DNN) and each of the two baseline models. First, we compute the binary cross-entropy loss for each observation (customer, product, time) in the test set for the DNN and the reference model. We then compare the loss values using a linear regression:

$$\begin{aligned}
 M1: \quad L_{ijtm} &= \alpha_0 + \alpha_c + \delta_{DNN}, \\
 M2: \quad L_{ijtm} &= \alpha_0 + \alpha_c + \delta_{DNN} + \delta_C + \delta_P + IPT_{ic} \\
 &\quad + \delta_{DNN} \times \delta_C + \delta_{DNN} \times \delta_P + \delta_{DNN} \times IPT_{ic},
 \end{aligned}$$

where m indexes the model (either DNN or the reference model), α_0 is the regression intercept and α_c are category-level fixed effects, and IPT_{ic} is the average customer-level category interpurchase time computed on the training data. The regression includes three indicator variables:

Table 3.6. Aggregate prediction performance (empirical application).

Model	Cross-Entropy Loss	Improvement vs. Best Uniform
Our Model	.0095	48.9%
LightGBM	.0116	37.6%
Binary Logit	.0126	32.3%
Best Uniform	.0186	-

Note: All differences are significant at $p < .01$, based on SEs computed using a nonparametric bootstrap with 100 replications.

$$\begin{aligned}\delta_{DNN} &= \mathbb{1}_{ijtm} (m = DNN), \\ \delta_C &= \mathbb{1}_{ijtm} (d_{itj} > 0), \\ \delta_P &= \mathbb{1}_{ijtm} \left(\left[\sum_{k \in C} b_{i,t-1,k} \right] > 0 \right),\end{aligned}$$

i.e., δ_{DNN} identifies loss values corresponding to our model, δ_C marks observations (i, t, j) with a coupon, and δ_P is an indicator for observations with a category purchase in a previous period. We use the retailer's category definition to compute IPT_{ic} and δ_P .

The three interaction terms with the indicator variable for the neural network observations, δ_{DNN} , allow us to evaluate whether our model loss is particularly low for the given data partitions (low is good). For readability, we multiply all loss values by a factor of 100. We repeat the analysis for the LightGBM model and the Binary Logit. In total, this analysis produces four sets of regression coefficients (two nested model specifications: M1 and M2; two model comparisons: DNN vs. LightGBM and DNN vs. Binary Logit).

Table 3.7 depicts the regression results. The results for the four models are very similar. The neural network achieves a significantly lower binary cross-entropy loss than the LightGBM model and the Binary Logit model (δ_{DNN}). On average, predicting probabilities is more challenging for observations with coupons (δ_C), observations with a category purchase in the last week (δ_P), and smaller interpurchase times (IPT_{ic}).

The interaction terms reveal that the DNN model particularly improves the predictions for observations with a recent category purchase and observations characterized by smaller interpurchase times, beyond the average improvement

Table 3.7. Binary cross-entropy loss regression analysis.

	DNN vs. LightGBM		DNN vs. Binary Logit	
	(M1)	(M2)	(M1)	(M2)
Intercept α_0	2.121 *** (.019)	2.039 *** (.170)	2.279 *** (.021)	2.259 *** (.186)
DNN δ_{DNN}	-.376 *** (.027)	-.465 *** (.041)	-.534 *** (.030)	-.752 *** (.045)
Coupon δ_C		.626 *** (.087)		.682 *** (.095)
Category Purchase δ_P		2.269 *** (.053)		1.986 *** (.058)
Interpurchase Time IPT_{ic}		-.027 *** (.001)		-.032 *** (.001)
DNN \times Coupon		-.135 (.123)		-.193 (.134)
DNN \times Category Purchase		-.899 *** (.074)		-.612 *** (.081)
DNN \times Interpurchase Time		.009 *** (.001)		.013 *** (.001)
Category Fixed Effects	Yes		Yes	

Notes: To simplify exposition, we scaled the binary cross-entropy loss values by a factor of 100. Sig. label: *** $p < .01$.

measured by δ_{DNN} . The DNN binary cross-entropy loss also tends to be smaller than the loss in the reference models for coupon observations, but the effect is not statistically significant. These findings are based on the empirical data and confirm the simulation-based analysis presented in Sections 3.5 and 3.6.

3.7 Conclusion

Retailers collect high-quality data about the customer choice and the effectiveness of marketing channels. However, leveraging these data for target marketing is challenging. Large assortments and customer bases require prediction and optimization methods to scale to high-dimensional inputs and large data sets.

In this paper, we have developed a nonparametric model to predict product choice for the entire assortment of a large retailer. The model is motivated by the coupon optimization problem. Given coupon assignments and customer purchase histories, our model predicts individual product choice probabilities for every product in the assortment. It is based on a custom deep neural network architecture and can be directly applied to transaction data from their loyalty card program. Our

model eliminates the need for data preparation and assumptions about category delineation and cross-product effects.

We have evaluated the prediction performance of the proposed model in simulations. The model significantly outperforms the machine learning benchmarks. We demonstrate that the model is able to approximate own- and cross-product coupon and inventory effects out-of-sample. The model recovers cross-product effects by identifying product similarities from the training data.

The model's higher prediction accuracy leads to a better performance for coupon personalization. We have verified this for the cases of one coupon and five coupons per customer. In the simulation, coupon optimization methods achieve substantially higher revenue gains when using purchase probabilities predicted by our model compared to the baseline prediction methods (e.g., up to 74% revenue increase vs. LightGBM).

The empirical application based on data from a leading German grocery retailer verifies the prediction accuracy results from the simulation study. The prediction performance improvements are particularly large for three types of observations (i.e., customer, time, product combinations): observations with a recent category purchase, observations characterized by smaller interpurchase times, and observations with coupon assignments (although the latter is not statistically significant). We conclude that the proposed model is a suitable solution to large-scale coupon optimization.

Target marketing and product choice modeling in retail provide a rich context for future research. Machine learning and deep learning are active areas in computer science and marketing literature. Our architecture applies convolutional filters and bottleneck layers to model choice based on loyalty card data. Promising alternative architectures include the WaveNet and recurrent neural networks. Moreover, deep neural networks are capable of processing unstructured data, so image data and customer reviews are promising sources for extending the model input. Additional, qualitatively different inputs can further improve the model's prediction accuracy. Finally, field experiments can provide additional empirical validation and new insights regarding the performance of the proposed product choice model in target marketing, including pricing, product recommendations, and promotion personalization.

3.8 Appendix

3.8.1 Parameters for Adam Optimizer

Table 3.8. Adam optimizer parameters (PyTorch).

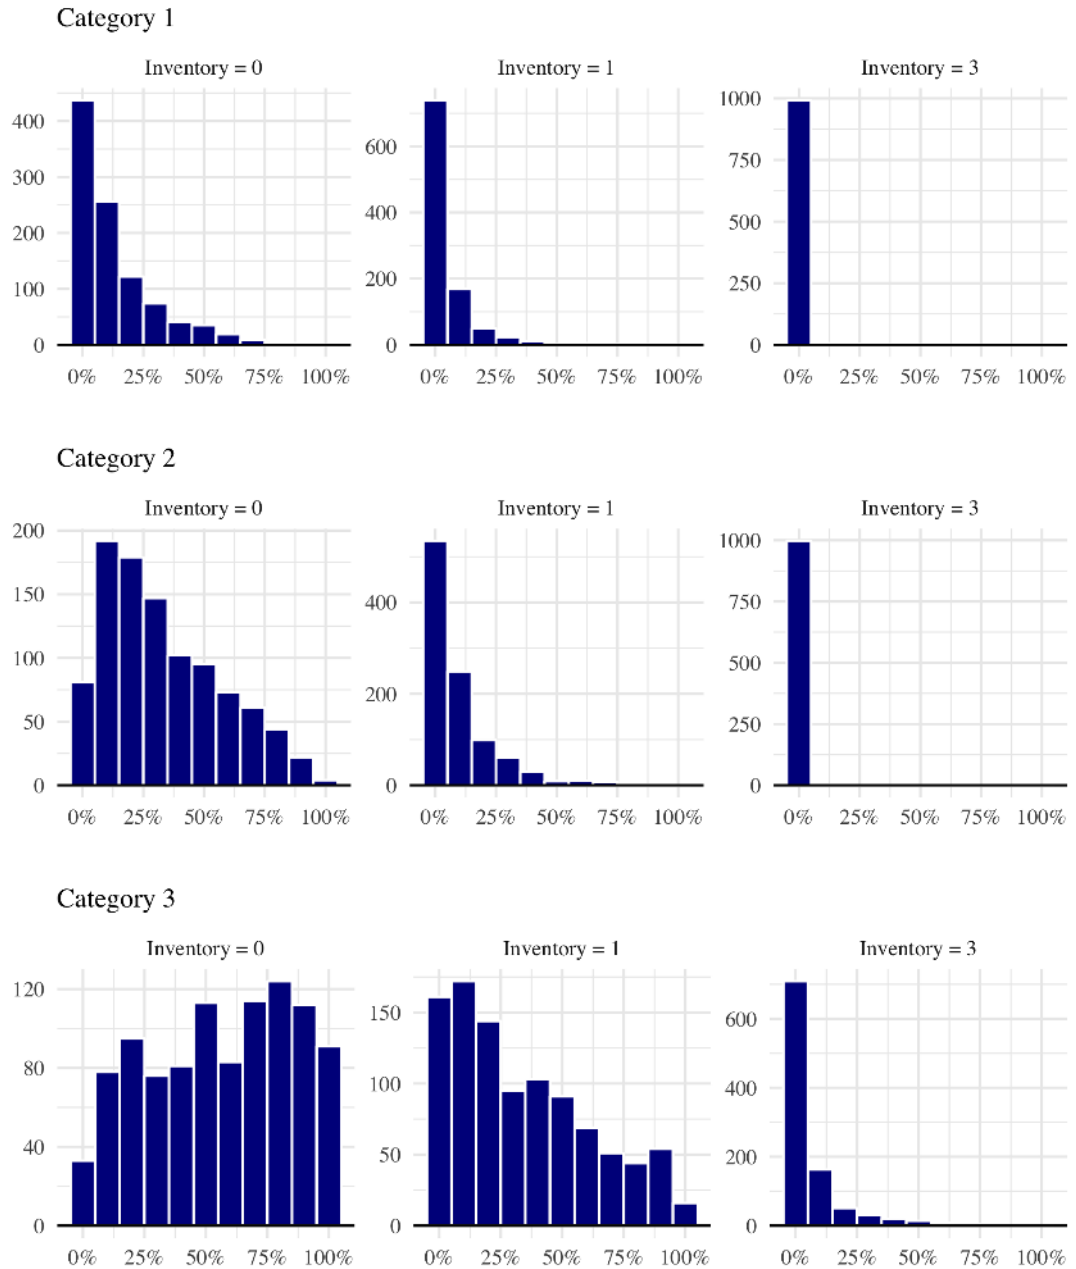
Parameter	Description	Value
lr	Learning rate	.001
betas	Coefficients used in the computations of the running averages and squared average of the gradient	[.9, .999]
eps	Constant added to the denominator to improve numerical stability	1e-8
weight_decay	Weight decay	0

3.8.2 Parameter Sampling for Simulation

Table 3.9. Simulation parameters.

Parameter	Description	Value
Category	γ_c	Category base utility $\gamma_c \sim U(-1.6, .2)$
	γ_c^{Inv}	Inventory sensitivity $\gamma_c^{Inv} \sim U(-1.2, -.6)$
	Inv_{ic}^0	Inventory initialization $Inv_{ic}^0 \sim Exp(.4)$
	$Cons_{ic}$	Consumption rate $Cons_{ic} = Cons_c(1 + W_i^{Cons})$ $W_i^{Cons} \sim U(-.2, .2)$ $Cons_c \sim U(.1, 1.4)$
	γ_{ic}^p	Category discount sensitivity $\gamma_{ic}^p \sim \gamma_c^p LN(0, .05)$ $\gamma_c^p \sim U(.5, 2.5)$
	Γ_c	Category similarity $\Gamma_c \sim MVN(0^{h \times 1}, I^{h \times h}), h = 20$
Product	B_{jc}	Product similarity $B_{jc} \sim MVN(0^{h \times 1}, I^{h \times h})$
	β_i^p	Discount sensitivity $\beta_i^p \sim LN(.6, .4)$
	$price_j$	Product price $price_j \sim (.5 + U(0, 1))price_c$ $price_c \sim LN(.7, .3)$

3.8.3 Effect of Inventory on Category Purchase Incidence

Figure 3.8. Category incidence probability histograms for different inventory levels for three product categories.

3.8.4 Independent Variables in the LightGBM Model

We provide a complete list of LightGBM features below. The features are unique for customer i , product j , and time t . We include two types of features: own-product and cross-product features.

1. Own-product discounts d_{itj} .
2. Own-product purchase frequencies \bar{b}_{itj} .
3. Own-product purchase histories B_{itj}^T .
4. Own-product moving window purchase frequencies

$$B_{itj}^h = \frac{1}{h} \sum_{k=1}^h b_{i,t-k+1,j} \quad (3.26)$$

with various window sizes

$$h \in \{2, 4, 8, 10, 15, 20, 25, 30\}. \quad (3.27)$$

This feature is motivated by the time filters in our deep neural network. It is designed to allow the LightGBM model to identify purchasing patterns along the time dimension of our panel data.

5. Cross-product discounts $d_{itk} \forall k \neq j$.
6. Cross-product purchase histories: Including the full purchase histories of all products would result in $(J \times T)$ -dimensional input. This is not feasible due to high dimensionality and data sparseness. We instead propose to use the cosine similarity between a customer's embedding u_i and a product j embeddings v_j to measure a product j 's attractiveness for customer i . This feature allows LightGBM to model preference correlations between products across the full assortment. We use the Product2Vec model to compute product embeddings, v_j , using market basket data (Gabel et al., 2019). We obtain customer embeddings, u_i , as an average of product embeddings for all products purchased by the customer in the past

$$u_i = \frac{\sum_{t=1}^{90} b_{itj} v_j}{\sum_{t=1}^{90} b_{itj}}. \quad (3.28)$$

7. Cross-product customer embeddings u_i .

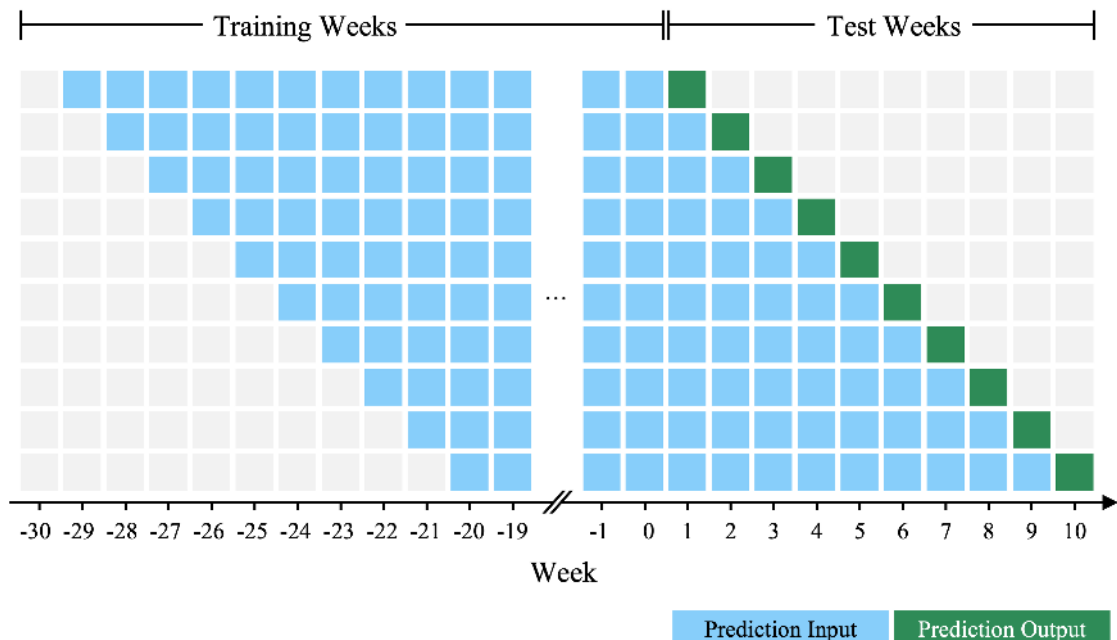
3.8.5 Test Set Definition

We use the models to predict future purchases (e.g., purchases at the next shopping trip), so we create the test data set by splitting the data along the time dimension. For a data set with 100 time periods we use the last ten periods as a test set and use the first 90 for model training and validation. For each customer, this approach yields $J \times 10$ predicted purchase probabilities.

Our data splitting approach avoids information transmission between the training data and the test data. Using more than one week in the test set increases the validity of the model evaluation and allows us to evaluate how well our model captures changes in purchase probabilities for a given customer and product over time.

Figure 3.9 illustrates how we predict purchase probabilities for a customer-product pair (i, j) . In predicting the purchase probability for test period $t = 1$ (green cell in row 1), we use the $T = 30$ time periods before $t = 1$ as model input (blue cells in row 1). The predicted purchase probabilities in week 2 are based on T time periods before $t = 2$ (including time period 1), etc. The cascading data structure ensures that the model always uses up-to-date information for prediction.

Figure 3.9. Data split for hold-out evaluation.



Note: Best viewed in color.


3.8.6 Loss Curves

We present the loss curves on the training, validation, and test data in Figure 3.10. The construction of the test data follows the description in Appendix 3.8. The validation data includes a fraction of the observations from time period $t = 89$.

We find that the test and validation losses converge after approximately 25 epochs. We compute 95% confidence intervals through nonparametric bootstrapping and do not observe significant differences in the loss between the three data sets. The same is true in the empirical application, so we conclude that the deep neural network is not overfitting.

We observe a large decrease in losses between epochs 10 and 15. This loss decrease occurs when the neural network learns the product embeddings W_d , W_∞ , and W_H . We elaborate more on this point in Appendix 3.8.

Figure 3.10. Loss curves for training, validation, and test data.



Note: Best viewed in color.

3.8.7 Comparison of True and Predicted Probabilities

For the coupon optimization problem, it is important that probabilities are scaled correctly. We visually verify this by plotting the predicted probabilities \hat{p}_{itj} against the true simulation probabilities p_{itj} for a subset of all categories (Figure 3.11) and products (Figure 3.12). Each point in the scatter plots is the probability for a single customer, week, and product. We do not find any systematic prediction errors. This further validates our model's predictions.

Figure 3.11. Probability scatter plots for six product categories.





Figure 3.12. Probability scatter plots for six products in category 1.

3.8.8 Additional Benchmarking Metrics for Simulated Data

We provide additional benchmarking metrics in Table 3.10

Area under receive operator characteristic curve (AUC):

$$AUC(b, \hat{p}) = \int_{-\infty}^{\infty} TPR(t) FPR(t) dt, \quad (3.29)$$

based on the predicted probability \hat{p} and the true purchase indicator b . TPR is the true positive rate and FPR denotes the false positive rate.

The KL divergence compares predicted probabilities \hat{p} and true probabilities p :

$$KL(p, \hat{p}) = \sum_{i,j,t} \left[p \log \left(\frac{\hat{p}}{p} \right) + (1 - p) \log \left(\frac{1 - \hat{p}}{1 - p} \right) \right]. \quad (3.30)$$

Table 3.10. Additional metrics for aggregate prediction performance (simulation).

	AUC	KL Divergence	Cross-Entropy Loss
True Probabilities	.9490	.0000	.0537
Our Model	.9393	.0032	.0563
LightGBM	.9242	.0128	.0589
Binary Logit	.9170	.0281	.0662

Note: All differences are significant at $p < .01$, based on SEs computed using a nonparametric bootstrap with 100 replications.

3.8.9 Nested DNN Model Specification

To better understand how the different components of the neural network architecture impact the model’s predictive performance, we compare four nested model specifications:

Full DNN Full model described in Section 3.3 and evaluated in Sections 3.5 and 3.6.

DNN without time filter To remove the time filters we set H to 1 and fix the time filter weights $\mathbf{w}_{h=1} = 1/T$. The time filter therefore simply averages the purchase histories for each product. We freeze the time filter weights and train the remaining weights of the network as usual.

DNN without bottleneck layers To remove the bottleneck layers, we replace the products of the weight matrices in the bottleneck layer (e.g., $W_d^\top W_d$) with $(J \times J)$ -dimensional identity matrices. We freeze the bottleneck layers and train the remaining weights of the network as usual.

Minimal DNN We remove both the time filter and the bottleneck layers.

Table 3.11 depicts the benchmarking scores (on the test set) for the four model specifications described above. The differences in cross-entropy loss are small but managerially relevant. The bottleneck layer improves the cross-entropy loss more than the time filter. We observe a significantly lower time correlation for the model without time filters. Only adding the time filter (but not using the bottleneck layers) produces correlation scores similar to the LightGBM baseline. In this specification, the DNN disregards that products from the same category are exchangeable (recall the Coke/Pepsi example). The predictions therefore fail to model consumption patterns adequately. The combination of the time filter and the product embedding increases the correlation by more than 3 times. Learning category structure is necessary to approximate purchase incidence. The results indicate that both components, the time filter and the bottleneck layers, significantly improve the model’s predictive performance and that the largest increase can be accomplished by using both components simultaneously.

To illustrate how the (hold-out) cross-entropy loss, the time correlation metric, and the product embeddings are related, we show the three outputs as a function of the training epochs in Figure 3.13. We observe that the neural network learns product embeddings between epochs 10 and 15. This coincides with an increase in the correlation scores and a decrease in the cross-entropy loss.


Table 3.11. Benchmarking results (test set) for four nested DNN models.

Model	Cross-Entropy Loss	Time Correlation
Full DNN	.0563	.5791
DNN w/o Time Filter	.0573	.0535
DNN w/o Bottleneck Layers	.0576	.1341
Minimal DNN	.0579	.0482
LightGBM	.0589	.1033


Note: Best scores per column in bold.

Figure 3.13. Test loss, correlation metric and product embedding W_H (products 1 to 40).


Product Embedding (Bottleneck Layer)



Time Correlation Between True and Predicted Probabilities



Test Loss



Note: Best viewed in color.

3.8.10 Analysis of Coupon Policies

We compute additional descriptive statistics to better understand (1) which customer-category-product combinations the policies target and (2) why the policy based on our proposed neural network leads to the highest revenue lifts. The starting point of the analysis are the coupons that the DNN, LightGBM, and Binary Logit coupon policies select for the 2,000 test customers (we focus on the case of one single coupon per customer). Each coupon is a unique combination of a customer and a product. We report the descriptive statistics about the coupons in Table 3.12. The statistics have different levels of variation. For example, the time since the last category purchase is category-specific, whereas the average inventory is category- and customer-specific. We report the average values for each coupon policy and the relative values for the LightGBM and Binary Logit policies relative to the DNN policy to simplify the comparison.

We find that the DNN policy targets more expensive products, categories that are purchased less frequently (lower γ_{ic}) and customer-category combinations with a lower inventory. The DNN policy therefore achieves a good balance in targeting categories. The policy identifies customer-category combinations for which base probabilities are small enough to increase revenue through coupons (i.e., potential for category purchase rate lift) but sufficient to maintain redemption rates.

This is supported by the observed category purchase incidence probability ($p_{itj}^{(2)}$), that is smaller than the average probability for the Binary Logit policy (lack of incrementality) but larger than the average probability for the LightGBM policy (lack of effectiveness). The higher price and the higher rate of incremental category purchases are two of the potential drivers for a higher revenue lift.

Table 3.12. Coupon policy analysis.

	DNN	LightGBM		Binary Logit	
	(abs.)	(abs.)	(rel.)	(abs.)	(rel.)
Average price of targeted product	\$14.5	\$13.6	94.0%	\$13.5	93.0%
Time since last category purchase	5.0	5.0	98.9%	3.9	77.8%
Category base utility (γ_c)	-.644	-0.658	102.2%	-.642	99.8%
Category base utility (γ_{ic})	.518	.524	101.3%	.703	135.7%
Category discount sensitivity (γ_{ic}^p)	2.105	1.737	82.5%	1.815	86.2%
Inventory	.784	1.014	129.3%	.877	111.9%
Category purchase incidence probability ($p_{itj}^{(2)}$)	23.7%	19.3%	81.5%	26.6%	112.1%
Product base utility (β_{ij}^0)	3.637	3.627	99.7%	4.037	111.0%
Incremental category purchase rate ^{a)}	19.5%	10.1%	51.9%	16.3%	83.3%
Revenue Uplift	2.26%	1.41%	62.4%	1.57%	69.5%

Note: a) The incremental category purchase rate is defined as the fraction of coupon observations with a category purchase in the simulation with a coupon treatment minus the fraction of coupon observations with a category purchase in the simulation without a coupon treatment (using the same simulation “state”).

3.8.11 Empirical Data–Category Statistics

Table 3.13. Category characteristics for loyalty card data set.

Category	Concentration	IPT (SD)	Category	Concentration	IPT (SD)
Beer	.17	2 (9.78)	Juices	.27	3 (11.86)
Butter	.22	2 (7.77)	Ketchup	.35	8 (13.84)
Cereal bars	.41	4 (12.50)	Milk	.17	1 (5.96)
Chewing gum	.24	4 (11.74)	Muesli / corn flakes	.13	3 (10.17)
Chips	.14	3 (9.83)	Paper towels	.56	4 (10.83)
Chocolate bars	.14	3 (10.52)	Pasta	.21	5 (11.47)
Chocolate spread	.56	6 (13.12)	Pasta (fresh)	.53	6 (13.42)
Coffee	.16	5 (10.47)	Pizza	.39	3 (9.86)
Coffee beans	.30	5 (10.82)	French fries	.36	6 (13.04)
Coffee capsules	.25	3 (9.06)	Salt sticks	.27	4 (12.19)
Coffee pads	.15	3 (9.03)	Shampoo	.12	9 (14.49)
Condensed milk	.31	2 (7.62)	Sliced cheese	.08	2 (7.99)
Cough drops	.48	3 (12.23)	Soft drinks	.16	1 (7.05)
Crisp bread	.21	3 (11.01)	Tea	.32	5 (11.64)
Detergent	.20	8 (13.04)	Toast	.32	2 (8.38)
Dishwashing liquid	.21	8 (13.59)	Toilet paper (wet)	.45	4 (10.02)
Dishwashing tabs	.33	12 (15.03)	Toilet paper	.33	5 (11.42)
Energy drinks	.16	2 (8.26)	Toothbrushes	.20	10 (15.94)
Fabric softener	.23	7 (12.94)	Toothpastes	.09	8 (13.16)
Milk	.34	2 (7.84)	Tuna	.35	4 (12.46)
Ice cream	.15	3 (12.19)	Water	.07	1 (6.94)
Iced tea	.15	2 (8.52)	Yogurt	.05	1 (6.61)
Jam	.17	3 (10.37)			

Note: Concentration measured by Herfindahl-Index (Adelman, 1969), based on product market shares.

3.8.12 Analysis of the Neural Network Product Embedding


Figure 3.14 depicts a two-dimensional t-SNE projection of the product embedding W_H trained on empirical data. Each dot represents a product, and the colors indicate retailer categories.

The product map is structured and contains several clusters. A closer inspection of the product clusters reveals three patterns:

1. Clusters can be perfectly aligned with a retailer category (e.g., cluster A: chocolate bars).
2. Clusters can contain products from several retailer categories (e.g., cluster B: soft drinks, juices, and (flavored) water).
3. A retailer category can be split into several product clusters (e.g., cluster C: coffee is split into regular ground coffee, coffee beans, and coffee capsules).

In contrast to the simulation (Figure 3.7, main text), we find that product clusters overlap. Possible explanations are that product categories in the real world are not as well-defined as in the simulation or that the manual category definitions are imperfect. Nonetheless, we see that the neural network learns a meaningful representation of products that should allow the model to efficiently incorporate cross-product effects.

Figure 3.14. t-SNE Projection of the DNN Product Embedding W_H .




Note: Best viewed in color.

3.8.13 Descriptive Analysis of DNN Probabilities

For the empirical application the true purchase probabilities are not known. A plot based on the dichotomous outcomes would be difficult to interpret, so we use an alternative approach to comparing the vector of true outcomes with the predicted probabilities. The first step is to create 20 equally large probability windows $w \in \{[0, 5\%], (5\%, 10\%], \dots, (95\%, 100\%]\}$. For each window we then compute the average predicted purchase probability \bar{p}_{itj}^w and the observed purchase rates \bar{b}_{itj}^w , defined as the average of the purchase indicators within the window. The averaging allows us to plot the purchase rates as a function of the average predicted purchase probability in a scatter plot. The Beta distribution is the conjugate distribution to the binomial distribution, so we compute 95% confidence interval for the purchase rate by sampling from the Beta(a, b), with the shape parameters $a = \sum_{itj} b_{itj}^w$ and $b = \sum_{itj} (1 - b_{itj}^w)$. For the predicted probability we use a nonparametric bootstrap with 100 replications. The scatter plot in Figure 3.15 shows that all points of the scatter plot are close to the bisecting line of the first quadrant. Most of the error bars include the line with slope 1 and intercept 0. We conclude that probabilities are scaled properly and that the learned probabilities reflect true purchase rates well.

Figure 3.15. Observed purchase rates as a function of predicted probabilities.



4 | The Impact of Personalization on Coupon Performance

Sebastian Gabel, Daniel Guhl, Daniel Klapper

Abstract

Coupons are an integral tool in the sales promotion mix of grocery retailers. It is therefore not surprising that the number of coupons has steadily increased over the last decade. In the face of decreasing redemption rates and coupon profitability, personalization through real-time offer engines is a viable approach to increase coupon effectiveness. To this end, the authors study a rich data set from a leading German grocery retailer that comprises market basket data, loyalty card data, and customer responses to 12 million personalized coupons across 1,116 brands in 115 product categories. For almost 1 million coupons, the brand and the discount were randomized, so the exogenous variation pertaining to both dimensions of the coupons facilitates an unbiased measurement of the effect of decision variables on customer responses. This study quantifies the effect of targeting on redemptions, revenues, and profits. Targeting increases redemption rates by 64.0%, revenues by up to 182.2% and profits by up to 111.8%, compared to non-targeted coupons. The impact of targeting on coupon effectiveness varies significantly across categories and brands, and much of the variance can be explained by brand and category characteristics, such as brand loyalty, price position, and purchase frequency. This research helps retailers to use targeting engines more efficiently. The results underline the benefits of sophisticated systems for automated one-to-one marketing and allow retailers to carefully compare the costs associated with implementing personalization engines to the financial benefits that such systems offer.

Keywords

targeted coupons, sales promotions, real-time offer engines, recommender systems

4.1 Introduction

In 2015, \$550 billion in coupons were distributed, an increase of 3.2% compared with the previous year (Valassis, 2016). In grocery retailing, only .6% of all distributed coupons are redeemed, and redemption rates decreased compared to 2017 (NCH Marketing Services, 2019). At the same time, coupon distribution is costly. Freestanding inserts (FSI) represent over 90% of the grocery coupons distributed in 2018, and the estimated distribution cost per coupon redemption is \$.35 (Biafore, 2016). Furthermore, increasing face values decrease coupon profitability (Valassis, 2016) and customers often redeem coupons for products for which they would have been willing to pay the regular price (Forrester, 2017).

These challenges motivate firms to seek new ways to address low coupon redemption rates and improve the impact coupons have on revenues and profits. Many retailers collect vast amounts of customer-level data, which they use to analyze customer purchasing habits (Blattberg et al., 2008; Bradlow et al., 2017). These data allow retailers to tailor coupons to customer segments or individuals, thereby enhancing coupon effectiveness (Rossi et al., 1996; Zhang and Wedel, 2009; Ailawadi and Gupta, 2014). Still, few retailers personalize product recommendations and discounts on a substantial scale. In grocery retailing, for example, segment and mass marketing (e.g., coupons and circulars) are still the dominant promotional strategies. Combining available customer data, especially those obtained through loyalty programs, with the advanced techniques offered by solution providers such as dunnhumby or Catalina Marketing, offers great potential for promotion personalization efforts (Rowley, 2005; Guillot, 2016). Real-time offer (RTO) engines allow retailers to leverage advanced analytics to derive personalized offers on the basis of real-time customer interactions and purchase histories. According to industry experts, “retailers will increasingly transfer promotional activity from traditional media [...] to more targeted and real-time offers” (Gartner, 2016, p. 20), especially considering that such RTO engines are used extensively in targeted online advertising (Chapelle et al., 2015).

The main challenges for retailers include the implementation of complex targeting algorithms, scaling these algorithms to the full assortment, and communicating personalized offers to individual customers in real-time (Naik et al., 2008; Gartner, 2016). At the same time, it is difficult to predict how RTO engines impact revenues and profits and whether these justify the implementation costs.

In this study, we seek to assess whether coupon personalization through RTO engines, similar to personalization in online advertising recommender systems, can increase coupon performance in grocery retailing. We base our study on a rich data set collected at a leading German grocery retailer. The data contains loyalty card transactions, market basket data, and 12 million coupons for 1,116 brands in 115

categories from the retailer’s RTO engine. For approximately one million coupons, the promoted brand and the discount were selected randomly, so the comparison of targeted and non-targeted coupons allows us to isolate the effect of targeting on coupon performance. The exogenous variation in the random coupon data makes it possible to measure the drivers of customer responses to coupons in regression models without bias. The models allow us to run policy simulations that lead to a deeper understanding of coupon performance. Specifically, we (1) quantify the impact of personalization on coupon redemption rates, revenues, and profits, (2) evaluate how brand and category characteristics affect the redemption probability uplift through personalization, and (3) compare the financial performance of personalized coupons to those of traditional promotional strategies.

With this article we contribute to prior research on promotion personalization and coupon effectiveness in three ways. First, we add to prior research that studies the effect of promotion personalization on coupon performance (Rossi et al., 1996; Zhang and Wedel, 2009; Venkatesan and Farris, 2012). In our data set, customers received targeted coupons (personalized in terms of both product selection and discount) or random coupons. A comparison across these conditions isolates the effect of personalization on coupon redemption rates without bias from confounders. In contrast with previous studies that analyze the effects of either brand personalization (Venkatesan and Farris, 2012; Osuna et al., 2016) or discount differentiation (Rossi et al., 1996; Zhang and Wedel, 2009), we simultaneously assess how brand and discount personalization influence coupon redemptions. And while previous studies conclude that the benefits of personalization in offline retailing are limited (e.g., Zhang and Wedel, 2009), we find clear evidence that personalization yields substantial increases in all performance metrics.

Second, the RTO engine studied here promotes products from many categories and brands. This allows us to analyze the impact of coupon personalization across the retailer’s whole assortment. While most studies focus on a small number of categories, the wider scope of our analysis supports a better generalizability of our findings. Additionally, we extend prior research that details how brand and category characteristics affect price elasticities, coupon redemption rates, and (incremental) sales (e.g., Narasimhan et al., 1996; Osuna et al., 2016) in that we uncover how brand and category characteristics explain variation in the impact of personalization.

Third, we provide novel insights on coupons distributed through in-store kiosks, which have not been studied in marketing literature. In contrast to coupons that are distributed in FSIs, at the checkout counter or through a mobile app, kiosk coupons are more similar to in-store “surprise coupons” (Heilman et al., 2002), which are intended to be used immediately.

With these measures, our research provides relevant insights for practitioners. We show that one-to-one marketing increases redemption rates, revenues, and profits, so personalization can be a valid answer to the challenges retailers face in the context of coupon optimization. The collaborating retailer and solution provider also noted the substantial costs associated with establishing loyalty programs, customer databases, and RTO engines, highlighting the need for a better understanding of coupon effectiveness to justify these investments. The insights gleaned from our study contribute to a better understanding of the value of one-to-one marketing and guide retailers in assessing the benefits of RTO engines. Accordingly, retailers can make more educated investment decisions when launching RTO engines and building effective personalized coupon programs. Understanding which brands and categories are most suitable for targeted coupons will help category and promotion managers select the right products to include in RTO engines, according to their specific goals. Insights pertaining to the link between the degree of personalization and the efficiency of RTO engines help retailers to assess the financial benefits of RTO engines. The close co-operation with the retailer allows us to analyze modern promotion technology in the field and enhances the external validity of our findings. The scale of the data set supports a precise measurement of the effects of personalization of redemption rates, revenues, and profits and contributes to the generalizability of our findings.

The remainder of the paper is structured as follows. After a literature review that highlights our research contribution, we present the data set and the model used to study customer responses to coupons. We then investigate the impact of personalized coupons on redemption rates, revenue, and profit. We conclude by summarizing our main findings, discussing managerial implications, and providing directions for further research.

4.2 Related Work and Contribution

Our study pertains to two literature streams in marketing: (1) studies that develop methods for targeted coupons and conceptualize/measure their effects on coupon performance and (2) studies of promotional effects and their drivers.

4.2.1 Targeted Couponing—Methods and Effects

In describing one-to-one target marketing, Rossi et al. (1996) show how to model, measure, and optimize price discounts in a brand choice setting, highlighting that household purchase histories are valuable to manufacturers for optimizing coupon profitability, according to their study in a single product category. Zhang and Wedel (2009) build on their work to study two kinds of customized checkout coupons (loyalty and competitive) at three levels of granularity (market, segment, and individual) in online and offline stores across two product categories. They find

that promotion optimization leads to substantial profit improvements and that loyalty (competitive) coupons are more effective in online (offline) stores. Yet, compared with segment-level promotions, the incremental profit of individual-level promotions appears small, in particular in offline stores. Instead of taking a brand perspective, we adopt the view of a retailer that sells many product categories, so personalization includes brand selection choices. We build on Rossi et al.'s (1996) results and add the dimension of product recommendations to the assessment of targeted coupons. Unlike Zhang and Wedel (2009), we provide reliable evidence of the value of individual-level personalization in offline retail settings. Combining kiosk coupons and full personalization (including brand selection) can overcome low redemption rates, which have been a “major impediment to the success of customized promotions” in offline stores (Zhang and Wedel, 2009, p. 204).

Heilman et al. (2002) examine the impact of in-store “surprise” coupons on total basket value. Because redemption rates are higher for surprise than for FSI coupons, they represent a promising promotion tool for retailers. According to Heilman et al. (2002), unexpected coupons have an income effect that increases basket size and unplanned (impulse) purchases in non-promoted but related categories. Kiosk coupons are similar in that they are distributed at the point of sale, and customers do not know in advance which coupons they will receive. Our setting provides further evidence of the value of in-store coupons. We add to the research by Heilman et al. (2002) in that we analyze a larger set of brands and categories, include the analyses of both planned and unplanned purchases, and assess the impact of brand selection and price differentiation on coupon performance.

Venkatesan and Farris (2012) present a conceptual framework for retailer-customized (email) coupon campaigns; in a quasi-experiment, they find that coupon exposure and redemption have positive effects on trip incidence and revenues. The positive exposure effect implies that sales increases might result even from non-redeemers. Sahni et al. (2016) evaluate the revenue effect of personalized email promotions in a field experiment on an online ticket resale platform. They find a 37.2% revenue increase that is especially strong for individuals who did not transact on the platform in the year before the experiment. The redemption itself does not explain the majority of the effect and the authors conclude that emailed offers also serve as “advertising” in addition to being a promotional tool. Although Venkatesan and Farris (2012) and Sahni et al. (2016) present clear empirical evidence of the economic consequences of customized coupon campaigns, they do not address the impact of discount personalization and the impact of targeting on redemption rates. We measure the increase in redemption rates, revenues, and profits for targeted versus non-targeted coupons and determine the drivers of difference across brands and categories.

In a related stream of literature, researchers studied the effects of personalization in the context of online advertising. Lambrecht and Tucker (2013) analyze dynamic retargeting of online advertisement, Tucker (2014) studies the impact of advertising personalization on Facebook, and Bleier and Eisenbeiss (2015) focus on the personalization of banner advertising. The RTO engine studied in our research is very similar to targeting approaches used in online advertising. Our research therefore creates a link between research on targeted coupons and research on recommender systems in online advertising. We show that RTO engines allow retailers to implement product recommendations and personalized discounts at scale, and we provide a first holistic assessment (across multiple categories and brands) of the effectiveness of such systems by disentangling the effects of the promotion channel (in-store coupons) and targeting. In the assessment of system effectiveness we go beyond the simple evaluation of redemptions (cf. clicks in online advertising) but study revenue and profit implications as well.

4.2.2 Promotional Effects and Their Drivers

A number of empirical marketing studies document promotional effects; the heterogeneity in these findings has motivated researchers to assess how promotional effectiveness depends on market, category, or brand characteristics. For example, in Bolton's (1989) study of the promotional price elasticities of twelve brands in four categories, brands with higher price elasticities exhibit less category and brand display activity, a lower market share, and more category couponing and feature activity. Raju (1992) analyzes the temporal variability of category sales for more than 200 brands from 25 categories and finds that greater discount magnitude (frequency) increases (decreases) sales variability. On the other hand, product categories that are bulky (which make stockpiling and transportation more difficult) and more competitive show less sales variability. Narasimhan et al. (1996) study the effects of product category characteristics on promotional price elasticities (price, feature, and display promotions) and find, beyond the effects of typical category characteristics (e.g., penetration), that promotional elasticities are higher in categories in which products are easier to stockpile and in "impulse" categories (though not to a significant extent). Bell et al. (1999) also investigate the effects of category, brand, and customer factors on price promotion effects (decomposed by primary and secondary demand). We extend this line of research by analyzing how category and brand characteristics influence the increase in redemption rates due to personalization. We leverage a data set that contains both targeted and non-targeted coupons. The exogenous variation in the random coupon data pertaining to both dimensions of coupons (i.e., brand and discount personalization) facilitate the unbiased measurement of the effect of decision variables on customer responses, contributing to the external validity of our findings.

Osuna et al. (2016) study the effects of brand and category characteristics on the performance of two types of checkout coupons (loyalty and cross-category), targeted such that eligibility to receive the coupons depend on the household’s purchase history. For 893 coupons, they fixed the discounts within each coupon type (10% for loyalty, 20% for cross-category). We study coupons that are targeted in both dimensions (discount and brand), and evaluate the effects of targeting on revenues and profits. Osuna et al. (2016) also highlight the need to study coupon effects for alternative distribution channels such as in-store kiosks, as we address herein.

4.3 Setup

4.3.1 Data Set

We obtain data from a leading German brick-and-mortar grocery retailer. The sole purpose of the retailer’s loyalty program is to collect customer-level data and distribute personalized coupons. A coupon is uniquely identified by the promoted brand and its discount value. To personalize coupons, the retailer and its target marketing solution provider implemented an RTO engine for 147 stores in one of Germany’s largest cities. Similar to CVS’s ExtraCare Coupon Center, customers scan their loyalty card at in-store kiosks and receive a printout that contains up to seven brand coupons. By collating the available coupons, customer-specific discounts, and corresponding (predicted) redemption probabilities, the RTO engine scores all brand-discount combinations for each user and selects coupons, with the goal of triggering additional purchases and increasing customer loyalty. In the context of this study, targeting therefore refers to selecting the *brand* and *discount* for each customer, based on past transaction data. In other words, the RTO engine determines which subset of the customer population should receive a given brand coupon and at what discount. Coupons are valid on the same shopping trip and are redeemed automatically if a customer purchases any of the promoted products and scans his or her loyalty card during checkout.

Before using the data to study customer responses to personalized coupons, we pruned the raw data in three steps. First, we removed observations for which the coupon printout occurred after the shopping basket was recorded. Coupons are only valid on the day of the printout so coupons printed just before customers leave the store have a redemption rate of zero by design. Second, we discarded observations for new loyalty card users, that is, customers without purchase histories. Without past transactions coupons cannot be personalized, so these observations are meaningless to our study. Third, we only keep the first observation for each customer/coupon combination. Table 4.1 summarizes the most important characteristics of the final data set. The data set spans over 72 weeks (11/2015 to 03/2017) and contains

Table 4.1. Summary of data set statistics.

Variable	Value
Time window (# of weeks)	11/2015 to 03/2017 (72)
# of stores	147
# of customers	217,299
# of distributed coupons (random)	11,697,018 (750,525)
Total coupon face value (redeemed coupon face value)	€7,105,989 (€257,581)
# of distinct brands	1,116
# of distinct categories	115
Average # of promoted brands per week (SD)	232.1 (32.4)
Discount range	[10%, 50%]

a total of twelve million coupons across a large number of brands and product categories. The minimum discount was 10% for all brands; the maximum discount varied between 30% and 50%, depending on the brands' average circular discount in the previous calendar year. These discount values are typical for coupons in grocery retailing. Based on the regular prices for the promoted brands (90% are between €0.75 and €3.99), the coupons had a total face value of €7.1 million. Eleven million coupons were targeted so only a small subset of the total customer population received the specific brand-discount combination.

The retailer promoted different brands at different points in time. On average, the RTO engine personalized 232.3 brands each week, and brands were promoted for 10.6 weeks. If a brand was featured in the retailer's weekly promotion circular or on in-store displays, coupons were deactivated for the time of the circular/display promotion. For our analysis, this means that we can measure customer reactions to coupons without the direct confounding effects of traditional promotion instruments. Spillover effects were avoided because coupons are only valid during the immediate shopping trip and with our focus on targeted coupons within the loyalty program, self-selection by customers is not an issue. Overall, the summary statistics underline the breadth and depth of the data set and support the generalizability of our results.

4.3.2 Targeting Policy

To better understand how the RTO engine targets coupons, we first analyze the retailer's coupon targeting policy. As a part of the coupon data, the retailer provided a brand-level score that aggregates the engine's understanding of the customers' (time-dependent) individual purchase and redemption likelihoods for brand coupons. This variable is built in the solution provider's recommender system from past coupon transactions, market basket data, and loyalty card data and is fundamental to the targeting algorithm. We use binary logistic regression to model

Table 4.2. Results for binary logit models to explain targeting policy.

	DV: coupon is targeted	DV: discount for targeted coupon is smaller than avg. discount for random coupon
Brand score	1.476 ***	2.008 ***
Brand fixed effect	yes	yes
N	718,068	685,647
Log-likelihood LL (LL0)	-82,958 (-100,257)	-165,785 (-262,353)

Note: Sig. label: *** $p < .01$.

(1) whether a customer receives a targeted coupon or a random (i.e., non-targeted) coupon and (2) whether the targeted discount is smaller than the average discount for random coupons of the same brand. The customers who receive a random coupon are a representative (random) subset of the customer population that does not receive a targeted coupon for the given brand. These customers are a good reference point for the analysis of the targeting mechanism. In addition to the brand score we include brand fixed effects. For the estimation, we randomly sample data from 50 brands and standardize the brand score variable within each brand.

Table 4.2 depicts the results for the targeting models. As expected, the effect of the brand score variable is positive in both models, such that a larger brand score (i.e., the proxy for brand preference) leads to a higher likelihood of receiving a targeted coupon for this brand and a higher likelihood of receiving a lower discount. This result is intuitive: The RTO engine targets customers with brands that fit the customers' preferences. Additionally, the engine takes into account that these customers should already have a higher willingness-to-pay and smaller discounts should be sufficient. This general mechanism is prototypical for targeting and price differentiation algorithms presented in marketing literature (e.g., Rossi et al., 1996). Therefore, there is good reason to believe that the results of our analysis generalize this application to other RTO engines that follow the same general mechanism for promotion personalization.

4.3.3 Descriptive Analysis of Redemption Rates

To begin the analysis of targeted coupons, we provide a descriptive analysis of redemption rates (Table 4.3). Across all random coupons, the average redemption rate is 1.528% (SE = .014%). Compared to this, the average redemption rate is 2.728 percent points (+178.4%) higher for targeted coupons (4.257%, SE = .006%). The difference is significant at $p < .01$. In this context, it is important to recall that targeting is endogenous. For random coupons, the distribution of printing

Table 4.3. Descriptive analysis of redemption rates.

Variable	All Coupons	Targeted Coupons	Random Coupons
Avg. redemption rate (SE)	4.082% (.006%)	4.257% (.006%)	1.529% (.014%)
# of coupons	11,697,017	10,946,493	750,525
Avg. discount (SD)	30.4% (11.1%)	30.9% (11.0%)	23.7% (10.0%)
Avg. # coupons per brand (SD)	10,481 (13,819)	9,809 (13,352)	673 (1,414)
Avg. discount per brand (SD)	29.1% (8.1%)	30.1% (8.3%)	23.0% (5.4%)

frequency across brands is uniform and independent of customer preferences and the discount distribution is approximately uniform (across the possible discount levels). Targeted coupons are directly tailored to customer preferences, in that brands preferred by customers are printed more often. For targeted coupons, the discounts depend directly on customers' preferences. Typically, only a small fraction of customers notes a strong preference for a given brand, so the distribution of discounts is skewed toward higher values. Average discounts of targeted coupons are thus 7.2 percent points higher than those of random coupons (+30.4%). In Section 4.4 we introduce a modeling approach that allows us to deepen this first descriptive analysis. Section 4.5 analyzes redemption rates for both types of coupons in more detail, also by directly accounting for differences in discounts through a modeling approach. Section 4.6 focuses on the financial impact of targeting coupons by evaluating implications for revenues and profits and a comparison to non-targeted (mass market) promotions.

As a side note, the observed redemption rates are much higher than the industry average for redemption rates of checkout coupons and coupons in freestanding inserts at the same retailer before the introduction of the loyalty program and the RTO engine—redemption rates were approximately .5%, similar to the values reported in NCH Marketing Services (2019). A likely reason for this is the coupon distribution channel. In-store coupons, in this case distributed through in-store kiosk systems, are known to have higher redemption rates than coupons distributed before the shopping trip (Heilman et al., 2002).

4.4 Approach

4.4.1 Model

To measure the impact of the RTO engine's coupon personalization on coupon redemptions, revenues, and profits, we make use of the fact that two different types of coupons were distributed to customers. Recall that for targeted coupons brands

and discounts are personalized according to the customers' purchase histories. For random coupons, the retailer randomized brands and discounts at the coupon level. The analyses in the following sections rely on results from two models, one for each type of coupon. In both data sets, we estimate a model that predicts redemption probabilities as a function of discounts and the brand score. We use binary logistic regression with random effects. The probability $pr^{R/T}$ that customer i redeems a random (R) or a targeted (T) coupon for brand b in store s at time t is

$$pr^{R/T} (y_{ibst}^{R/T} = 1) = \frac{1}{1 + \exp\{-u_{ibst}^{R/T}\}}, \quad (4.1)$$

where the utility function (for simplicity, we omit the data set labels here)

$$u_{ibst} = \alpha_0 + \alpha_s + \alpha_t + \alpha_b + \gamma bs_{ibt} + (\beta_0 + \beta_b)d_{ibst} \quad (4.2)$$

depends on the model intercept α_0 , the average discount effect β_0 , store random effects $\alpha_s \sim N(0, \sigma_s)$, year-week random effects $\alpha_t \sim N(0, \sigma_t)$, correlated brand random effects and random discount coefficients $[\alpha_b, \beta_b]' \sim MVN(0, \Sigma_b)$, the discount d_{ibst} , and the effect of the customer-, brand-, and time-specific brand score γbs_{ibt} . We estimate a separate model for each type of coupon (see Appendix 4.8 for a discussion of five nested model specifications).

Given that we only have two continuous covariates, binary logistic regression with random effects is a good model choice because it is the most parsimonious model that fully leverages the strength of our data set. It also offers ‘‘borrowing strength’’ across brands, which is essential for brands with few observations. Because we only use one observation for each customer/brand combination, we cannot estimate a model that accounts for *unobserved* heterogeneity. However, given that the brand score for each customer and brand (and time) is available in the data and we know from the analysis of the targeting mechanism, that this variable plays a key role in the RTO engine, it is well-suited for modeling *observed* customer heterogeneity. The brand score should be positively related to the coupon redemption probability. The discount variation for random coupons is exogenous, so price endogeneity is not an issue when we model redemptions. For targeted coupons, the discounts are endogenous and related to the customer's brand preferences. Including the brand score in the model for targeted coupons mitigates the endogeneity issue to some extent. More importantly, we only use the model for targeted coupons to predict redemption probabilities of observed coupons and discounts in-sample and we do not claim to estimate a causal effect. Differences across brands are accounted for by the brand-level random intercepts α_b , so we standardize brand scores within a brand to ensure that we only explain differences within brands. In the estimation,

Table 4.4. Estimation results for redemption models.

Variable	Random Data		Targeted Data	
	Est.	Sig.	Est.	Sig.
Intercept α_0	-5.579	***	-4.116	***
Discount β_0	3.928	***	1.430	***
Brand-score γ	.523	***	.492	***
SD(Brand) α_j	.896	***	.865	***
SD(Discount) β_j	1.183	***	1.612	***
Cor(α_j, β_j) ρ	-.784	***	-.782	***
SD(Year-Week) α_t	.222	***	.301	***
SD(Store) α_s	.133	***	.096	***
N	750,876		750,876	
LL	-55,370		-123,264	

Note: Sig. label: *** $p < 0.01$.

we randomly subsample the full data for targeted coupons to the same size as the random coupon data set to speed up the estimation and simplify the model comparison.

Table 4.4 summarizes the estimated coefficients for the redemption models. All model coefficients are significant at $p < .01$. For the model estimated on the random coupon data, the average discount effect is positive, as expected for price-offs. The average of the brand-specific price elasticities is -2.96 ($SD = .51$), with 90% of the values in $[-3.74, -2.12]$. This result is in line with the promotional price elasticities for grocery products (accounting for price endogeneity) reported in the meta-analysis by Bijmolt et al. (2005). As expected, the effect of the brand score is positive, such that higher brand scores result in higher redemption probabilities for the corresponding brands. The standard deviations (SD) of three random effects are all relevant in magnitude. We observe the largest heterogeneity in the brand dimension, followed by the dimensions store and week. It makes sense that redemption probabilities vary over stores and weeks, but the variation over brands should be larger, given that we analyze 1,116 brands from 115 categories. Customers are known to be less price sensitive when it comes to attractive brands (Bolton, 1989), so the negative correlation between the brand random effect and the random price coefficient is intuitive.

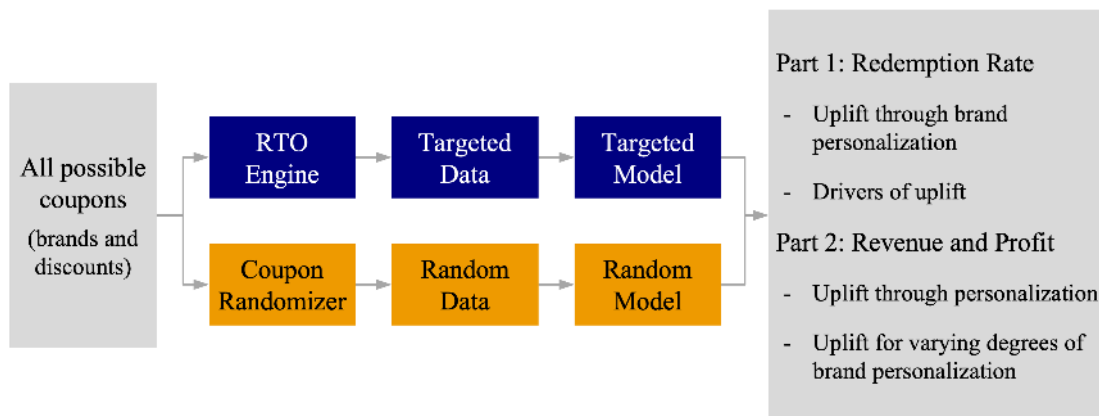
The results for the targeted coupons are quite similar. The signs for the estimates are the same compared to the model for random coupons. It is noteworthy that the discount parameter is considerably lower in magnitude, which translates into

average brand-specific price elasticities of about $-.94$ ($SD = .60$). This is in line with the results in Section 4.3.2: the RTO engine sets prices according to the brand preferences and the price sensitivities of the customers, such that customers with higher brand scores receive lower discounts. Hence, even after controlling for brand preferences via the brand score variable, the observed reaction to discounts is lower for targeted coupons, compared to random coupons. As mentioned above, we use the model based on targeted coupons only for in-sample predictions, so this downward “bias” is irrelevant. The other estimates have a similar magnitude as the estimates in the model that is based on random data.

4.4.2 Analysis Overview

Combining the predictions of the two models is the basis for studying the effect of personalization by comparing the outcomes for different targeting mechanisms (e.g., targeted and random). The clean and exogenous variation in the random coupon data is the foundation for evaluating outcomes for unobserved coupon policies. Using a parametric model (instead of a nonparametric approach) enables us to control for confounding factors and to run simulations of promotion policies which were not observed in the data. We leverage this in deepening our understanding of coupon personalization.

Figure 4.1 systematically summarizes the main steps of our approach. In part 1 (Section 4.5) we use the estimated models to compare the coupon redemption probabilities for targeted coupons with the redemption probabilities of the random baseline. In the coupon data set, two factors lead to higher redemption rates in the case of targeted coupons: the targeting itself and the higher average discount (see Table 4.3). The exogenous variation of discounts in the random coupon data set enables us to control for the latter by stratifying the discount distribution for random coupons, so it equals the discount distribution for targeted coupons. This isolates the redemption rate uplift through targeting. We extend the redemption probability analysis by investigating the systematic differences in redemption probabilities across brands and categories. In part 2 (Section 4.6), we focus on revenues and profits, thereby measuring the financial impact of targeting. Both metrics directly penalize for larger discounts so we can analyze both dimensions of targeting (brand and discount) simultaneously. In doing so, we compare the RTO engine targeting (i.e., one-to-one marketing) to mass marketing policies for which all customers receive the same coupons and discounts. We then explicate how the selectiveness of brand targeting affects financial outcomes by systematically decreasing the size of the sub-population that is targeted with brand coupons. As in the first part of our analysis, we rely on the exogenous variation in the random coupon set to evaluate outcomes under unobserved pricing and brand targeting policies.

Figure 4.1. Analysis overview.

Note: Best viewed in color.

The two parts differ in their outcome variables (redemption probabilities in part 1, revenues and profits in part 2) but they share one important similarity: The first step in each part focuses on the performance of the *specific* RTO engine that produced the targeted coupons studied here. The second step then *generalizes* the insights by studying how mediator variables explain the variation in redemption rates and how the degree of personalization affects financial outcomes. This widens the applicability of our findings.

4.5 Part 1: Redemption Rate Analysis

4.5.1 Redemption Rate Uplift Through Brand Personalization

When comparing redemption rates between random coupons and targeted coupons it is important to keep in mind that random coupons and targeted coupons can have very different discount distributions. Discounts for random coupons are sampled from all allowed discount values for each brand b , using uniformly distributed weights. Each discount level $d_l^b \in \{d_1^b, \dots, d_{n_b}^b\}$ is therefore observed with a frequency of approximately $1/n_b$, where n_b is the number of distinct discount levels for b . The face values of targeted coupons are selected from the same set of discounts, but the RTO engine picks discount levels with different frequencies. As discussed in Section 4.3.3, the observed average discounts are 23.0% for random coupons and 30.1% for targeted coupons. Higher discounts lead to higher redemption probabilities, so it is necessary to stratify the discount distributions to ensure a fair comparison between the two coupon types.

To accomplish this, we use the models presented in Section 4.4 to predict redemption probabilities for both types of coupons according to Equation 4.1. The key idea of this approach is that random and targeted coupons are distributed to different customer sub-populations. The part of the overall population that is

Figure 4.2. Approach for redemption rate comparison with stratified discount distributions.

$$\begin{array}{l}
 \text{Targeted} \\
 \text{redemption} \\
 \text{probabilities} \\
 \widehat{pr}_{ibst}^T = \text{Inverse} \\
 \text{Logit} \left(\begin{array}{l} \alpha_{0bst}^T + \Gamma_{ibst}^T + d_{ibst}^T \times (\beta_0^T + \beta_b^T) \end{array} \right) \\
 \\
 \text{Random} \\
 \text{redemption} \\
 \text{probabilities} \\
 \text{(stratified discounts)} \\
 \widehat{pr}_{ibst}^R = \text{Inverse} \\
 \text{Logit} \left(\begin{array}{l} \alpha_{0bst}^R + \Gamma_{ibst}^R \text{ (sampled)} + d_{ibst}^T \times (\beta_0^R + \beta_b^R) \end{array} \right)
 \end{array}$$


Targeted
Random

Note: Best viewed in color.

exposed to random coupons produces the random data set and vice versa. By training two separate models we infer two sets of model parameters for the two sub-populations. The price variation in the random data set is exogenous, so the measured discount sensitivity captures the effect of discounts on redemption probabilities. This makes it possible to predict redemption probabilities for the random coupons population, assuming that prices are distributed as they are in the population that receives targeted coupons. The difference in redemption probabilities isolates the effects of brand targeting (identified by the different responses to coupons with the same discounts in the two sub-populations).

To simplify the notation, we group all variables except the utility contribution of discounts d_{ibst}^T and the customer-specific brand score $\Gamma_{ibst}^{T/R}$ (see Figure 4.2). We denote the utility offset that does not depend on customer variables or discounts $\alpha_{0bst}^{T/R}$. For both models, we derive the predictions using the discounts selected by the RTO engine, d_{ibst}^T . The brand scores $\Gamma_{ibst}^{T/R}$ and discount sensitivities $\beta_0^{T/R} + \beta_b^{T/R}$ are either based on the random data (R) or on the targeted data (T). The size of the random data set is smaller than the size of the targeted data set, so we sample a value for each discount d_{ibst}^T from Γ_{ibst}^R , using uniform sampling weights. The prediction then yields 10,946,493 predicted redemption probabilities $pr_{ibst}^{T/R}$ for both random and targeted coupons.

The scatter plot in Figure 4.3 depicts the predicted redemption probabilities for both coupon types. The bubbles represent the brands' average redemption probabilities and the horizontal/vertical lines indicate the average redemption probabilities across all observations. The average redemption probabilities for random coupons and targeted coupons are 2.59% and 4.25%, respectively (all SE < .0001). Targeting (in the brand dimension) leads to a redemption probability

Figure 4.3. Redemption probability comparison.

Note: Best viewed in color.

increase of 64.0%. The model-based analysis allows us to explicitly control for differences caused by the respective discount distributions, so the redemption probability increase is smaller than the direct comparison of redemption rates on the raw data (178.5%, Table 4.3). It is not very surprising that the RTO engine makes coupon redemptions more likely because coupons are personalized based on the customers' past purchases. Nonetheless, the results underline that coupon personalization is feasible at scale, even for a large number of categories and brands.

Figure 4.3 also shows that targeted coupons have higher average redemption probabilities than random coupons at the brand level, and we note only a few exceptions. For most brands, redemption rates are a factor 1.5 higher than for targeted coupons, and for a significant number of brands we even observe more than three times larger redemption rates. The setup of the system studied here requires that each customer receives eight coupons on a coupon printout. That almost all brands have higher redemption probabilities for targeted coupons underlines that the pool of available brands in the RTO engine (i.e., on average 232.1 brands per week) is large enough to find eight brands for each customer that align with his

or her preferences. The increase in redemption probabilities through targeting is particularly strong for brands that appeal to a narrower target audience and are therefore distributed to fewer customers (smaller N in Figure 4.3). Regressing the brand-level ratio of average probabilities between targeted coupons and random coupons on the number of distributed coupons (divided by 1,000) yields a slope of $-.018$ ($p < .01$).

4.5.2 Drivers of Redemption Rate Uplift

The comparison of random and targeted coupons reveals that the redemption probability increase resulting from brand personalization (i.e., the vertical distances of bubbles to the 45° line in Figure 4.3) varies significantly across brands. To explicate which category and brand characteristics affect this uplift (and to what extent), we study the differences in redemption probabilities across brands and categories using linear regression.

The *dependent variable* in the linear model, y_{bc} , is the difference between the redemption log odds for targeted and random coupons (i.e., the log odds ratio), averaged by brand. The number of observations differs across brands, so we use the inverse squared standard errors (SE) as weights in the regression analysis to account for the varying precision of the log odds ratios (Schwarzer et al., 2015). We calculate SEs using a nonparametric bootstrap with 1,000 iterations.

The explanatory variables used in prior studies of price elasticities and coupon redemption rates (e.g., Bell et al., 1999; Osuna et al., 2016; Narasimhan et al., 1996) are similar and highly correlated, so we use the 13 category- and brand-level variables that are relevant in our context and that do not create empirical issues in the model estimation (see Table 4.5). We derive the brand variables (X_{bc}) and category variables (Z_c) from the retailer's sales and loyalty card data and measure the stockpile and impulse scores on the scales from Narasimhan et al. (1996) in an Amazon Mechanical Turk survey. Further details regarding the variable operationalization are available in Appendix 4.8.

For most categories, we observe multiple measurements. Rather than treating measurements as independent, we follow (Bijmolt and Pieters, 2001) and use a random effects model that can account for the nested structure of the data. The full regression model is given by

$$y_{bc} = \alpha_0 + \beta X_{bc} + \gamma Z_c + \sum_{t=1}^5 \delta_t + \alpha_c + e_{bc}. \quad (4.3)$$

The two error components α_c and e_{bc} are normally distributed with zero mean and SDs of σ_c and e_{bc} . Note that σ_c and e_{bc} vary on different levels. The first

Table 4.5. Meta-regression estimation results.

Variable	Operationalization	Est.	SE	Sig.	
Intercept		1.238	.198	***	
Brand	Loyalty ¹⁾	Avg. number of purchases of the brand by users of brand	.066	.022	***
	Penetration	Fraction of customers who have bought brand product	-.075	.024	***
	Brand score range	RTO engine score range ($p_{5\%} - p_{95\%}$)	.422	.058	***
	Price position	Avg. brand price divided by weighted avg. across brands	.141	.045	***
	Deal depth ¹⁾	Avg. percentage promotion discount of products in brand	-.621	.061	***
	Promotion frequency	Promotion sales of brand products divided by total sales	.015	.014	
Category	Purchase frequency	Fraction of all trips in which category is purchased	-.069	.063	
	Private label share	Market share of private labels/generic brands in category	.497	.149	***
	Competition	Herfindahl index (brand market shares) in category	-.155	.051	***
	Price dispersion	Ratio of maximum and minimum regular price in category	-.001	.047	
	Price	Avg. dollars spent in category per shopping trip	.141	.059	**
	Stockpile score	Ability-to-stockpile scale score for category	-.005	.049	
	Impulse score	Impulse buying scale score for category	-.182	.051	***
SD(category random effect)			.142	***	
SD(residuals)			5.869	***	
Quarter fixed effect			yes		
LL			-852.999		
N			969		
R^2			.328		

Notes: 1) Normalized in category. Sig. labels: ** $p < .05$, *** $p < .01$.

error term, α_c , accounts for (random) variation between categories, whereas e_{bc} serves as an error term pertaining to the level of brands within categories. On each level of the model, we relate the log odds ratios y_{bc} to our explanatory variables. Five effect-coded year-quarter dummies δ_t (2015-4 to 2016-4, with 2017-1 as the reference) that indicate the (main) time window in which coupons for a brand were printed control for changing market conditions. We estimate the model coefficients by likelihood maximization (Hox et al., 2010).

Table 4.5 summarizes the estimation results. The R^2 value of 32.8% suggests that the model explains the variance in the log odds ratio well. The value is comparable to those reported by Osuna et al. (2016), who fit their models without category random effects. The SDs of the random components of the models show that (unexplained) variation between categories is lower than that within categories. Likelihood-ratio-tests for (nested) model versions that include no mediators or only brand- or category-specific variables reveal that both groups of variables are jointly significant ($p < .01$), and the proposed model is the best one.

In the discussion of the drivers, we focus on the variables that have a statistically significant effect on the redemption probability uplift. The effect of *brand loyalty* on the log odds ratio is positive and highly significant. Brands with higher (customer) loyalty typically have lower price elasticities in brand choice (Krishnamurthi and Raj, 1991), and promotions have greater potential to evoke purchases (Bell et al., 1999). For brands with high loyalty it is more important to reach appropriate customers, so a positive effect of brand loyalty on the uplift through personalization is plausible.

We observe a lower redemption probability uplift in the case of coupons for brands with a higher *customer penetration*. A larger customer penetration increases the pool of targetable customers, so redemption probability for both coupon types should be larger, all else being equal. At the same time, a larger target audience reduces the benefit of targeting, supporting the negative effect.

The opposite effect is true for the *RTO engine brand score range*. This variable can be interpreted as a proxy variable for the heterogeneity in a brand's attractiveness. More diverse customer preferences provide a better potential for personalization and increase the risk of reaching the wrong user in the case of random coupons, so the variable's impact on the redemption probability uplift should be positive. This is clearly the case.

The impact of a brand's *price position* (in a given category) is positive, such that the uplift in redemption probabilities through personalization is higher for expensive brands. Brands that have established a higher price position than other brands in the category should draw more customers when they promote, leading to higher primary and secondary demand effects (Bell et al., 1999). However, due to the surprise character of the in-store coupons, coupons for more expensive brands might have lower redemption probabilities *ceteris paribus*, and non-brand buyers might feel that the risk of buying the wrong brand is higher (Narasimhan et al., 1996). Given that the RTO engine targets customers according to their prior purchases, the negative impact of higher prices will be lower for targeted coupons, resulting in the observed positive effect of targeting.

The negative effect of the variable *deal depth* on the redemption probability uplift is not surprising. A higher percentage discount improves the quality-per-dollar equivalent of a brand and should induce primary and secondary demand effects (Raju, 1992; Bell et al., 1999). High discounts make offers more attractive and customers should be willing to redeem coupons even if the brand is not targeted well. On the other hand, the quality of targeting becomes (even) more important if the discount is low.

The *market share of private labels* within a category increases the measured redemption probability uplift. Marketing literature provides mixed results regarding the effect of the private label share on promotion effectiveness (Narasimhan et al., 1996). Nonetheless, we expect that categories with a high private label share should have higher redemption probabilities because coupons are more attractive for value-conscious customers. However, to switch such customers away from attractive private label products, good targeting is a prerequisite. In line with this, Osuna et al. (2016) find a significant positive effect for reward coupons but not cross-category coupons.

A similar argument holds for the *degree of competition*. In categories with low competitive intensity (i.e., highly concentrated categories, reflected by a high Herfindahl index), customers have well-established preferences, and it is harder to stimulate brand switching (Raju, 1992). This is particularly true for targeted coupons, because the potential pool of good brands is smaller, resulting in a negative relationship between the redemption probability uplift and the degree of competition.

The positive effect of *price* is in line with the brand-level variable price dispersion. Higher prices increase the perceived risk of buying the wrong product (Narasimhan et al., 1996). Yet, targeted coupons fit customer preferences so the RTO engine can counter the negative effect risk associated with higher prices.

The nature of the coupon channel can explain the negative relationship between *impulse score* and the redemption probability uplift. In-store coupons lead to category expansion effects due to the surprise character of the coupons (Narasimhan et al., 1996; Heilman et al., 2002). The surprise effect for targeted coupons should be smaller. In other words, the gains in redemption probabilities due to personalization are lower in impulse-buying categories, not because personalization is ineffective, but because even random surprise coupons work reasonably well.

In summary, the structured analysis of the redemption probability differences between random and targeted coupons reveals that redemption probabilities differ significantly across brands and that a number of brand and category characteristics impact the redemption probability uplift induced by targeting. Findings offer

face validity and can be explained in the context of prior research on the effect of promotions on consumer decisions.

4.6 Part 2: Revenue and Profit Analysis

4.6.1 Revenue and Profit Uplift Through Personalization

In addition to the analysis of redemption rates, it is important to evaluate revenues and profits as outcomes. This sheds more light on the monetary benefits that one-to-one marketing and RTO engines bring to retailers when replacing mass marketing promotion strategies. To make the results comparable between targeted and mass marketing promotions (e.g., circulars), we compare revenues and profits at the customer level and we assume that both promotions use the same distribution channel, that is kiosk systems.

For the targeted coupon, we use the redemption model from Section 4.4 and predict the redemption probabilities in-sample. We then use the predicted probabilities to compute the expected revenues and profits for each offer in our data set. The (expected) revenue r_{ibst} for customer i and brand b at time t is calculated by

$$r_{ibst} = \widehat{pr}(d_{ibst})p_b[1 - d_{ibst}]. \quad (4.4)$$

Here, p_b is the regular price of brand b and $\widehat{pr}(d_{ibst})$ is the predicted redemption probability as a function of the customer-specific discount d_{ibst} . The (expected) profit π_{ibst} is given by

$$\pi_{ibst} = \widehat{pr}(d_{ibst})p_b[1 - d_{ibst} - c_b]. \quad (4.5)$$

We set the cost factor c_b to the maximum allowed discount for each brand. The results for targeted coupons are benchmarked against the revenue and profit values computed under three mass market reference policies:

Circular with fixed discount All customers receive the same promotions, discounts are set to 25%, which (approximately) equals the average circular discount at the retailer.

Circular with revenue-maximizing discount All customers receive the same promotions and discounts are set to the revenue-maximizing value for each brand (within the possible bounds of the system).

Circular with profit-maximizing discount All customers receive the same promotions and discounts are set to the profit-maximizing value for each brand (within the possible bounds of the system).

Table 4.6. Comparison of RTO engine revenue with mass market promotions.

Policy	Brand	Discount	Revenue in Euro cents	SD	Uplift through RTO Engine
RTO Engine	Individual	Individual	6.863	2.577	-
Circulars with revenue maximizing discount	Mass market	Mass market	3.112	1.639	+120.5%
Circulars with fixed discount	Mass market	Mass market	2.432	1.376	+182.2%

Table 4.7. Comparison of RTO engine profit with mass market promotions.

Policy	Brand	Discount	Profit in Euro cents	SD	Uplift through RTO Engine
RTO Engine	Individual	Individual	.771	.643	-
Circulars with profit maximizing discount	Mass market	Mass market	.504	.321	+53.0%
Circulars with fixed discount	Mass market	Mass market	.364	.241	+111.8%

Giving coupons to a random subset of all customers produces the same results as giving coupons to all customers, so we can use the redemption rate model fitted on the random data to calculate revenues and profits for the benchmark policies. To make sure that the measured revenue and profit values for each customer are meaningful, we select customers that have at least five observations for random coupons, and between 5 and 50 observations for targeted coupons. The final sample consists of 3,453 customers. After that, we obtain all measures at the customer/coupon level by averaging values within each customer and policy.

Tables 4.6 and 4.7 depict the results for the revenues and profits comparison, averaged across customers. Targeted RTO engine coupons achieve the highest revenues ($r_{ibst}^T = 6.558$ Euro cents) and profits ($\pi_{ibst}^T = .736$ Euro cents). On the other hand, the fixed discount policy (policy 1) leads to the worst results ($r_{ibst}^{P1} = 2.321$ Euro cents and $\pi_{ibst}^{P1} = .348$ Euro cents). For policies 2 and 3, revenues and profits are higher compared to the fixed discount policy, but still lower than for the targeted policy. Interestingly, the fixed discount policy leads to the smallest SD across customers whereas the targeted policy has the largest SD. This indicates that targeting exploits the heterogeneity in customer preferences and that the RTO engine takes into account that customers are not equally profitable.

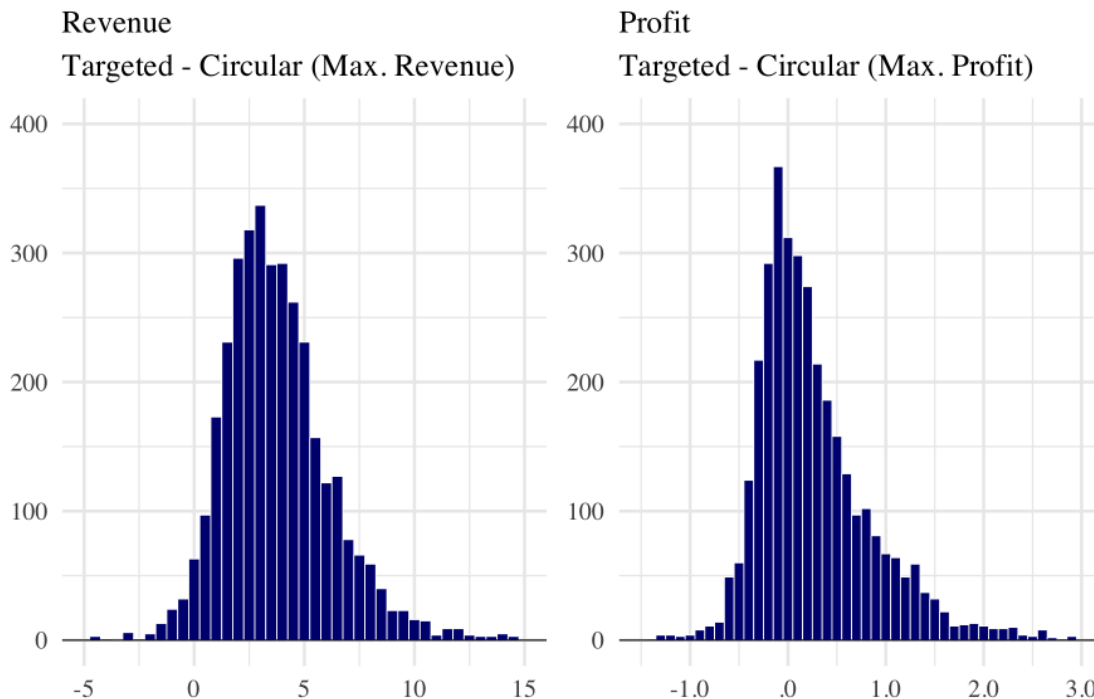
Figure 4.4. Revenue and profit per customer/coupon vs. circular.

Figure 4.4 depicts the revenue uplift through targeting compared to the circular policy with revenue-maximizing discounts ($r_{ibst}^T - r_{ibst}^{P2}$) and the profit uplift compared to the circular policy with profit-maximizing discount ($\pi_{ibst}^T - \pi_{ibst}^{P3}$) across the customers in our data sets. For the revenue differences, most values are centered around the value of about 3 to 6 cents. The distribution is slightly right-skewed with a revenue uplift of 8 to 10 cents for a significant number of customers. We find a decrease in revenues for only a small fraction of customers. The result for profits is very similar although the distribution is more skewed and we find negative differences for a larger fraction of customers. Nonetheless, the overall profit uplift is positive. This result seems to indicate that the RTO engine studied here focuses more on increasing revenues than profits.

Overall, we observe that targeting leads to significantly higher revenues and profits, in addition to higher redemption probabilities (see Section 4.5). Revenues are 120.5% larger than for mass market promotions with a revenue-maximizing price, and profits can be increased by +53.0% in comparison to mass market promotions that use a profit-maximizing price.

4.6.2 Revenue and Profit Uplift for Varying Degrees of Personalization

To understand how the size of the targeted population (in other words, the degree or quality of personalization) affects financial metrics, we extend the analysis of

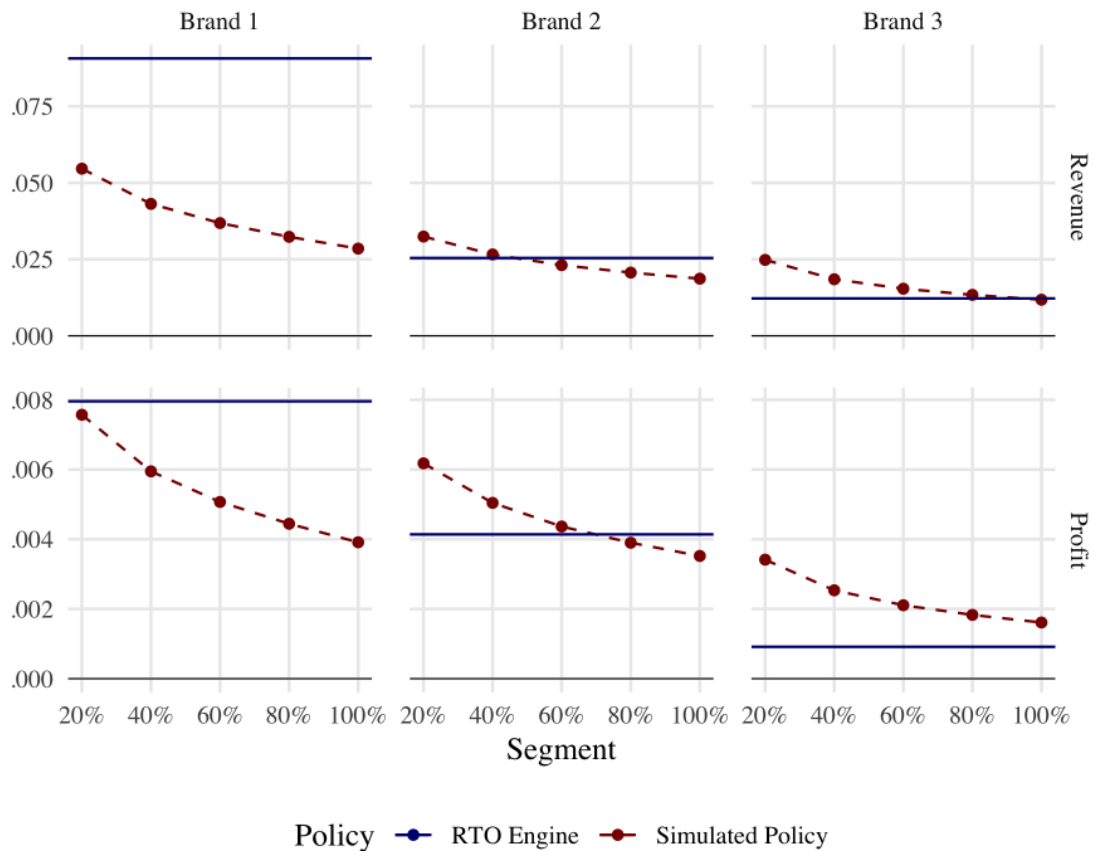
personalization to revenues and profits. Similar to the analysis of redemption probability drivers in Section 4.5.2, this analysis contributes to the generalizability of our findings. Various factors might reduce the degree or quality of personalization (e.g., the type of retailer, the specific RTO engine implementation, the number of available brands for targeting), so the results presented here explicate how sensitive revenues and profits are to external factors.

To this end, we use the redemption model trained on random data and predict redemption probabilities, revenues, and profits for each customer and brand. Discounts are based on policies 2 and 3, that is mass marketing promotions with brand-specific revenue or profit-maximizing discounts. The key difference in this analysis is that we distribute coupons only to a (varying) subset of the customer population. The customers are selected based on the brand score variable used in the RTO engine. Specifically, we vary the degree of brand personalization in this analysis by focusing on customers that are within the top 100%, 80%, 60%, 40%, or 20% quantile for the brand score within each brand. We expect revenues and profits to be higher for customers with relatively high values for the brand score because this should lead to higher redemption probabilities. Furthermore, the top 100% group refers to no personalization and is therefore a logical benchmark in the analysis. As a second benchmark, we compare the results to the values resulting from targeted coupons as derived by the RTO engine.

The analysis provides consistent results across all brands. Revenues and profits increase with the degree of personalization and the average relative uplifts of 13% (top 80% cohort) to 95% (top 20% cohort). Assuming that the pool of potential customers and available brands for targeting is large enough, we see that brand personalization results in significant uplifts.

The three brands in Figure 4.5 are representative for the larger group of brands in our data set, such that the revenue and profit uplifts are very similar. The figure highlights that a higher degree of personalization (i.e., segments containing a smaller number of customers with higher values for the brand score) leads to higher values for revenues (upper panel) and profits (lower panel). However, potential uplifts for revenues and profits can differ across brands and are related to the heterogeneity of the customers' brand preferences. Intuitively, more heterogeneity allows for higher uplifts.

The figure also contains the average expected values of revenues and profits for each brand based on the targeted cohort as horizontal lines. In some cases, revenues and/or profits in the targeted group are above the value for the top 20% cohort (i.e., the segment with the highest degree of personalization). This indicates that the RTO engine leads to better results. On the other hand, in some cases, particularly for profits, the results for the targeted coupons are somewhere between

Figure 4.5. Revenue and profit for varying degrees of personalization.

Note: Best viewed in color.

the top 20% segment and the “no personalization” case (i.e., 100%). This suggests that even though the RTO engine individualizes coupon offers, the personalization could still be improved and the discounts are most likely not solely set to maximize revenues or profits.

Lower revenues and profits for the RTO engine can be explained by external constraints not in the control of the RTO engine (e.g., product availability) or an overdistribution of brand coupons. Targeting too many customers with a given brand leads to lower revenues and profits. Nonetheless, it is important to note that the performance of the RTO engine is comparable to mass marketing policies even in the worst case.

4.7 Conclusion

Although coupons are essential to the retailers’ sales promotion mix, research on promotion personalization through RTO engines in grocery retailing has been limited. We base our study on data collected at a leading German grocery retailer. The data comprise loyalty card transactions, market basket data, and 12 million (brand) coupons for 1,116 brands in 115 categories. For almost 1 million

coupons, the brand and the discount were randomized, so the exogenous variation in both dimensions of targeting (i.e., brand and discount) facilitates an unbiased measurement of the effect of targeting on redemption rates, revenues, and profits.

The results reveal that the targeted brand coupons have (on average) 64.0% higher redemption rates than non-targeted coupons. We observe significant variation across categories and brands, much of which can be explained by brand and category characteristics, such as brand loyalty, price position, and purchase frequency in a second-stage regression model. At the same time, the RTO engine increases the per customer/coupon revenue by up to 182.2% and profit by up to 111.8% compared to mass market price promotions. We further show that the coupon performance is directly linked to the quality of the targeting algorithm (reaching the right customer), such that a smaller degree of distribution leads to significantly higher revenues and profits.

This research offers several pertinent implications for sales promotion management. Most importantly, the effectiveness of targeted coupons is significantly higher than that of non-targeted coupons (e.g., FSI or mass market checkout coupons). The increase in redemption rates due to coupon personalization underlines the value of RTO engines (in addition to efficiency gains that result from using kiosk systems). The analysis also shows that RTO engines offer tangible economic benefits. Targeted coupons increase the expected revenue per coupon and customer by 3.75 Euro cents and the expected profit per coupon and customer by .27 Euro cents. Assuming that customers use the kiosk 40 times per year (and each print contains eight coupons), this translates into a revenue increase of €12 million and an annual profit increase of approximately €1 million per 1 million loyalty card customers.

A better understanding of the mechanics of RTO engines empowers retailers to use these complex target marketing tools appropriately. For promotion management, for example, our analysis of redemption rates reveals substantial differences across brands. Even brands with small redemption rates for non-targeted coupons can be highly relevant, because coupon personalization can lead to high redemption rates for a subset of the total customer population. When it comes to coupons, retailers typically gravitate toward brands with the highest average impact (e.g., in terms of redemption rates). Our results suggest that retailers might benefit from a more customer-centric approach (Shah et al., 2006). Individual-level promotion management is not feasible unless it is automated (Kannan et al., 2017), so this step requires retailers to give up some control by relying on RTO engines.

The structured analysis of redemption rate heterogeneity is important for retailers deciding which brands to promote. In all our analyses, brand heterogeneity explains more variation in coupon effectiveness than does category heterogeneity. This

finding highlights the gains that are possible from managing RTO engines using a brand-based view. It is unrealistic to expect that all products can be promoted in RTO engines, but retailers should exploit the full potential of RTO engines by ensuring that a broad range of brands is available to cater to the diverse preferences of individual users. The value of targeting is particularly evident for highly specialized, unique brands. Moreover, all else being equal, coupons for brands with higher (customer) loyalty achieve higher redemption rates. Because the positive effect of a brand's penetration is stronger for non-targeted coupons, brands with above-average penetration should be included if the potential for personalization is low (e.g., when cold-starting the system). Brands with a broader range in preferences also facilitate personalization, as is intuitive. But if personalization is not feasible, managers should rely on popular brands with less diverse preferences. The RTO engines should primarily include brands promoted less frequently. Yet categories with higher purchase frequencies invoke higher redemption rates, due to the lower perceived risk associated with redeeming a (customized) surprise coupon in those cases. Categories characterized by more competition and strong private labels are also favorable for targeting. Finally, the uplift-effect is lower in impulse categories, presumably because redemption rates are higher in these categories even without targeting. A lack of personalization could be offset by promoting such categories.

It is also important to acknowledge the strategic implications of the relationship between the degree of personalization of the promotion channel (for a particular brand) and the efficiency of RTO engines. Instead of distributing a smaller set of brands to many customers, brand coupons should be strongly differentiated, such that a larger set of brands is distributed to more specific target audiences. The increase in revenues and profits comes along with an increased complexity of executing and analyzing promotions and promotion automation requires retailers to use and trust RTO engines. A smaller customer reach for certain brands also might influence the retailer's negotiations with manufacturers. But regardless of the degree of distribution, the use of RTO engines for coupon personalization produces redemption rates, revenues and profits that are equal to or higher than those of non-targeted coupons.

Our results also emphasize the value of collecting user-specific purchase history data through loyalty programs. Beyond using customized coupons for marketing activities, retailers can function as custom data intermediaries, by leveraging their customer data and RTO engines to offer targeted coupon capabilities to manufacturers that sell through their stores (Pancras and Sudhir, 2007). Personalized coupons are not only a way to increase revenues and profits but also have the potential to be a new revenue stream for retailers in the form of programmatic target marketing platforms (Pathak, 2017; Chen and Friesz-Martin, 2018).

Finally, our research points to the value of kiosk systems for in-store couponing. The measured response to price discounts through kiosk coupons (average price elasticity = -2.9) is similar to responses to classical price promotions, so the results presented here validate this promotion channel and emphasize its value for effective price discrimination.

The breadth of the brand and category dimensions of the data support the external validity of our results. At the same time, we note some promising opportunities for further research. It would be interesting to substantiate the generalizability of our findings further and analyze similar RTO engines across different retail settings (e.g., supermarkets vs. discount stores) or different engines at the same retailer (e.g., using different algorithms). A comparison of kiosk coupons with mobile coupons might reveal potential differences in the effects (e.g., redemption rates, sales, search) that arise from the distinct distribution mechanisms. A natural extension of the study of financial implications of RTO engines is the analysis of long-term and cross-category effects of RTO engines. A decomposition of the effects of (targeted) coupon redemption versus exposure (Venkatesan and Farris, 2012) represents another promising avenue for research. To inform retail strategy, it would be interesting to investigate the implications of RTO engines for loyalty programs, analyze the interaction of customized coupons with classical promotion instruments (e.g., displays and features), or address the effects of RTO engines on the retailer's image.

4.8 Appendix

4.8.1 Model Specification

Table 4.8 contains the estimation results for five nested model specifications for the random data. M5 is the full model as specified by Equations 4.1 and 4.2; M1 to M4 are simpler models in which we have systematically omitted specific terms. All model coefficients are significant with $p < .01$. The implied average price elasticity for M5 is -2.96 . The elasticities for the other model specifications are very similar, with the exception of model M1. Only using a global intercept does not model redemption rate heterogeneity across brands adequately and price effects are biased toward 0. The random effects (i.e., brand, year week, and store) are significant across all model specifications and improve the log-likelihood and the Akaike information criterion (AIC).

The biggest improvement is achieved by controlling for brand heterogeneity which is in line with our argument that brands are the most relevant source of heterogeneity in the results. Model coefficients are not significantly different across model specifications M2 to M4, which underlines the robustness of the estimated models. The most flexible model, model M5, has the best log-likelihood value, so we use this specification as the basis for the predictions and policy simulations. On a side note, for random coupons, the R^2 of M5 for log odds on the brand-level is .870, and the correlation is .967. Therefore, we believe that M5 is also well suited for our analyses in Sections 4.5 and 4.6.

In Table 4.9 we show the estimation results for the same five nested model specifications estimated on the targeted data. The results are quite similar to the results discussed above. In particular, not accounting for heterogeneity leads to a downward bias in the discount effects. Also for the targeted data, models with random effects fit the data significantly better, and the full model (M5) outperforms all other models. R^2 and correlations values between fitted redemption probabilities and redemption rates are also excellent on the brand-level (.932 and .967, respectively). Hence, we use model M5 for the (in-sample) predictions in our analyses.

Table 4.8. Model results based on random data.

Variables	Random Model				
	GLM	GLMM1	GLMM2	GLMM3	GLMM4
Intercept α_0	-5.086	-5.465	-5.576	-5.467	-5.579
Discount β_0	3.070	3.903	3.925	3.906	3.928
Brand score γ	.503	.509	.525	.507	.523
SD(Brand) α_j		.902	.897	.902	.896
SD(Discount) β_j		1.175	1.181	1.176	1.183
Cor (α_j, β_j) ρ		-.786	-.784	-.786	-.784
SD(Year-Week) α_t			.222		.222
SD(Store) α_s				.132	.133
LL	-56,960	-55,517	-55,410	-55,476	-55,370
AIC	113,927	111,046	110,834	110,967	110,757
N			750,876		

Note: All coefficients are significant with $p < 0.01$.

Table 4.9. Model results based on targeted data.

Variables	Targeted Model				
	GLM	GLMM1	GLMM2	GLMM3	GLMM4
Intercept α_0	-3.354	-3.962	-4.111	-3.965	-4.116
Discount β_0	.478	1.289	1.423	1.297	1.430
Brand score γ	.448	.485	.492	.484	.492
SD(Brand) α_j		.921	.869	.918	.865
SD(Discount) β_j		1.740	1.618	1.735	1.612
Cor (α_j, β_j) ρ		-.718	-.785	-.715	-.782
SD(Year-Week) α_t			.301		.301
SD(Store) α_s				.095	.096
LL	-128,402	-123,718	-123,334	-123,648	-123,265
AIC	256,810	247,449	246,683	247,311	246,546
N			750,876		

Note: All coefficients are significant with $p < 0.01$.

4.8.2 Variables for Redemption Rate Analysis

Table 4.10 lists the variables for our second stage analysis on the brand-level. Many of the variables from prior studies (e.g., Bell et al., 1999) are similar and highly correlated, so we use the 13 variables that are relevant in our context but that do not create empirical issues for the model estimation (see also the discussion regarding the correlations between the variables below).

Table 4.10. References and descriptive statistics for explanatory variables.

Variable	Reference	Mean	SD
Loyalty ¹⁾	Bell et al. (1999)	2.423	3.233
Penetration	-	.115	.140
RTO engine brand score range	-	.220	.075
Price position	Bell et al. (1999)	1.268	.606
Deal depth ¹⁾	Bell et al. (1999)	.993	.216
Promotion frequency	Osuna et al. (2016)	.099	.096
Purchase frequency	Bell et al. (1999)	.141	.126
Private label share	Narasimhan et al. (1996)	.387	.175
Competition	Osuna et al. (2016)	.218	.110
Price dispersion	Osuna et al. (2016)	7.755	9.703
Price	Narasimhan et al. (1996)	2.080	1.049
Stockpile score	Bell et al. (1999)	-.143	.802
Impulse score	Narasimhan et al. (1996)	.295	1.060

Note: 1) Normalized in category.

For loyalty, price position, purchase frequency, private label share, competition, price dispersion, and price, we rely on the household panel data available from a German panel data provider. By using household panel data, we can measure these variables with great precision and enhance the generalizability of our results. We limit these data to panelists who were observed at least once every two weeks. For the variables penetration, RTO engine brand score range, deal depth, and promotion frequency, we turn to the retailer's sales and loyalty card data. Both stockpile score and impulse scores are measured on scales from Narasimhan et al. (1996). We tested several operationalizations including discretized versions, and the results remained consistent. Dummy coding based on median splits of the raw scores yielded the best results. We log-transform right-skewed variables (competition, loyalty, penetration, RTO engine brand score range, price position, promotion frequency, purchase frequency, price dispersion, and price) to reduce the effect of extreme values.

Table 4.11. Correlation between brand and category characteristics.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Loyalty (1)												
Penetration (2)	.26											
Brand score range (3)	.17	.00										
Price position (4)	-.27	-.16	-.09									
Deal depth (5)	-.13	-.20	-.16	.02								
Promotion frequency (6)	.06	.08	-.20	.07	-.05							
Purchase frequency (7)	.34	.47	.24	-.09	-.08	-.26						
Private label share (8)	-.24	-.06	-.04	.13	.08	-.20	-.07					
Competition (9)	.07	.04	.00	.18	-.03	-.06	.12	.38				
Price dispersion (10)	.03	-.47	.18	.05	.07	.01	-.34	-.06	-.04			
Price (11)	-.07	-.32	-.10	-.13	-.02	.18	-.31	-.19	-.24	.14		
Stockpiling score (12)	-.20	-.18	-.23	-.05	.04	.17	-.38	.28	-.04	.00	.23	
Impulse score (13)	-.11	.24	-.09	-.09	-.01	.15	.13	-.11	-.25	-.47	.02	.15

Table 4.11 shows the lower-triangle of the correlation matrix of the transformed variables (i.e., how the data is used in the regression analysis). All correlations are in $[-.5, .5]$, and most correlations are $< .1$ (in absolute terms). Therefore we conclude that multicollinearity is not an issue in our data set. Adding more variables from prior studies to the model is not reasonable because these would lead to higher correlations ($> .7$ in absolute terms). Because several variables are quite similar and measure closely related constructs (e.g., purchase frequency and interpurchase time), our selection is complete in that we cover critical dimensions for explaining redemption rate heterogeneity over campaigns.

Lastly, the data for the impulse score and the stockpiling score for the product categories used in Section 4.5.2 were collected in July 2017 on MTurk. We followed recent guidelines and recommendations for designing MTurk surveys (Goodman and Paolacci, 2017). Each of the 614 respondents rated eight (randomly chosen) categories. By keeping the effort low, we minimized wearout effects during the survey. We only used a respondent's rating if he or she purchased the category at least once in the last six months. Therefore, the number of observations differs across categories (min = 7, median = 41, max = 58).

Table 4.12 summarizes the descriptive statistics for the four items from Narasimhan et al. (1996). Items 1 and 3 should be related to stockpiling, while items 2 and 4 are hypothesized to be related to impulse. Based on a principal components analysis, this pattern is also apparent from the structure of the resulting loadings using varimax rotation (Table 4.12). The two-component solution (component 1 =

Table 4.12. Items for impulse and ability-to-stockpile scales.

	Descriptive Statistics		Loadings	
	Mean	SD	Comp 1	Comp 2
It is easy to store extra quantities of this product in my home.	5.149	1.854	-.106	.953
I often buy this product on a whim when I pass by it in the store.	3.621	2.026	.967	-.140
I like to stock up on this product when I can.	4.408	2.060	-.340	.886
I typically like to buy this product when the urge strikes me.	3.986	2.056	.940	-.269

Notes: Statistics measured on a 7-point agree/disagree-scale. Loadings derived through PCA.

Table 4.13. Ten highest and lowest ranked categories for the impulse and stockpiling scores.

Impulse		Stockpiling	
High	Low	High	Low
Candy	Dog Food	Pasta	Fish
Salty Snacks	Cat Food	Coffee (Single Pack)	Fresh Bakery Products
Chocolate	Dishwashing	Coffee	Cream
Frozen Pizza	Detergent, Oil/Vinegar	Canned Fish	Milk Drinks
Cookies	Bags/Wraps/Disposable Containers	Soap	Convenience Salads
Cocoa	Dish Care	Salty Snacks	Milk
Tea	Shaving Needs	Herbs	Cake
Ice cream	Eggs	Body Care	Frozen Desserts
Desserts	Female Care	Coffee Filters	Eggs
Cereal bars	Fabric Softener	Female Care	Cream Cheese

impulse, component 2 = stockpiling) explains 93.3% of the variance of the items. We standardize both variables to simplify their interpretation. The face validity of the results is supported by taking a closer look at the ten highest and lowest ranked categories for each variable (Table 4.13). High impulse-buying categories include candy and ice cream, whereas low impulse-buying categories include pet food and fabric softener. Categories with a high ability-to-stockpile are pasta and coffee. Fresh products such as fish and convenience salads are among the categories

with a low value for the stockpiling score. These results are intuitive and very similar to the findings of Narasimhan et al. (1996).

To ensure that the MTurk results apply to data collected in Germany, we let ten German retailing experts rate 20 categories on the four items of Narasimhan et al. (1996). We calculated simple sum scores for each construct and correlated these scores with the results from the MTurk sample for the corresponding categories. We obtained positive and highly significant correlations for both variables: .860 for the impulse score ($t = 7.15$, $df = 18$, $p < .01$) and .788 for the stockpiling score ($t = 5.44$, $df = 18$, $p < .01$). These results establish the usefulness of the MTurk sample and the validity of the measured scales.

As a final way to validate the results, we used the same MTurk survey to collect data for the “should-minus-want” score of Milkman et al. (2010). Being able to replicate existing research findings adds to the validity of surveys (Laurent, 2013). We found a highly significant positive correlation of .827 ($t = 10.19$, $df = 48$, $p < .01$) with their results across the 50 categories that we were able to match.

5 | The Impact of Personalized Coupons on Loyalty Program Usage

Sebastian Gabel, Denis Vuckovac, Daniel Guhl

Abstract

Many customers are members in loyalty programs (LP) and the number of LP memberships is increasing steadily. At the same time, LP usage is at an all-time low. LPs send too many communications, it takes too long to earn points for rewards, and LPs do not provide relevant rewards. Retailers respond to these challenges by designing new ways to interact with and reward customers. One example is the LP of a leading German grocery retailer studied here. At an in-store kiosk system close to the entrance of each store, customers can check their loyalty point balance, redeem loyalty points for free grocery products of their choice and receive personalized coupons. This setup is very suitable for studying the impact of rewards on LP usage. A rich longitudinal data set that contains data from more than 7,000 customers over a period of 60 weeks makes it possible to (1) analyze how personalized coupons affect LP usage, (2) compare the effect of personalized coupons to that of classic LP rewards, (3) study differences in effectiveness across customer segments, and (4) derive pertinent implications for reward design and ways to increase LP usage. We conduct an additional Mechanical Turk experiment to support our findings and conclusions. Prior research has studied the revenue and profit implications of LPs extensively, yet little is known about how the LP design and LP rewards influence LP usage. Our study is a first step in this direction. At the same time, we provide insights into the interaction between two of the most fundamental aspects of retail management: LP design and price promotions.

Keywords

loyalty programs, rewards, one-to-one marketing, personalized coupons, duration analysis

5.1 Introduction

As of 2016, 3.8 billion individual loyalty program (LP) memberships exist in the US and this number has nearly tripled within the last ten years (Fruend, 2017). 74% of all customers who belong to an LP use at least one program from a grocery store (Collins, 2017), so (grocery) retailing remains one of the most relevant industries for LP research. For customers, LPs are an important way to save money, both through discounts and coupons (Collins, 2017). Retailers continue to trust in the strategic importance of LPs. Although the direct effects of LPs on the customers' purchase behavior and the retailer's revenue remain debated (Zhang and Breugelmans, 2012), customer-centric firms understand that personalized marketing leads to substantially higher profits (Rust and Verhoef, 2005; Fader, 2012). Retailers collect vast amounts of customer-level data that they use to analyze customer purchasing habits, so LPs help firms to improve their position in target segments (Arora et al., 2008; Blattberg et al., 2008; Bradlow et al., 2017; Palmatier and Sridhar, 2017). Target and Safeway's, for example, personalize circulars and coupons according to customers' individual shopping histories (Bleier et al., 2018).

Nonetheless, LPs face significant challenges, especially in the offline world. In fact, more than half of all LP memberships in the United States are inactive (Fruend, 2017). The 2017 Colloquy loyalty census reveals three main reasons why a growing number of customers stop using LPs: (1) LPs send too many communications, (2) it takes too long to earn points for rewards, and (3) the LPs do not provide relevant rewards and offers (Fruend, 2017). It comes as no surprise that retailers are looking for new LP designs that mitigate these problems. We study one such example, an LP that a leading German grocery retailer introduced in 2015. An in-store kiosk system at the entrance of each store allows the retailer to communicate with customers. Customers can check their loyalty point balance (points are collected proportionally to their spending) and redeem loyalty points for free grocery products of their choice. In addition to these LP rewards, the retailer uses a real-time offer (RTO) engine to personalize coupons that are distributed through the kiosks. These components of the LP directly address the challenges mentioned above: Instead of overwhelming the customer with information by "pushing" content to the customer, the customer can control the information flow by deciding when to "pull" information (Marketing Science Institute, 2016). Loyalty points are redeemed for grocery products so the reward turn-around is quick. Additionally, offers are highly relevant because the RTO engine personalizes coupons and customers can choose the LP rewards.

A rich, longitudinal data set that comprises purchase histories and LP transactions is the basis for modeling LP usage—defined as the time between two consecutive kiosk usage events (inter-usage time, or IUT)—as a function of (past)

personalized promotions and LP rewards in a continuous-time latent class proportional hazard model (PHM). We (1) evaluate how personalized promotions affect LP usage, (2) compare the effect of personalized promotions to that of classic LP rewards, (3) study differences in effectiveness across customer segments, and (4) derive pertinent implications for reward design.

With these goals, our work is relevant to both researchers and practitioners. Prior research has studied the revenue and profit implications of LPs extensively, yet little is known about how the LP design and LP rewards influence LP usage. “Because underuse of LPs by consumers has a detrimental effect on firm performance, practitioners [...] have called for academic insights on measuring membership participation” (Breugelmans et al., 2015, p. 132). The rich cross-sectional and longitudinal data set collected in the LP mentioned above is well-suited for this kind of analysis. While it might seem that the LP studied here has uncommon if not unique features compared to LPs previously studied in the literature, in-store kiosks are increasingly popular among practitioners. In fact, a growing number of retailers (e.g., CVS, Ahold) use kiosk systems and two of the four largest retailers in Germany have introduced kiosk systems in the last two years. Yet, research on such systems is scarce (Grewal et al., 2011; Osuna et al., 2016; Inman and Nikolova, 2017) and knowledge about usage behavior remains locked within retailers and solution providers. New retail technologies provide exciting opportunities for LP research (Breugelmans et al., 2015). The results from this study support the general positive perception of kiosk systems and can guide retailers in the design, development, and roll-out of LPs. With “pull” marketing gaining more importance (Marketing Science Institute, 2016), it is essential to understand how customers respond to kiosk systems and usage becomes a relevant proximal behavioral outcome that retailers need to actively manage. At the same time, our research directly combines two topics that are fundamental to retail management: LP design and price promotions. We study how LPs can be integrated with price promotions and what the effect of such combinations is (see LP research agenda in Bijmolt and Verhoef, 2017).

For practitioners, our research provides pertinent insights in that we show to increase LP usage. In doing so, we point to a new use case for price promotions. Promotions typically function as an instrument to increase store traffic, sales, and profits (Venkatesan and Farris, 2012). We find that personalized coupons boost LP usage, even more so than classic LP rewards. This insight is relevant because personalized coupons are cheaper than free products in most settings. Leveraging LPs as a platform for personalized, programmatic promotions also opens up additional revenue streams for retailers (Pathak, 2017; Chen and Friesz-Martin, 2018).

The remainder of this paper is structured as follows: We first review related literature on LPs and highlight our research contribution. Next, we introduce the data used in this study. After discussing our empirical analysis and main results, we conclude by detailing the key findings, managerial implications, and opportunities for further research.

5.2 Related Work and Contribution

In the discussion of related work and our relative contribution, we focus on three main streams of literature: customer responses to LPs, LP rewards, and the effects of personalized coupons.

5.2.1 Customer Responses to LPs

LPs aim to strengthen the long-term relationship between customers and firms by providing program rewards or access to exclusive services in return for repeat purchases (Berry, 1995; Liu, 2007; Bijmolt et al., 2011). In frequency reward programs three mechanisms increase customer value (Blattberg et al., 2008; Taylor and Neslin, 2005): (1) The feeling of being close to obtaining a reward (points-pressure mechanism) increases the likelihood of additional purchases. (2) The act of rewarding (rewarded-behavior mechanism) reinforces customer attachment to the firm. (3) The exploitation of personalized data obtained by means of the LP for marketing (personalized marketing mechanism) triggers desired customer responses.

Bijmolt et al. (2011) further differentiate between two types of customer responses to LPs, namely attitudinal and behavioral loyalty. The authors argue that commitment and satisfaction are fundamental to behavioral loyalty and point out that a firm can increase attitudinal loyalty by enhancing LP and reward attractiveness (Demoulin and Zidda, 2009; Keh and Lee, 2006). Despite divergent findings, the majority of studies demonstrate a positive effect of LPs on behavioral responses such as purchase frequency, sales, share of wallet, and customer retention (Liu and Yang, 2009; Leenheer et al., 2007; Minnema et al., 2017; Verhoef, 2003; Meyer-Waarden, 2007).

Most studies focus on the effects of LPs on the firm's financial performance, that is sales and profits. While it remains the ultimate goal of LPs to impact financial metrics, retailers face the problem that many customers use LPs irregularly and finally churn. As pointed out in the introduction, more than half of all LP memberships in the United States are inactive (Fruend, 2017). Given the importance of LP usage for firm performance and many industry reports on decreasing LP usage, research is needed to address usage and participation (Breugelmans et al., 2015). In this study, the LP setup and the available data set make the analysis

of LP usage possible: To receive LP rewards, customers have to engage with the LP by interacting with the kiosk. Customers need to decide whether to inform themselves about rewards, so the time delta between two prints directly measure LP usage. This differentiates our study setup from other settings in which system usage is passive and in fact unobserved, for example, when customers receive emails or snail mail. In such systems researchers directly depend on observing a click or redemption, so usage depends on the attractiveness of the offer. The advantage of the data set used in this study is that LP usage, measured as the time between two consecutive kiosk usage events, is independent of the current interaction with the customer because only prior interactions with the system influence usage. The longitudinal nature of our data set facilitates a clean and unbiased measurement of LP usage and the moderating impact of rewards on LP usage, eliminating additional confounding effects. Insights, therefore, can inform the design of LPs and LP rewards.

5.2.2 LP Rewards

Research on LP design has shown that the interplay of multiple design components such as the overall program setup and the reward structure directly affect LP effectiveness. A growing body of literature has thus taken a reward-centered focus. Table 5.1 provides an overview of relevant studies that analyze LP rewards.

Prior studies have focused on different dependent variables, including behavioral responses analyzed in observational studies such as purchase quantities or frequencies (Lewis, 2004; Taylor and Neslin, 2005; Zhang and Breugelmans, 2012), or questionnaire measured loyalty constructs and explicit reward preferences (Kivetz and Simonson, 2002; Keh and Lee, 2006; Meyer-Waarden, 2015). While research has demonstrated the positive impact of novel loyalty schemes on purchase size and frequency (Zhang and Breugelmans, 2012; Minnema et al., 2017; Bijmolt and Verhoef, 2017), empirical studies have examined only a limited number of effects of reward mechanisms and design components of LPs (Bijmolt and Verhoef, 2017).

Studies suggest that in low engagement industries LP members prefer immediate rewards (Yi and Jeon, 2003; Meyer-Waarden, 2015) and that in line with this, necessity rewards are preferred when the effort to achieve rewards is lower (Kivetz and Simonson, 2002). While immediate direct rewards (rewards that are linked to the product/service, e.g., preferential price discounts for LP members) in grocery retail can increase the satisfaction of LP members (Söderlund and Colliander, 2015), only a handful of studies deal with such rewards, and no study evaluates direct rewards empirically. This is surprising as the effectiveness of direct rewards in low-involvement industries (e.g., groceries) is well established and customers typically prefer direct rewards (Yi and Jeon, 2003; Keh and Lee, 2006; Meyer-Waarden, 2015). And while some studies have highlighted the importance of accounting for

Table 5.1. Prior studies on LP rewards.

Author (Year)	Study			Reward		
	Industry	Format	Dependent Variable	Timing	Type	Personalized
Kivetz and Simonson (2002)	car rental, airline, hospitality	experiment	reward preferences	Del.	Ind.	No
Yi and Jeon (2003)	beauty, restaurants	experiment	perceived value	Imm./Del.	Dir./Ind.	No
Lewis (2004)	(online) grocery	empirical	purchases, lifetime	Del.	Ind.	No
Taylor and Neslin (2005)	retail	field study	sales	Del.	Dir.	No
Keh and Lee (2006)	bank, restaurant	experiment	loyalty	Imm./Del.	Dir./Ind.	No
Kivetz et al. (2006)	convenience	field study	purchase observations	Del.	Dir.	No
Zhang and Breugelmans (2012)	(online) grocery	empirical	incidence, spending, LP membership	Del.	Ind.	No
Dorotic et al. (2014)	retail	empirical	balance of points	Del.	Ind.	Yes
Meyer-Waarden (2015)	grocery, perfumery	survey	(store) loyalty	Imm./Del.	Dir./Ind.	No
Söderlund and Colliander (2015)	retail	experiment	satisfaction, repatronage intention	Imm.	Dir.	No
Minnema et al. (2017)	grocery	empirical	incidence (trip/category), quantity	Imm.	Ind.	No
Breugelmans and Liu-Thompkins (2017)	convenience	empirical/experiment	incidence, spending	Del.	Dir.	No
This study	grocery	empirical	Inter-usage time	Imm./Del.	Dir.	Yes

Notes: Timing: Del. = delayed and Imm. = immediate. Type: Dir. = direct and Ind. = indirect.

customer heterogeneity in the analysis of responses to LPs (Lewis, 2004; Taylor and Neslin, 2005), little research addresses the effects of reward personalization (e.g., Dorotic et al., 2014), and no study evaluates delayed and immediate direct rewards simultaneously.


We contribute to prior research on LP rewards through our empirical analysis of two reward mechanisms: (1) Direct and delayed rewards in the form of free products in exchange for loyalty points and (2) direct and immediate rewards in

the form of personalized coupons. We show how the two reward types impact LP usage and how their effectiveness varies over customer segments. The studied LP setup and the available data set facilitate a clean effect measurement. The direct comparison of the reward mechanisms allows us to compare the costs of the studied LP rewards, thereby providing relevant insights for researchers and practitioners when it comes to designing reward mechanisms for LPs. An additional online experiment on Amazon Mechanical Turk (MTurk) provides context for our findings. To our knowledge, this is the first empirical study to assess (and compare) immediate and delayed direct, personalized rewards and their impact on LP usage.

5.2.3 Effects of Personalized Coupons

Studies on personalized coupons typically focus on the effect of coupons on revenue and sales. Rossi et al. (1996) show how to personalize price discounts in a brand choice setting, highlighting that household purchase histories are valuable to manufacturers for optimizing coupon profitability. Zhang and Wedel (2009) study personalized checkout coupons in online and offline stores across two product categories. They find that promotion optimization leads to substantial profit improvements and that loyalty (competitive) coupons are more effective in online (offline) stores. The authors point to low redemption rates of checkout coupons as a “major impediment to the success of customized promotions” in offline stores (Zhang and Wedel, 2009, p. 204). Heilman et al. (2002) examine the impact of in-store “surprise” coupons on total basket value. In addition to positive effects on spending, they find that in-store coupons are well received by customers and prompt up to ten times higher redemption rates than coupons from freestanding inserts. Nevertheless, home-sent coupons are still the predominant means of providing personalized coupons (Bijmolt and Verhoef, 2017). Venkatesan and Farris (2012) present a conceptual framework for personalized email coupon campaigns. In a quasi-experiment, they find that coupon exposure and redemption have positive effects on trip incidence and revenues.

We add to prior research on personalized coupons by outlining one potential application of personalized promotions as (exclusive) rewards in the context of LPs. Although researchers have demonstrated that personalization can act as a loyalty-building mechanism (Bijmolt et al., 2011; Meyer-Waarden, 2007; Verhoef, 2003), such opportunities remain underutilized in practice and understudied in academia (Bijmolt and Verhoef, 2017). We show how LPs can “be combined or even integrated with other marketing-mix instruments” (Bijmolt and Verhoef, 2017, p. 161). The effectiveness of in-store promotions and the “pull”-based nature of the customer interaction through kiosk systems are key to high customer responsiveness (Grewal et al., 2011; Marketing Science Institute, 2016).

Figure 5.1. LP usage cycle.

5.3 Loyalty Program Setup and Data

To evaluate how personalized promotions affect LP usage and to compare the effect of personalized promotions to that of classic LP rewards, we conduct an empirical study using data from a large German brick-and-mortar grocery retailer. In the twelve months before the study, the retailer’s average market share was 5.6% and 51.8% of all customers who had one of the retailer’s stores in their neighborhood visited the retailer at least once. Despite its excellent reach, customers spent most of their money at competitors. In 2015, the retailer therefore introduced an LP that combined classic LP rewards with personalized price promotions.

Barcodes on the back of credit card–sized loyalty cards function as user IDs. Privacy concerns are a known issue in LPs (Bijmolt and Verhoef, 2017), so the retailer decided to avoid a formal registration that requires personal information of customers. Personalized coupons are based on individual-level data, such as purchase histories, loyalty points, usage and redemption behavior, and distributed through in-store kiosks that are similar to CVS’s ExtraCare Coupon Center. Customers use the kiosk when entering the store to receive (up to eight) personalized coupons. The coupons are valid only on that same day (i.e., a customer receives a new set of personalized coupons every day) and redeemed automatically if a customer purchases the promoted products and scans the loyalty card during checkout. When customers scan their cards at checkout, they collect loyalty points proportional to the value of the goods they have purchased. Customers can check their loyalty point balance and select rewards (free grocery products) at the kiosk system. Figure 5.1 summarizes a customer’s complete usage cycle. Personalized coupons are selected according to customer preferences, using an RTO engine that leverages market basket data, LP transaction data, data about past coupons, and reward redemptions, all linked through the customer’s loyalty card. The personalization algorithm is regularized logistic regression with batch


updates (Chapelle et al., 2015). Hand-made expert features are complemented by deep learning features that capture product relationships within and across categories (Gabel et al., 2019). At any given time, 25% to 50% of all category-brand combinations in the assortment were used for personalized promotions.

The LP design is well suited for studying and comparing the impact of classic LP rewards and personalized coupons on LP usage for three reasons. First, to receive rewards, customers must interact with the kiosk, so the time between two prints is a direct measure of LP usage. This setup differs from research settings in which the customers' passive system usage is actually unobserved, for example, when customers receive coupons via email. In these settings, researchers depend on observing a click or redemption and when offers are not redeemed, it remains unknown whether they have been noticed at all. In our setting, we can observe LP usage independent of redemptions. Also, a click or redemption directly depends on the attractiveness of the offer, so this confounding effect makes it challenging to separate the effects of past and current rewards on LP usage.

Second, we measure the effects of personalized promotions and free product rewards on LP usage simultaneously and with a high frequency. LP rewards are grocery products with a relatively low monetary value, so the number and frequency of reward redemptions is higher than in other LP settings (the frequency of coupon redemptions is even higher).

Third, the data collected are well suited for assessing the relationships between the dependent variable (usage) and explanatory variables (rewards). As we illustrate in Figure 5.2, we obtain longitudinal and cross-sectional data about system usage and reward redemptions, so we measure system usage and quantify the impact of rewards on usage while controlling for customer heterogeneity and external influences. The customer cross-section further makes it possible to analyze how the effectiveness of rewards varies across customers.

The data that is the basis for our empirical analysis tracks 15,103 customers who joined the system in the 60 weeks between September 2015 and October 2016. The retailer did not carry out any marketing campaigns during this time window to promote the loyalty program and its usage. To calculate the variables that we use to explain heterogeneity in the latent class model, we specify the first eight weeks after a customer joins the LP as an initialization window. The time after this initialization window is used to model system usage. We consider customers with at least three purchases and two printouts in the initialization window and six printouts in the model window. This is necessary for calculating the variables used in the model and ensures sufficient precision. Further, we only study customers who remained in the system for at least 200 days, so we ensure that we observe system and reward usage sufficiently well. The final sample thus consists of 7,373

Figure 5.2. Longitudinal and cross-sectional data on print and reward events.

Note: Best viewed in color.


customers, whose average basket size is 9.3 items and average inter-purchase time is 7.3 days. The small basket size and low inter-purchase time are representative of a German retailer in a metropolitan area.

5.4 Empirical Analysis

The focus of the empirical analysis is to study the effect of redemptions of personalized coupons and LP rewards on future system usage, measured as the time between two subsequent print events. We use duration models to assess how rewards increase the likelihood of future prints.

5.4.1 Descriptive Analysis

The IUT in days is the dependent variable in our analysis. The data contains 95,285 (uncensored) observations from 7,373 customers (on average, 13 IUTs per

Figure 5.3. Median customer inter-usage time.

customer). IUTs range from 1 to 244 days with an average value of 14 (median 7). As Figure 5.2 shows, many customers redeem coupons and free products, but there is considerable heterogeneity across customers and time. On average, customers saved €9.2 by redeeming 13.1 personalized coupons and €2.6 by receiving 1.0 free products per year. Figure 5.3 depicts the median IUT for each customer in the model time window (i.e., excluding initialization). Most IUTs are smaller than 10 days, but many customers have a median IUT of several weeks, indicating considerable heterogeneity in print behavior. We expect that rewards increase the likelihood of a next print event and thus, *ceteris paribus*, should decrease the IUT.

To test this hypothesis, we first regress $\log(IUT)$ on dummies that indicate a reward redemption (coupons or free product) during the previous print occasion in a linear model (Table 5.2). To control for unobserved heterogeneity in both dimensions in our panel data we use random effects for customers and dates. The results show (significant) negative effects for both lagged reward dummies even after controlling for unobserved heterogeneity, with meaningful effect sizes: Coupons reduce the IUT by 7.3%, and free product rewards decrease it by 4.7%. The preliminary descriptive analysis reveals the expected reward effects, but it does not account for heterogeneity in reward effects or intertemporal dynamics in print behavior (beyond rewards effects). Most importantly, the regression approach does not allow for a changing probability of LP usage over time, given the time elapsed since the last usage event. To overcome these limitations, we extend this first analysis by modeling durations with a continuous-time latent class PHM.

Table 5.2. Effect of lagged rewards on $\log(IUT)$.

	Linear Model		Linear Mixed Model	
	Est.	SE	Est.	SE
Intercept	2.012 ***	.005	2.316 ***	.011
Lagged coupon	-.036 ***	.007	-.073 ***	.018
Lagged free product	-.093 ***	.019	-.047 ***	.007
SD(customers)			.560 ***	
SD(dates)			.125 ***	
SD(residuals)	1.085		.948	
Log-Likelihood (LL)	-143,005.9		-135,999.3	
N	95,285		95,285	

Notes: The dependent variable is $\log(IUT)$. The lagged reward dummies indicate whether a particular reward was redeemed on the last print occasion. Sig. label: *** $p < .01$.

5.4.2 Duration Analysis

Hazard rate modeling. Hazard models are a popular choice in marketing for duration analysis (e.g., Allenby et al., 1999; Venkatesan and Kumar, 2004; Manchanda et al., 2006). Extant studies offer in-depth details regarding the model specification and estimation (e.g., Seetharaman and Chintagunta, 2003), so we only summarize the main steps here. Further details regarding the model implementation, including estimation details and the derivation of the gradient for the latent class PHM, can be found in Appendices 5.6 and 5.6. Our dependent variable is the IUT. The hazard rate is customer i 's instantaneous probability of printing, conditional on the time t (in days) since the last print. The hazard rate is modeled using two main components, a baseline hazard $h_0(t)$ and the effect of covariates x_i that shift this hazard proportionally:

$$h(t, x_i) = h_0(t) \exp(x_i \beta). \quad (5.1)$$

In line with the shape of the IUT histogram (Figure 5.3), we opted for a log-logistic baseline hazard,

$$h_0(t) = \frac{\gamma \alpha (\gamma t)^{\alpha-1}}{1 + (\gamma t)^\alpha}, \quad (5.2)$$

with shape parameters α , $\gamma > 0$. This functional form allows for decreasing and inverted U-shaped hazards and has worked well in other applications of hazard models to purchase timing data (e.g., Seetharaman and Chintagunta, 2003). In our

application, a log-logistic baseline hazard also provides a better fit than alternative specifications such as Weibull or Erlang-2.

Each customer has J_i duration intervals (*spells*) for print events j , so the panel structure of the data set makes it possible to account for unobserved heterogeneity in all parameters with a latent class approach, thereby avoiding biased parameter estimates (Wedel and Kamakura, 2012). We estimate the locations and masses of the multivariate discrete distribution, flexibly from the data, without imposing functional assumptions about the distribution of the parameters. For the prior probability that customer i belongs to class c , we use a multinomial logit model

$$\lambda_{ic} = \frac{\exp(z_i \theta_c)}{\sum_{c'=1}^C \exp(z_i \theta_{c'})}. \quad (5.3)$$

The vector z_i contains customer-specific covariates (“concomitant variables”, see Gupta and Chintagunta, 1994) to determine customer class membership (for details, see Appendix 5.6).

We specify the proportional part of the hazard function and the logit model for the class membership as functions of covariates x and concomitant variables z . Table 5.3 contains the variable operationalization. Their selection and definition follows the literature (e.g., Allenby et al., 1999; Venkatesan and Kumar, 2004) and allows us to control for heterogeneity and structural differences between customers, but also for dynamic effects across spells. The focal variables, $Stock_{ij}^{Coupon}$ and $Stock_{ij}^{FP}$, capture the value of past rewards. We build reward stocks using exponential smoothing, such that

$$Stock_{ij}^{Coupon} = \chi Stock_{ij-1}^{Coupon} + (1 - \chi) \log(1 + Coupon_{ij-1}) \quad (5.4)$$

and

$$Stock_{ij}^{FP} = \phi Stock_{ij-1}^{FP} + (1 - \phi) \log(1 + FreeProduct_{ij-1}), \quad (5.5)$$

where χ and ϕ are smoothing constants, and $Coupon_{ij-1}$ and $FreeProduct_{ij-1}$ are lagged monetary reward values. We initialize both stocks at zero as we observe customers since they enter the system. Our results are based on a value of .85 for both smoothing constants, which is consistent with estimated values in the literature (Dorotic et al., 2014). The stock formulation of the rewards is based on the idea that rewards create a goodwill reservoir toward the LP (or memory of personalized promotional savings). The log-transformation leads to a diminishing return of rewards. We expect the effects of both stock variables to be positive, such that past rewards increase the hazard of using LPs. All variables are mean

Table 5.3. Variable operationalization and descriptive statistics.

	Variable	Operationalization	Mean	SD
<i>y</i>	IUT	Time in days between two usage events	13.534	19.242
<i>x</i>	StockCoupon	Weighted average of past stock and log of coupon value at last usage event	.191	.156
	StockFreeProduct	Weighted average of past stock and log of free product value at last usage event	.027	.070
	LaggedLogIUT	Logarithm of the IUT measured at the previous usage event	1.983	1.083
	FreeProductDummy	Dummy variable that indicates whether a customer used free products	.365	.481
<i>z</i>	MedianBasketValue	Median value of baskets in Euro	11.777	11.400
	AvgInterpurchaseTime	Average time in days between purchases	7.120	5.133
	NProducts	Number of unique products purchased	56.182	41.648
	BasketPrintRate	Ratio of the number of prints and the number of purchases	.466	.185
	RedeemedCouponStart	Indicator variable that customer has redeemed a coupon on the first usage occasion	.075	.263
	TimeEnteredSinceStart	Time in years between the date when a customer enters the LP and its launch	.215	.145
	LogTourists	Logarithm of yearly tourists (thousand) in the area of a customer's main store	6.465	1.026

Note: To obtain the same order of magnitude for all parameters, we multiply BasketPrintRate by 10 and divide NProducts by 100 in the estimation.

centered, so we can interpret the results of the baseline hazards as sample averages. Table 5.3 provides descriptive statistics before mean centering. The large standard deviations (SD) indicate considerable heterogeneity (customer and time). The coupon stock is higher than the stock of free products; only 36.5% of households redeem free products.

Estimation results. Table 5.4 contains the parameter estimates for a three-class log-logistic PHM with concomitant variables. Because $\log(\alpha) > 0$ in all classes, all baseline hazards have inverted U-shapes. The magnitudes of α and γ are similar to other applications in marketing (e.g., Seetharaman and Chintagunta, 2003). The turning points of the baseline hazards for the three classes, $tp = (\alpha - 1)^{1/\alpha}/\gamma$, and the values of the expected IUT, $E[IUT] = \pi/(\alpha\gamma \sin(\pi/\alpha))$, offer face validity. All confidence intervals and standard errors of transformations are computed using parametric bootstrapping with 10,000 draws from the estimated joint distribution of the coefficients (King et al., 2000). The three classes exhibit very different usage behavior, and our model captures this unobserved heterogeneity well.

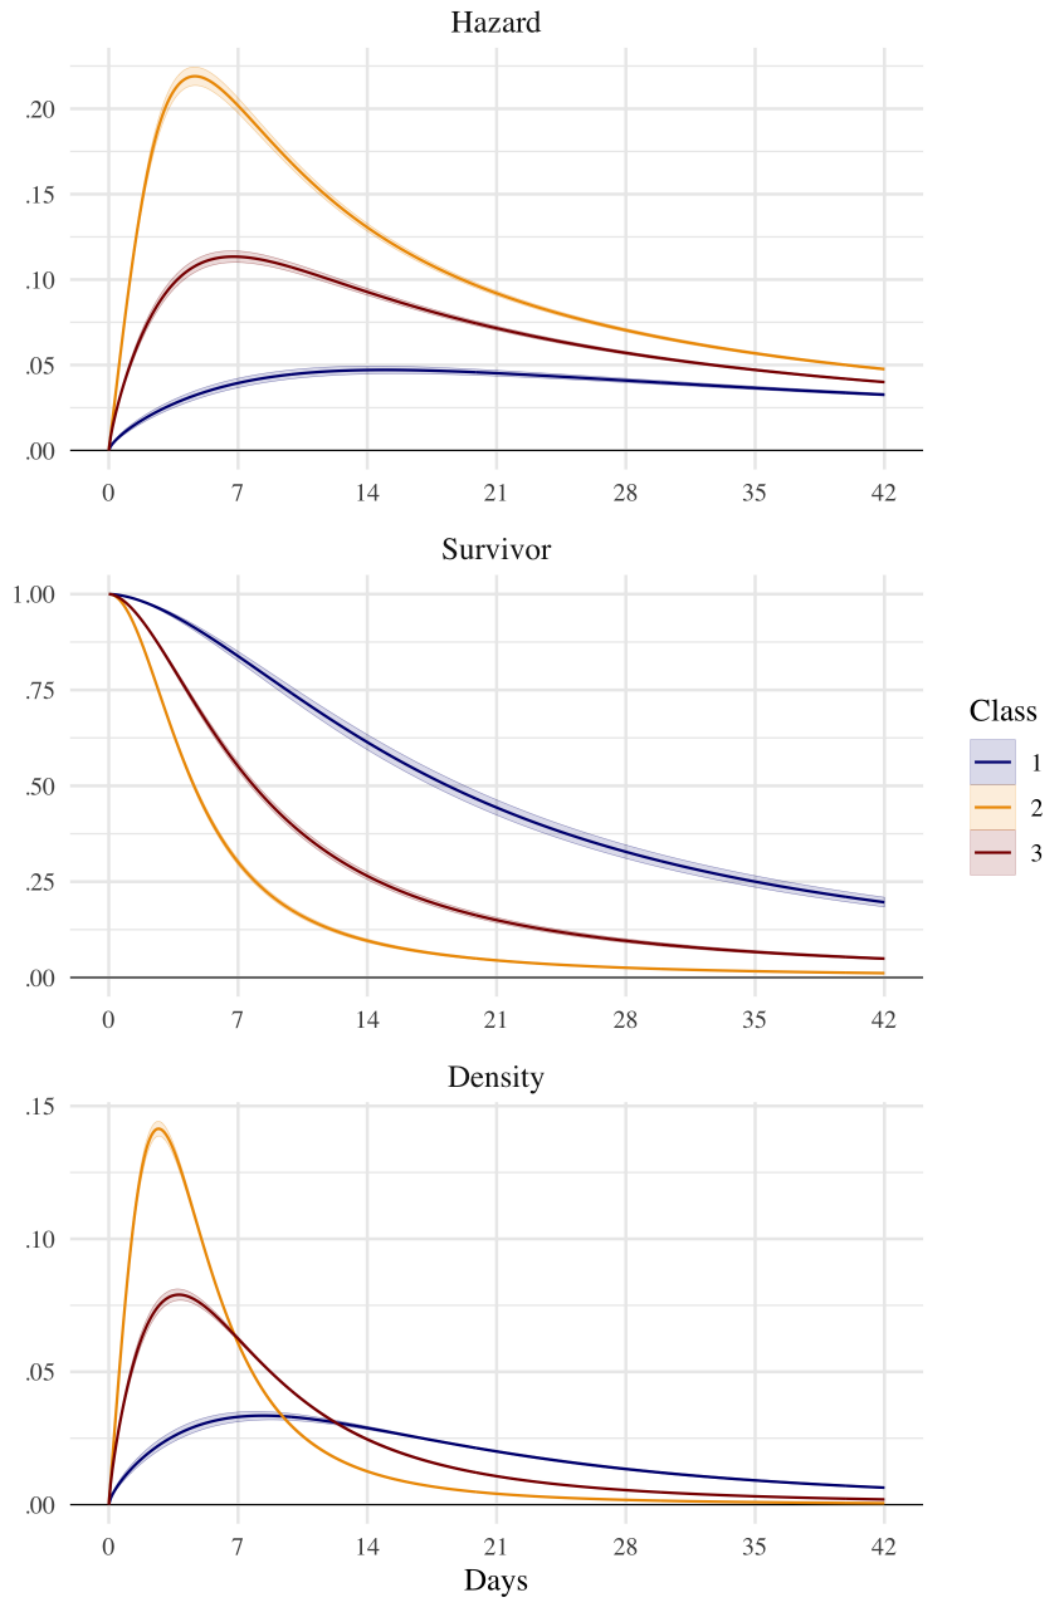
Table 5.4. PHM model estimation results.

	Class 1	Class 2	Class 3	
h_0	$\log(\alpha)$.533 (.009) ***	.704 (.007) ***	.569 (.006) ***
	$\log(\gamma)$	-2.911 (.024) ***	-1.529 (.009) ***	-2.061 (.013) ***
x	StockCoupon	1.091 (.062) ***	.580 (.045) ***	.881 (.039) ***
	StockFreeProduct	.470 (.174) ***	-.090 (.096)	.291 (.089) ***
	LaggedLogIpt	.058 (.010) ***	-.132 (.008) ***	-.089 (.006) ***
	FreeProductDummy	.111 (.024) ***	.130 (.015) ***	.169 (.014) ***
z	Intercept		-1.130 (.099) ***	.329 (.064) ***
	MedianBasketValue		-.038 (.007) ***	-.021 (.004) ***
	AvgInterpurchaseTime		-.240 (.022) ***	-.045 (.009) ***
	NProducts		.760 (.163) ***	.510 (.140) ***
	BasketPrintRate		.553 (.028) ***	.198 (.023) ***
	RedeemedCouponStart		-.780 (.173) ***	-.380 (.138) ***
	TimeEnteredSinceLaunch		-.250 (.314)	.490 (.256) *
	LogTourists		.082 (.044) *	.067 (.036) *
Class Size	.389	.156	.455	
Turning Points (SE)	14.968 (.458)	4.662 (.040)	6.757 (.093)	
Expected IUT (SE)	35.179 (.889)	7.174 (.089)	14.280 (.242)	
N Parameters		34		
LL		-325,251.5		

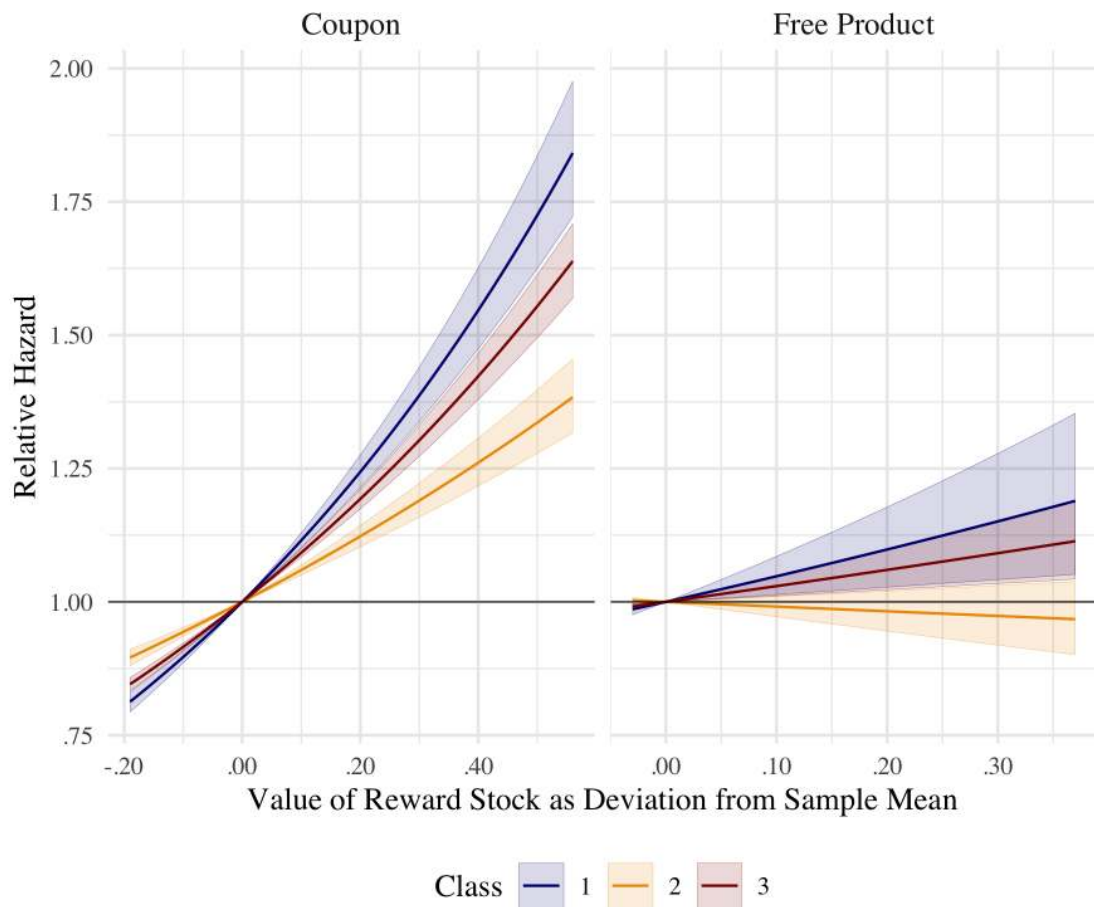
Notes: SE in parentheses. Sig. labels: * $p < .10$ and *** $p < .01$.

To glean further insights from the baseline hazards and implied survivor and density functions for each class, we plot the functions with 95% confidence intervals for IUT values up to 42 days (Figure 5.4). The hazard for class 2 reveals the highest values and is particularly peaky; the hazard for class 1 is rather flat and indicates the lowest values. The hazard of class 3 falls between the other two classes. After three weeks, the hazards of all classes become quite similar. The probability of “surviving” (i.e., not printing) until a specific day decreases fastest for class 2, followed by classes 3 and 1. After a week, more than two-thirds of class 2 customers used the system, but in class 3 less than half did, and in class 1 only roughly 15% used the promotional kiosk. The density functions offer similar insights: The hazards of the three classes have the same order in their values. After some time has passed since the last usage (e.g., 14 days) fewer customers of class 3 and class 1 print, so the order changes for the densities (i.e., unconditional probabilities of LP usage on a certain day) over time. Therefore, it is more likely to observe high IUTs

Figure 5.4. Estimated baseline hazard, survivor and density functions.



Note: Best viewed in color.

Figure 5.5. Effect of reward stock values on relative hazard.

Note: Best viewed in color.

for customers from these two classes. These findings underline the importance of accounting for customer heterogeneity when modeling IUTs. Neglecting such differences can lead to biased estimates of reward effects. Customers who generally redeem free product rewards (*FreeProductDummy*) have a higher usage hazard, with percentage differences or relative hazard values ($100(\exp(\beta_{3c}) - 1)$) of about 12%, 14%, and 18% in the three classes. The effect of *LaggedLogIUT* is negative in classes 2 and 3, such that longer durations since the last two usage events decreases the hazard of the next usage. This finding is in line with the observation that IUTs are shorter in classes 2 and 3.

The focal variables of our study—the two types of LP rewards—have positive, significant effects, except for the insignificant effect of free products in class 2. To better understand how the two types of rewards affect IUT, Figure 5.5 depicts the relative hazards as a function of reward stocks. We use the range observed in the data sample; the reference points are 0 and represent the sample averages as we mean center variables before estimation. The plot shows that the effects on the relative hazards are strong (but still reasonable), with hazards changing by -20% to $+80\%$. The coupon reward effect is strongest in class 1, followed by classes 3

and 2. For free product rewards, the order is the same, though the magnitude of the effects is lower than that for coupon rewards. Again, the effect of free products in class 2 is insignificant. For classes 1 and 3, we obtain hazard increases of 10% to 20%. The largest effect for both reward types appears in class 1, the class with the flattest baseline hazard. That is, even though the hazard of usage is rather low, higher usage can be explained by higher reward stocks.

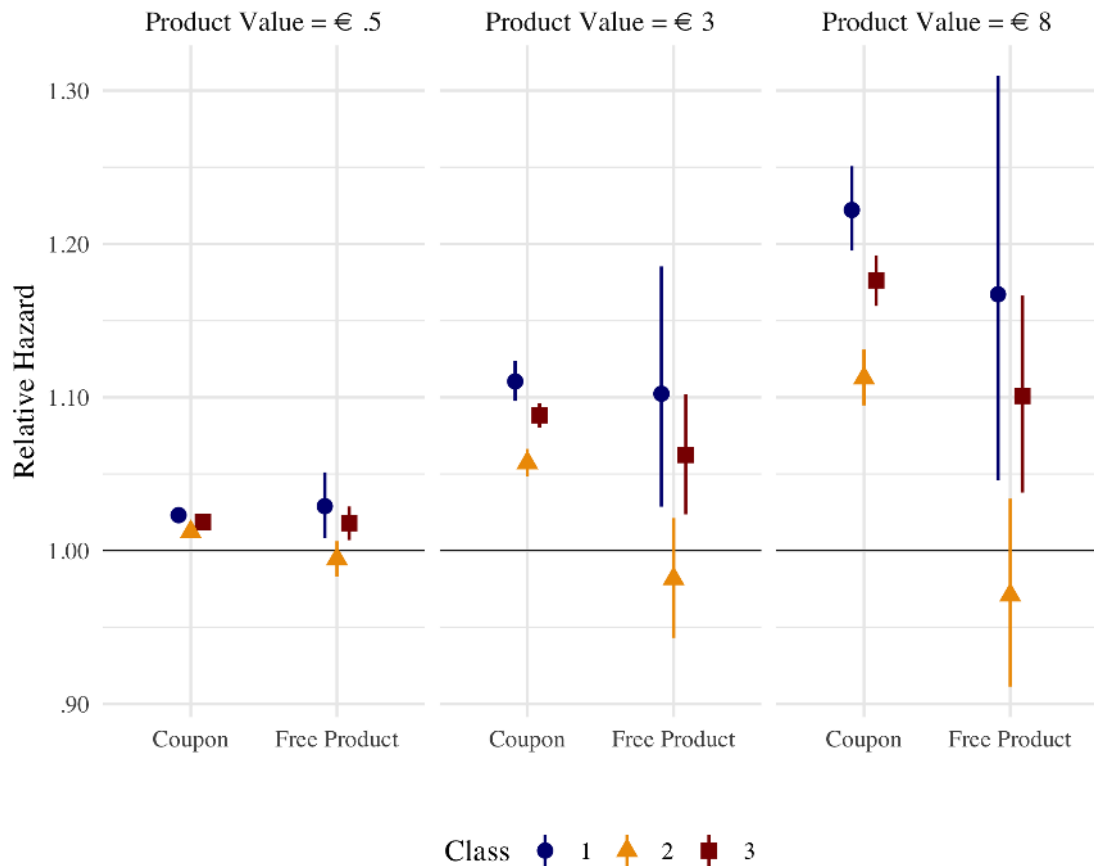
Most of the concomitant variables have significant effects and explain some heterogeneity in the prior class probabilities. For example, it is more likely that a customer who has redeemed a coupon right at the start of entering the LP or with higher value baskets, higher IUTs, or less unique products, *ceteris paribus*, belongs to class 1. On the other hand, a high BasketUsageRate increases the likelihood of being in classes 2 or 3 and late adopters tend to belong to class 3. Conditioning on the observed IUTs, we compute the posterior class probability with Bayes' theorem, given the estimated parameters (Wedel and Kamakura, 2012), and assign each customer to the class with the highest value (Gupta and Chintagunta, 1994).

This results in 2,871, 1,149, and 3,353 customers in the three classes. For 41.3%, 33.1%, and 37.5% of all print events, customers in classes 1, 2, and 3 redeem (at least) one coupon. These values are substantially higher than the industry average for coupons (Osuna et al., 2016), likely because the RTO engine works well, the kiosks are located in stores, and the timing of the reward is appropriate (Heilman et al., 2002). The sum of redeemed coupons per customer ranges from approximately 7.4 (class 1) over 12.3 (class 3) to 21.3 (class 2). For the sum of redeemed free products, this pattern remains the same, but the values are much lower (.5, .9, and 1.7).

Comparing these values with the estimated effects of reward stocks in the model reveals an interesting pattern. For example, class 1, which exhibits the highest effects, has the highest value for the percentage of usage events that lead to a coupon redemption, but the lowest number of coupon and free product redemptions. As the base hazard in the class is rather flat, the rewards are effective means for stimulating LP usage. For class 2, we observe the opposite. This class has the highest number of coupon redemptions, but the lowest effect of the rewards on the IUT, and the effect of free products is not significant. Class 3 lies between classes 1 and 2.

5.4.3 Discussion of Findings and Cost Evaluation

Overall, our findings illustrate the value that “pull”-based in-store promotions have for customers. All classes show a positive relationship between system usage and personalized coupons, though the effect sizes differ significantly. When customers already exhibit low IUTs it is more difficult to use rewards to change usage

Figure 5.6. Effect of reward values on relative hazards.

Note: Best viewed in color.

behavior. If IUTs are short, the long-term benefits of rewards do not materialize as strongly. For both reward types, the impact of rewards on LP is higher for customers who have larger baskets and redeem coupons on their first kiosk usage. The same is true for late adopters, who already have less prior engagement with the retailer (Demoulin and Zidda, 2009).

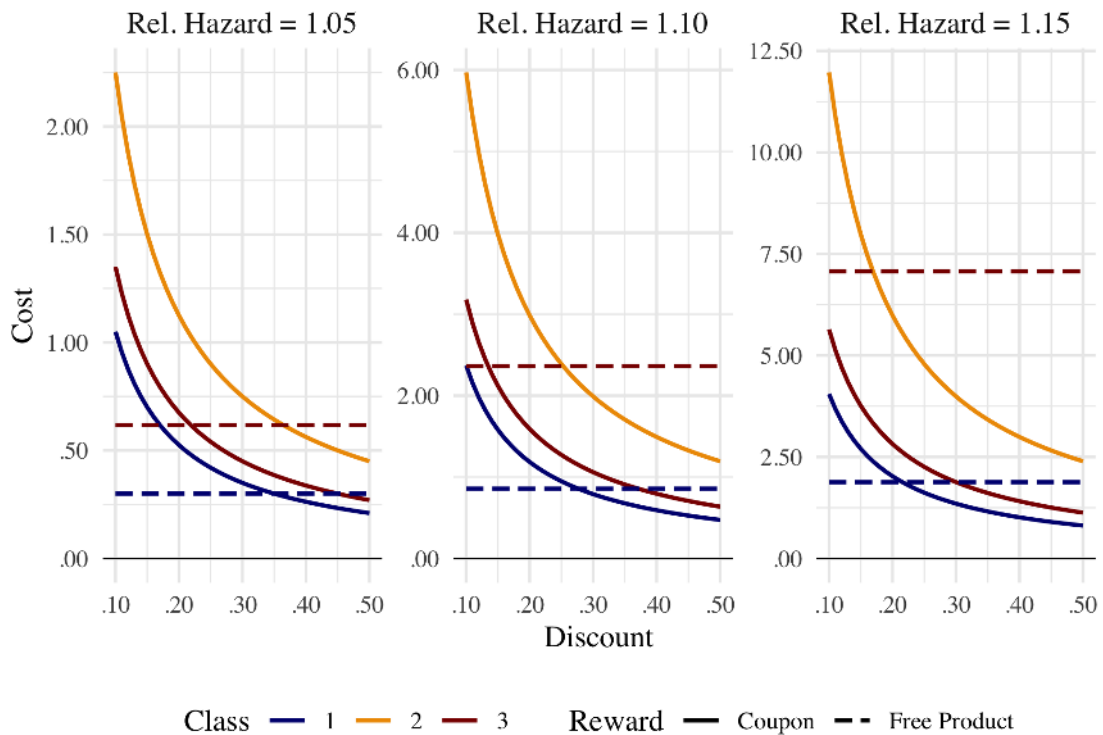
To gain a deeper understanding of the reward effects, Figure 5.6 visualizes the effect of coupons and free products on hazards. That is, instead of analyzing the effect of the stock, we plot hazards uplift as a result of rewards with different magnitudes and compare the effects over classes and reward types. We use products with prices of € .5, € 3, and € 8, which is in line with the values of products used in the LP. Examples for the three price tiers are the product categories yogurt, toothpaste, and detergent. For coupon rewards, we assume a discount of 30%, a value close to the average discount observed in the data. Therefore, a coupon only provides 30% of the product's price as a reward, while the reward of a free product is the full price. As before, we evaluate effects as deviations from the sample average that is relative to the baseline hazard.

A coupon for an €8 product leads to relative hazards of 22%, 6%, and 9% in classes 1, 2, and 3, respectively. Interestingly, the effects are somewhat lower for free product rewards, even though the full value of the product enters the stock variable for this type of reward. The relative hazards are 17% (class 1) and 10% (class 3), while class 2 has no significant effect. The 95% confidence intervals for the relative hazards of free products are rather large and do overlap the confidence intervals of the corresponding coupon rewards. The results for €3 products (middle panel) and €5 products (left panel) as rewards are similar. However, the magnitudes of relative hazards and the differences between classes are smaller. The relative hazards for a reward based on a €3 price range from 6% to 11% and for a price of €5 from 2% to 3%. Both reward types increase the hazard of (future) LP usage and, more importantly, even single rewards have meaningful effects.

For analyzing the rewards from the retailer's point of view, we next look at the costs of products that are necessary to shift hazards by a particular amount. Hence, we fix the effect to a desired hazard increase, solve the Equations 5.4 and 5.5 of the stock variables for the monetary value of a reward, and transform it into costs for the retailer by assuming a cost rate of .3. Figure 5.7 shows the results for relative hazards of 1.05, 1.1, and 1.15. For coupons, we vary discount levels between 10% and 50%. The discounts do not affect free product rewards (dashed lines).

To achieve a relative hazard of 1.1 (middle panel) a free product reward costs the retailer about €86 in class 1. The cost of a product for the corresponding coupon reward depends on the discount and varies between €47 (50% discount) and €2.37 (10%). The retailer needs products with considerably higher (lower) costs in the case of a low (high) discount to realize a specific reward value. In this example, products for coupon rewards have lower costs than for free product rewards if the discount exceeds 28%. In class 3, the costs for products are higher for both reward types compared to class 1, which is driven by the lower estimates of effects in the model. As soon as discounts are higher than 14%, coupons have lower costs compared to free products. The plots for relative hazards of 1.05 and 1.15 look similar. Lower (higher) values of the relative hazard are also related to lower (higher) costs of products for the rewards. Typical cost values for relative hazards 1.05 range from €5 to €2.5. These values increase up to over €12 in the case of a relative hazard of 1.15. The discount level at which the costs of products for coupon and free product rewards are equal is a function of the relative hazard. Higher (lower) relative hazards lead to lower (higher) discount limits. An interesting case can be observed for a hazard increase of 15%. Here, coupons are always cheaper than free products for class 3, no matter the discount.

The analysis provides interesting insights for retailers regarding designing and managing their LPs. Although retailers cannot directly affect whether and which rewards customers redeem, they can influence which rewards are offered to the

Figure 5.7. Costs of products to achieve a certain relative hazard value.

Note: Best viewed in color.

customers. Given that a retailer has a particular goal in terms of relative hazards in mind (which in turn translate into shifts in IUT), the retailer can assess the value and costs of the rewards. A better targeting algorithm, for example, might lead to lower discounts, but this might reduce the usage of the system because of lower rewards. Also, even though the estimation results highlight that the effects of free products are lower (see Table 5.4 as well as Figures 5.5 and 5.6), the differences are less pronounced when it comes to costs. Coupons have a higher effect, but only a fraction of the value (determined by the discount) contribute to the stock variable. Which type of reward is more cost-effective from a cost perspective depends on the class-specific reward effects and the discount levels. In most cases, coupons are cheaper than free product rewards because their effect on LP usage is significantly stronger than the effect of free products.

5.4.4 Supporting Insights from an Online Experiment

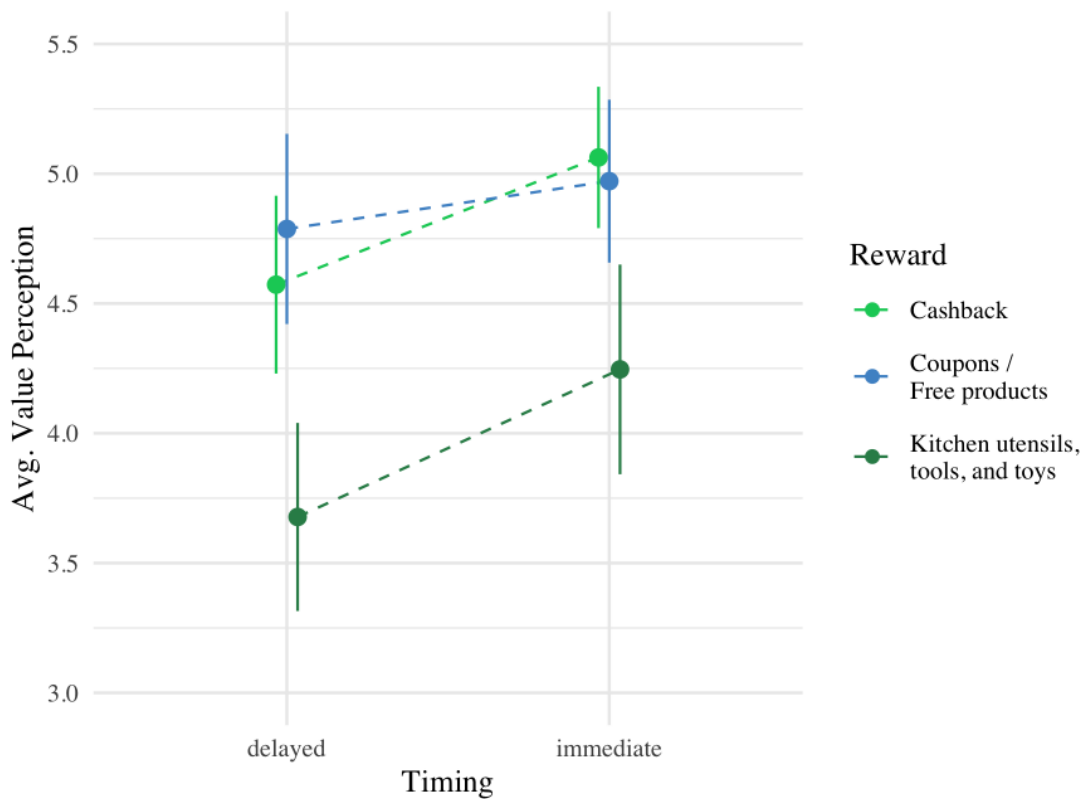
To put our results into perspective, test their robustness, and increase the generalizability of our findings, we conducted an additional online experiment on Amazon MTurk in July 2017. Following the experimental design of Yi and Jeon (2003), 410 respondents rated their value perceptions of different LP rewards. Each respondent was randomly assigned to one reward type that we described in a vignette (see Appendix 5.6).

The five-item value perception scale was based on O'Brien and Jones (1995), measured on a seven-point scale. All rewards were framed in a fictive setting, such that customers imagined spending \$100 per shopping trip on grocery purchases at a store where members of the LP saved an average of 3%. The rewards were manipulated in a 3×2 design (three reward types and two reward timings), featuring cashback, personalized coupons/free grocery product rewards, or non-grocery products such as kitchen utensils, tools, and toys, as well as immediate rewards at the end of each shopping trip, or delayed rewards after every tenth shopping trip.

The rewards in the empirical study thus span direct and immediate personalized coupons, as well as direct and delayed rewards in the form of free grocery products. The other four conditions serve two purposes. First, we seek to replicate Yi and Jeon (2003) and Keh and Lee (2006) results in a grocery retailing context. The replication of prior research in the context of grocery retailing also supports the validity of the MTurk sample (Laurent, 2013). Second, the (conceptual) replication improves understanding of how LPs work in general and provides a valuable context for the findings in this study. Given that the focal retailer only uses the two reward types, the online study enables us to compare our results with results for other popular reward types in a grocery setting.

Figure 5.8 shows the average of scores of the perceived value scale across rewards. In general, immediate, direct rewards evoke higher perceived values, consistent with previous findings (Yi and Jeon, 2003; Keh and Lee, 2006; Meyer-Waarden, 2015). The between-subjects ANOVA shows that these differences are significant (type $F(1, 406)$, 35.556, $p < .01$; timing $F(1, 406)$, 8.655, $p < .01$), but their interaction is not ($F(1, 406)$, .533, $p = .466$). Immediate cashback have a slightly higher value than personalized coupons. Interestingly, this effect is the opposite for delayed rewards. Free products (direct delayed rewards) have higher perceived value than delayed cashback (although the differences are not statistically significant) so the value of free product rewards does not seem to suffer from the delayed timing. An intuitive explanation for this observation is that receiving a product for free is more noticeable and memorable than receiving cashback of the same value.

These findings support the results of the duration analysis in Section 5.4.2 in that they underline the usefulness of personalized coupons and free products as LP rewards. Personalized coupons and immediate cashback are perceived as almost equally valuable and free products yield even higher value scores than delayed cashback. In both studies, the behavioral effect of personalized promotions is stronger than the effect of free products. Even though the value of a single reward is not high, personalized coupons are redeemed more frequently (i.e., 40% of usage events have a coupon redemption). Waiting for rewards of a similar monetary value reduces their attractiveness.

Figure 5.8. Perceived value of LP rewards.

Note: Best viewed in color.

5.5 Conclusion

In this paper, we have empirically studied an LP at a major German grocery retailer. The retailer uses in-store kiosks, an increasingly popular retail technology, at the entrance of each store to communicate with its customers and to distribute two types of rewards: vouchers for free products in exchange for loyalty points and (exclusive) personalized coupons. In this “pull”-based channel (Marketing Science Institute, 2016), customers control the information flow so we introduce system usage as a key proximal outcome. Given that research on kiosk systems and LP usage is scarce, this study provides relevant insights for researchers and practitioners. A rich longitudinal data set, that contains data for more than 7,000 customers over a period of 60 weeks, makes it possible to (1) analyze how personalized coupons affect LP usage, (2) compare the effect of personalized coupons to that of classic LP rewards, (3) study differences in effectiveness across customer segments, and (4) derive pertinent implications for reward design.

We find that both reward types increase usage (i.e., decrease the time between usage events). For example, even small rewards such as a 30% discount for a typical grocery product can increase the hazard of using the system by up to 25%. This provides clear empirical support for the value of “pull”-based in-store

promotions. Consequently, managers can use both types of rewards, free products and personalized discounts, to increase LP usage.

Interestingly, personalized coupons have a stronger effect on usage than classic LP rewards. This might be driven by the fact that coupon rewards are redeemed more often (i.e., 40% of usage events have a coupon redemption) and waiting for free product rewards is less attractive. Also, the surprise character of personalized coupons works in favor of this type of reward (Heilman et al., 2002). Given these results, coupon rewards appear to be more appealing to retailers. The cost analysis has shown that personalized coupons are cheaper than free product rewards in most situations, given that only a fraction of the price is discounted. Additionally, coupons serve other retail goals as well, as they are designed to increase sales and profits. The reward effect of personalized coupons is therefore an additional benefit, and hence their costs should be treated accordingly. Moreover, costs for coupons can be subsidized by brands, which is typically not possible for classical LP rewards, such as free products in exchange for LP points.

In the analysis of LP usage, we have identified three, sizable customer segments with very different usage patterns. The effects of rewards vary significantly across customers, emphasizing the need to account for customer heterogeneity. Customers with relatively high IUTs are affected more strongly by rewards compared to customers with short IUTs. As these classes are easy to identify, also based on the results for the concomitant variables, retailers can use our results to design their LPs accordingly.

The additional MTurk experiment supports the findings from the duration analysis. The study establishes that personalized coupons and free products as LP rewards provide value to customers and that their value is similar to that of cashback and higher than indirect rewards in the form of non-grocery products. Direct rewards are attractive to a broader set of customers and can increase customer satisfaction (Keh and Lee, 2006). While immediate rewards are clearly preferred by customers, free products are the most valued delayed reward. This underlines the usefulness of both types as LP rewards. A key benefit of direct rewards in grocery retailing is a reduced supply chain complexity which further materializes in cost savings.

Finally, personalized coupons are not only a cost-effective way to increase LP usage but also have the potential to be a new revenue stream for retailers in the form of programmatic target marketing platforms (Pathak, 2017; Chen and Friesz-Martin, 2018). Based on our study we recommend retailers to tightly integrate LPs and personalized coupons. As rewards increase usage, this also indirectly leads to more customer data and thus increases the quality of the targeting in RTO engines.

We identify several fruitful avenues for future research. First, we have studied in-store kiosk systems as technology to provide customers with information and “pull”-based personalized promotions. It would be interesting to study other channels to reach customers in offline retailing. For example, apps are gaining more popularity so the interplay with kiosk systems as well as the comparison of reward effects across channels appear to be promising directions for future research. Second, we have employed observational data and data from an online experiment to obtain our results. Given that we have documented robust (and heterogeneous) reward effects, it would be very useful to conduct a comparison of rewards in a controlled field experiment. Lastly, we have focused only on the effect of rewards on time between usage events but have not investigated the impact of rewards on revenues and profits. Further research can also evaluate the interplay between usage and financial performance metrics.

5.6 Appendix

5.6.1 Details Regarding the Estimation of the Log-Logistic Latent Class PHM

Extant studies offer in-depth details regarding the model specification and estimation (e.g., Seetharaman and Chintagunta, 2003), so we summarize the main steps here. The hazard function can also be written as $h(t, x_i) = f(t, x_i)/S(t, x_i)$, where $f(t, x_i)$ and $S(t, x_i)$ are the probability density and survivor function, respectively. The former is the unconditional probability of using the kiosk at t ; the latter is the probability that the customer has not used the kiosk (i.e., “survived”) until t . As we opt to account for heterogeneity using a latent class approach, all parameters and, therefore, hazard functions are class-specific. The corresponding log-likelihood (LL) function of the model is

$$LL = \sum_{i=1}^I \log \left(\sum_{c=1}^C \lambda_{ic} \left\{ \prod_{j=1}^{J_i} f_c(t_{ij} - t_{ij-1}, x_{ij}) S_c(T_i - t_{iJ}, x_{iJ+1}) \right\} \right). \quad (5.6)$$

Each customer has J_i duration intervals (“spells”), and c indexes the latent classes. Let t_{ij} denote the calendar time of the j -th print observation of customer i , such that $t_{ij} - t_{ij-1}$ is the IUT; x_{ij} is customer i ’s covariate vector for spell j ; and T_i is the calendar time of a right-censored observation of customer i . We compute the maximum likelihood estimates of the model parameters using the gradient-based Broyden-Fletcher-Goldfarb-Shanno (BFGS) algorithm (see Appendix 5.6 for the derivation of the LL and its partial derivatives that are used in the BFGS algorithm). To avoid convergence to local optima, we estimate all models with 100 random starting values each and retain the solution with the highest LL value (Wedel and Kamakura, 2012).

To select the number of classes C , we ran multiple models with one to five classes and compared the results. In particular, we re-estimated the models using 80% of the customers and calculated the LL values in the holdout sample. We repeated this procedure 10 times to reduce variance in the results and to prevent edge cases. Information criteria using the full sample point to a five-class solution, but a closer inspection of the average hold-out LL values reveals that the increase in fit is marginal ($< .2\%$) for solutions for more than three classes. Furthermore, all models with more than three classes had at least one class with very few customers ($\ll 5\%$) and such small segments provide little relevant insights and might be an artifact of the method employed to derive the classes (DeSarbo and DeSarbo, 2001). Indeed, the additional, extremely small classes were very similar in terms of shape for the baseline hazard. Therefore, we chose the model with

$C = 3$, which provides meaningful results and is a good compromise between fit and complexity. Other applications of PHMs use similar numbers of segments (Seetharaman and Chintagunta, 2003). Our proposed full model with three classes (LL = -325,251.5) clearly outperforms the simpler three-class models, namely, those without concomitant variables (LL = -325,721.6), with only baseline hazards (LL = -326,826.2), and without unobserved heterogeneity (LL = -330,428.1).

5.6.2 Derivation of the Partial Derivatives

5.6.2.1 Homogeneous Model

Using the hazard function

$$h_0(t) = \gamma\alpha(\gamma t)^{\alpha-1} \cdot (1 + (\gamma t)^\alpha)^{-1}$$

and the survival function

$$S(t) = (1 + (\gamma t)^\alpha)^{-1}$$

of the log-logistic PHM, we can specify the probability density function as

$$f(t, x_i) = h_0(t) \cdot \exp(x_i\beta) \cdot S(t).$$

With $a_i = \exp(x_i\beta)$, $b_i = (\gamma t)^\alpha$, and $m = \sum_i \delta_i$, the log-likelihood LL is (ignoring multiple observation per person for simplicity):

$$\begin{aligned} LL = & m \log(\alpha) + m \alpha \log(\gamma) + (\alpha - 1) \sum_i \delta_i \log(t_i) \\ & - \sum_i \delta_i \log(1 + b_i) + \sum_i \delta_i \log(a_i) - \sum_i a_i \log(1 + b_i). \end{aligned}$$

We define

$$c_i = \frac{\log(b_i)}{1 + b_i}$$

and

$$d_i = \frac{b_i}{1 + b_i}.$$

Differentiating LL w.r.t. α , β , and γ leads to:

$$\partial_\alpha LL = \frac{m}{\alpha} + m \log(\gamma) + \sum_i \delta_i \log(t_i) - \sum_i \delta_i \frac{b_i c_i}{\alpha} - \sum_i a_i \frac{b_i c_i}{\alpha},$$

$$\partial_{\beta_j} LL = \sum_i \delta_i x_j + \sum_i x_j a_i \log(1 + b_i)$$

and

$$\partial_{\gamma} LL = \frac{m\alpha}{\gamma} - \frac{\alpha}{\gamma} \sum_i \delta_i d_i - \frac{\alpha}{\gamma} \sum_i \alpha_i d_i.$$

5.6.2.2 Latent-Class Model

The log-likelihood with multiple latent classes is:

$$LL = \sum_i \log\left(\sum_c \lambda_{ic} \exp(LL_{ic})\right)$$

with

$$\lambda_{ic} = \frac{\exp(z_i \theta_c)}{\sum_c \exp(z_i \theta_c)}$$

and

$$\begin{aligned} LL_{ic} &= m \log(\alpha_c) + m \alpha_c \log(\gamma_c) + (\alpha_c - 1) \sum \delta_{ic} \log(t_i) \\ &\quad - \sum \delta_{ic} \log(1 + b_{ic}) + \sum \delta_{ic} \log(a_{ic}) - \sum a_{ic} \log(1 + b_{ic}). \end{aligned}$$

Defining $\phi_{ic} = \exp(LL_{ic})$ and differentiating LL w.r.t. α_c , β_c , γ_c , and θ_c results in:

$$\partial_{\alpha_c} LL = \sum_i \frac{1}{\sum_c \lambda_{ic} \phi_{ic}} \partial_{\alpha_c} \left\{ \sum_c \lambda_{ic} \phi_{ic} \right\} = \sum_i \frac{\lambda_{ic} \phi_{ic}}{\sum_c \lambda_{ic} \phi_{ic}} \partial_{\alpha_c} LL_{ic},$$

$$\partial_{\beta_{jc}} LL = \sum_i \frac{1}{\sum_c \lambda_{ic} \phi_{ic}} \partial_{\beta_{jc}} \left\{ \sum_c \lambda_{ic} \phi_{ic} \right\} = \sum_i \frac{\lambda_{ic} \phi_{ic}}{\sum_c \lambda_{ic} \phi_{ic}} \partial_{\beta_{jc}} LL_{ic},$$

$$\partial_{\gamma_c} LL = \sum_i \frac{1}{\sum_c \lambda_{ic} \phi_{ic}} \partial_{\gamma_c} \left\{ \sum_c \lambda_{ic} \phi_{ic} \right\} = \sum_i \frac{\lambda_{ic} \phi_{ic}}{\sum_c \lambda_{ic} \phi_{ic}} \partial_{\gamma_c} LL_{ic}$$

and

$$\partial_{\theta_{kc}} LL = \partial_{\theta_{kc}'} \left\{ \sum_i \log\left(\sum_c \lambda_{ic} \exp(LL_{ic})\right) \right\} = \sum_i \frac{1}{\sum_c \lambda_{ic} \phi_{ic}} \sum_c \phi_{ic} \partial_{\theta_{kc}'} \lambda_{ic}.$$

The partial derivative of λ_{ic} w.r.t. θ_{kc} is:

$$\begin{aligned} \partial_{\theta_{kc}} \lambda_c &= \partial_{\theta_{kc}'} \left\{ \exp(z_i \theta_c) \right\} \frac{1}{\sum_c \exp(z_i \theta_c)} + \exp(z_i \theta_c) \partial_{\theta_{kc}'} \left\{ \frac{1}{\sum_c \exp(z_i \theta_c)} \right\} \\ &= \frac{\mathbb{1}^{c'} \exp(z_i \theta_c) z_k}{\sum_c \exp(z_i \theta_c)} - \frac{\exp(z_i \theta_c)}{(\sum_c \exp(z_i \theta_c))^2} \exp(z_i \theta_c) z_k \\ &= z_k \lambda_{ic} (\mathbb{1}^{c'} - \lambda_{ic}) \end{aligned}$$

with

$$\mathbb{1}^{c'} = 1 \text{ if } c = c' \text{ and } 0 \text{ otherwise.}$$

Intermediate steps of gradient calculation and an implementation in Python are available on request.

5.6.3 Vignettes for the Online Experiment

The following instruction was presented to the respondents in the MTurk study as part of the online experiment. Each respondent was assigned to one of six groups. The general presentation of the LP description was identical for all six groups, but we systematically varied the reward mechanism.

Please imagine the following situation. Your average shopping basket at the retailer is approx. \$100 and you never spend less than \$50 per shopping trip. Assume that the retailer introduces a new loyalty program (until now there was no loyalty program). If you participate, you receive a loyalty card (you can choose between a key ring and a plastic card) which is scanned at the checkout, so the retailer knows how much you spend during your shopping trips. The loyalty program rewards you receive depend on your revenue at the retailer. As a reward for participating in the program, you receive

Group 1 a 3% discount on every shopping trip (e.g., \$3 for a \$100 shopping basket).

Group 2 a \$30 discount after spending \$1,000 (e.g., after 10 shopping trips with an average basket size of \$100). This equals a 3% discount.

Group 3 exclusive access to coupons at an in-store kiosk system. The coupons are personalized to your preferences and based on your purchase history. Consumers save on average 3% on every shopping trip (e.g., \$3 for a \$100 shopping basket).

Group 4 5 loyalty points for every \$100 you spend at the retailer. You can exchange loyalty points for free products after you have collected 50 points (e.g., after 10 shopping trips with an average basket size of \$100). For 50 points consumers receive free products that have a value of \$30. This equals a 3% discount.

Group 5 a non-grocery product such as a free kitchen utensil (e.g., bowl, plate, knife), a tool (e.g., screwdriver, hammer, wrench) or a toy (e.g., card game, stuffed animal) on every shopping trip. The rewards have a value that equals 3% of your last shopping basket (e.g., \$3 for a \$100 shopping basket).

Group 6 free non-grocery products such as kitchen utensils (e.g., bowls, plates, knives), tools (e.g., screwdrivers, hammers, wrenches) or toys (e.g., card games, stuffed animals) with a total value of \$30 after spending \$1,000 (e.g., after 10 shopping trips with an average basket size of \$100). The total value of the rewards equals 3% of your revenue.

After reading this description, we would like to understand how much you like the proposed loyalty program. Please state how much you agree with the following statements. When responding, please think about the value of the rewards and how much effort on your part (e.g., remembering to bring your loyalty card and showing it at the checkout) is needed.

Bibliography

- Adelman, M. A. (1969). Comment on The "H" Concentration Measure as a Numbers-Equivalent. *The Review of Economics and Statistics*, pages 99–101.
- Aguirre, E., Mahr, D., Grewal, D., de Ruyter, K., and Wetzels, M. (2015). Unraveling the Personalization Paradox: The Effect of Information Collection and Trust-Building Strategies on Online Advertisement Effectiveness. *Journal of Retailing*, 91(1):34–49.
- Ailawadi, K. L. and Gupta, S. (2014). Sales Promotions. In *History of Marketing Science*, pages 463–497. Now Publishers.
- Alain, G. and Bengio, Y. (2014). What Regularized Auto-Encoders Learn From the Data-Generating Distribution. *The Journal of Machine Learning Research*, 15(Jan):3563–3593.
- Allenby, G. M., Leone, R. P., and Jen, L. (1999). A Dynamic Model of Purchase Timing With Application to Direct Marketing. *Journal of the American Statistical Association*, 94(446):365–374.
- Ansari, A. and Mela, C. F. (2003). E-Customization. *Journal of Marketing Research*, 40(2):131–145.
- Archak, N., Ghose, A., and Ipeirotis, P. G. (2011). Deriving the Pricing Power of Product Features by Mining Consumer Reviews. *Management Science*, 57(8):1485–1509.
- Arora, N., Dreze, X., Ghose, A., Hess, J. D., Iyengar, R., Jing, B., Joshi, Y., Kumar, V., Lurie, N., Neslin, S., et al. (2008). Putting One-to-One Marketing to Work: Personalization, Customization, and Choice. *Marketing Letters*, 19(3):305.
- Bell, D. R., Chiang, J., and Padmanabhan, V. (1999). The Decomposition of Promotional Response: An Empirical Generalization. *Marketing Science*, 18(4):504–526.
- Belloni, A., Freund, R., Selove, M., and Simester, D. (2008). Optimizing Product Line Designs: Efficient Methods and Comparisons. *Management Science*, 54(9):1544–1552.

- Bengio, Y., Lamblin, P., Popovici, D., and Larochelle, H. (2007). Greedy Layer-wise Training of Deep Networks. In *Advances in Neural Information Processing Systems*, pages 153–160.
- Bergstra, J. and Bengio, Y. (2012). Random Search for Hyper-Parameter Optimization. *Journal of Machine Learning Research*, 13(Feb):281–305.
- Berry, L. L. (1995). Relationship Marketing of Services—Growing Interest, Emerging Perspectives. *Journal of the Academy of Marketing Science*, 23(4):236–245.
- Biafore, M. (2016). How to Measure and Compare the Real Distribution Costs of Promotions. <http://insights.revtrax.com/how-to-measure-compare-the-real-distribution-costs-of-promotions> (Accessed 2019-04-03).
- Bijmolt, T. and Pieters, R. (2001). Meta-Analysis in Marketing When Studies Contain Multiple Measurements. *Marketing Letters*, 12(2):157–169.
- Bijmolt, T. H., Dorotic, M., Verhoef, P. C., et al. (2011). Loyalty Programs: Generalizations on Their Adoption, Effectiveness and Design. *Foundations and Trends® in Marketing*, 5(4):197–258.
- Bijmolt, T. H., Heerde, H. J. v., and Pieters, R. G. (2005). New Empirical Generalizations on the Determinants of Price Elasticity. *Journal of Marketing Research*, 42(2):141–156.
- Bijmolt, T. H. and Verhoef, P. C. (2017). Loyalty Programs: Current Insights, Research Challenges, and Emerging Trends. In *Handbook of Marketing Decision Models*, pages 143–165. Springer.
- Blattberg, R. C., Kim, B.-D., and Neslin, S. A. (2008). *Why Database Marketing?* Springer.
- Blattberg, R. C. and Neslin, S. A. (1990). *Sales Promotion: Concepts, Methods, and Strategies*. Prentice-Hall.
- Blei, D. M., Ng, A. Y., and Jordan, M. I. (2003). Latent Dirichlet Allocation. *Journal of Machine Learning Research*, 3(Jan):993–1022.
- Bleier, A., De Keyser, A., and Verleye, K. (2018). Customer Engagement Through Personalization and Customization. In *Customer Engagement Marketing*, pages 75–94. Springer.
- Bleier, A. and Eisenbeiss, M. (2015). Personalized Online Advertising Effectiveness: The Interplay of What, When, and Where. *Marketing Science*, 34(5):669–688.
- Bolton, R. N. (1989). The Relationship Between Market Characteristics and Promotional Price Elasticities. *Marketing Science*, 8(2):153–169.

- Bordes, A., Chopra, S., and Weston, J. (2014). Question Answering With Subgraph Embeddings. *arXiv Preprint arXiv:1406.3676*.
- Bradlow, E. T., Gangwar, M., Kopalle, P., and Voleti, S. (2017). The Role of Big Data and Predictive Analytics in Retailing. *Journal of Retailing*, 93(1):79–95.
- Breugelmans, E., Bijmolt, T. H., Zhang, J., Basso, L. J., Dorotic, M., Kopalle, P., Minnema, A., Mijnlief, W. J., and Wunderlich, N. V. (2015). Advancing Research on Loyalty Programs: A Future Research Agenda. *Marketing Letters*, 26(2):127–139.
- Breugelmans, E. and Liu-Thompkins, Y. (2017). The Effect of Loyalty Program Expiration Policy on Consumer Behavior. *Marketing Letters*, 28(4):537–550.
- Chapelle, O., Manavoglu, E., and Rosales, R. (2015). Simple and Scalable Response Prediction for Display Advertising. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 5(4):61.
- Chen, T. and Friesz-Martin, L. (2018). Quotient Technology and Albertsons Companies Launch the Retailer’s New Data-Driven Media Platform. <https://www.quotient.com/quotient-technology-albertsons-companies-launch-retailers-new-data-driven-media-platform> (Accessed 2019-02-20).
- Chintagunta, P., Hanssens, D. M., and Hauser, J. R. (2016). Editorial—Marketing Science and Big Data. *Marketing Science*, 35(3):341–342.
- Collins, E. (2017). How Consumers Really Feel About Loyalty Programs. <http://www.oracle.com/us/solutions/consumers-loyalty-programs-3738548.pdf> (Accessed 2019-02-20).
- Collobert, R., Weston, J., Bottou, L., Karlen, M., Kavukcuoglu, K., and Kuksa, P. (2011). Natural Language Processing (Almost) From Scratch. *Journal of Machine Learning Research*, 12(Aug):2493–2537.
- Covington, P., Adams, J., and Sargin, E. (2016). Deep Neural Networks for Youtube Recommendations. In *Proceedings of the 10th ACM Conference on Recommender Systems*, pages 191–198. ACM.
- Demoulin, N. T. and Zidda, P. (2009). Drivers of Customers’ Adoption and Adoption Timing of a New Loyalty Card in the Grocery Retail Market. *Journal of Retailing*, 85(3):391–405.
- DeSarbo, W. S. and DeSarbo, C. F. (2001). A Generalized Normative Segmentation Methodology Employing Conjoint Analysis. In *Conjoint Measurement*, pages 447–478. Springer.

- Dorotic, M., Verhoef, P. C., Fok, D., and Bijmolt, T. H. (2014). Reward Redemption Effects in a Loyalty Program When Customers Choose How Much and When to Redeem. *International Journal of Research in Marketing*, 31(4):339–355.
- Dubé, J.-P. H. and Misra, S. (2017). Scalable Price Targeting. Available at SSRN: <https://ssrn.com/abstract=2992257>.
- Einav, L. and Levin, J. (2014). Economics in the Age of Big Data. *Science*, 346(6210).
- eMarketer (2016). Why Marketers Don't Personalize Content. <https://www.emarketer.com/articles/print.aspx?r=1013768> (Accessed 2019-02-11).
- Erhan, D., Bengio, Y., Courville, A., Manzagol, P.-A., Vincent, P., and Bengio, S. (2010). Why Does Unsupervised Pre-Training Help Deep Learning? *Journal of Machine Learning Research*, 11(Feb):625–660.
- Fader, P. (2012). *Customer Centricity: Focus on the Right Customers for Strategic Advantage*. Wharton Digital Press.
- Fader, P. S. and Hardie, B. G. (1996). Modeling Consumer Choice Among SKUs. *Journal of Marketing Research*, pages 442–452.
- Forrester (2017). Demystifying Price and Promotion: Shoppers Bust Long Held Myths on Pricing and Promotions. <http://www.parkeravery.com/media/forrester-study-demystifying-price-and-promotion.pdf> (Accessed 2019-04-03).
- Fruend, M. (2017). An In-Depth Analysis of Where Loyalty Is Now ... And Where It's Headed. <https://www.loyalty.com/home/insights/article-details/2017-colloquy-loyalty-census-report> (Accessed 2019-04-04).
- Gabel, S., Guhl, D., and Klapper, D. (2019). P2V-MAP: Mapping Market Structure for Large Retail Assortments. *Journal of Marketing Research* (forthcoming).
- Gartner (2016). Hype Cycle for Retail Technologies, 2016. Technical report, Gartner.
- Goldfarb, A. and Tucker, C. E. (2011). Privacy Regulation and Online Advertising. *Management Science*, 57(1):57–71.
- Goodfellow, I., Bengio, Y., and Courville, A. (2016). *Deep Learning*. MIT Press. <http://www.deeplearningbook.org>.
- Goodman, J. K. and Paolacci, G. (2017). Crowdsourcing Consumer Research. *Journal of Consumer Research*, 44(1):196–210.
- Green, P. E. and Krieger, A. M. (1985). Models and Heuristics for Product Line Selection. *Marketing Science*, 4(1):1–19.

- Grewal, D., Ailawadi, K. L., Gauri, D., Hall, K., Kopalle, P., and Robertson, J. R. (2011). Innovations in Retail Pricing and Promotions. *Journal of Retailing*, 87:S43–S52.
- Grewal, D., Roggeveen, A. L., and Nordfält, J. (2017). The Future of Retailing. *Journal of Retailing*, 93(1):1–6.
- Guadagni, P. M. and Little, J. D. (1983). A Logit Model of Brand Choice Calibrated on Scanner Data. *Marketing Science*, 2(3):203–238.
- Guillot, C. (2016). Ahold Increases Loyalty and Basket Size Personalized Coupons. <https://www.nrf.com/blog/ahold-increases-loyalty-and-basket-size-personalized-coupons> (Accessed 2019-02-12).
- Gupta, S. and Chintagunta, P. K. (1994). On Using Demographic Variables to Determine Segment Membership in Logit Mixture Models. *Journal of Marketing Research*, 31(1):128–136.
- Hanssens, D. M. (2014). Econometric Models. In *The History of Marketing Science*, pages 99–128. Singapore: World Scientific Publishing.
- Hawkins, G. (2012). Will Big Data Kill All but the Biggest Retailers? *Harvard Business Review*.
- Heilman, C. M., Nakamoto, K., and Rao, A. G. (2002). Pleasant Surprises: Consumer Response to Unexpected In-Store Coupons. *Journal of Marketing Research*, 39(2):242–252.
- Hinton, G., Deng, L., Yu, D., Dahl, G. E., Mohamed, A. R., Jaitly, N., Senior, A., Vanhoucke, V., Nguyen, P., Sainath, T. N., et al. (2012). Deep Neural Networks for Acoustic Modeling in Speech Recognition: The Shared Views of Four Research Groups. *IEEE Signal Processing Magazine*, 29(6):82–97.
- Hox, J. J., Moerbeek, M., and van de Schoot, R. (2010). *Multilevel Analysis: Techniques and Applications*. Routledge, 2nd edition.
- Inman, J. J. and Nikolova, H. (2017). Shopper-Facing Retail Technology: A Retailer Adoption Decision Framework Incorporating Shopper Attitudes and Privacy Concerns. *Journal of Retailing*, 93(1):7–28.
- Jacobs, B. J., Donkers, B., and Fok, D. (2016). Model-Based Purchase Predictions for Large Assortments. *Marketing Science*, 35(3):389–404.
- Jannach, D., Zanker, M., Felfernig, A., and Friedrich, G. (2011). *Recommender Systems: An Introduction*. Cambridge University Press.

- Johnson, J., Tellis, G. J., and Ip, E. H. (2013). To Whom, When, and How Much to Discount? A Constrained Optimization of Customized Temporal Discounts. *Journal of Retailing*, 89(4):361–373.
- Kannan, P. et al. (2017). Digital Marketing: A Framework, Review and Research Agenda. *International Journal of Research in Marketing*, 34(1):22–45.
- Ke, G., Meng, Q., Finley, T., Wang, T., Chen, W., Ma, W., Ye, Q., and Liu, T.-Y. (2017). LightGBM: A Highly Efficient Gradient Boosting Decision Tree. In *Advances in Neural Information Processing Systems*, pages 3146–3154.
- Keh, H. T. and Lee, Y. H. (2006). Do Reward Programs Build Loyalty for Services?: The Moderating Effect of Satisfaction on Type and Timing of Rewards. *Journal of Retailing*, 82(2):127–136.
- King, G., Tomz, M., and Wittenberg, J. (2000). Making the Most of Statistical Analyses: Improving Interpretation and Presentation. *American Journal of Political Science*, 44(2):347–361.
- Kingma, D. P. and Ba, J. (2014). Adam: A Method for Stochastic Optimization. *arXiv Preprint arXiv:1412.6980*.
- Kivetz, R. and Simonson, I. (2002). Earning the Right to Indulge: Effort as a Determinant of Customer Preferences Toward Frequency Program Rewards. *Journal of Marketing Research*, 39(2):155–170.
- Kivetz, R., Urminsky, O., and Zheng, Y. (2006). The Goal-Gradient Hypothesis Resurrected: Purchase Acceleration, Illusionary Goal Progress, and Customer Retention. *Journal of Marketing Research*, 43(1):39–58.
- Krishnamurthi, L. and Raj, S. P. (1991). An Empirical Analysis of the Relationship Between Brand Loyalty and Consumer Price Elasticity. *Marketing Science*, 10(2):172–183.
- Krizhevsky, A., Sutskever, I., and Hinton, G. E. (2012). Imagenet Classification With Deep Convolutional Neural Networks. In *Advances in Neural Information Processing Systems*, pages 1097–1105.
- Lambrecht, A. and Tucker, C. (2013). When Does Retargeting Work? Information Specificity in Online Advertising. *Journal of Marketing Research*, 50(5):561–576.
- Laurent, G. (2013). EMAC Distinguished Marketing Scholar 2012: Respect the Data! *International Journal of Research in Marketing*, 30(4):323–334.
- LeCun, Y., Bengio, Y., and Hinton, G. (2015). Deep Learning. *Nature*, 521(7553):436–444.

- Leenheer, J., Van Heerde, H. J., Bijmolt, T. H., and Smidts, A. (2007). Do Loyalty Programs Really Enhance Behavioral Loyalty? An Empirical Analysis Accounting for Self-Selecting Members. *International Journal of Research in Marketing*, 24(1):31–47.
- Lewis, M. (2004). The Influence of Loyalty Programs and Short-Term Promotions on Customer Retention. *Journal of Marketing Research*, 41(3):281–292.
- Lindsay, K. (2014). Personalization Payoff ROI Getting Personal. <https://theblog.adobe.com/personalization-payoff-roi-getting-personal> (Accessed 2019-02-13).
- Liu, D.-R., Lai, C.-H., and Lee, W.-J. (2009). A Hybrid of Sequential Rules and Collaborative Filtering for Product Recommendation. *Information Sciences*, 179(20):3505–3519.
- Liu, L., Dzyabura, D., and Mizik, N. (2018). Visual Listening In: Extracting Brand Image Portrayed on Social Media. Available at SSRN: <https://ssrn.com/abstract=2978805>.
- Liu, X., Lee, D., and Srinivasan, K. (2017). Large Scale Cross Category Analysis of Consumer Review Content on Sales Conversion Leveraging Deep Learning. Available at SSRN: <https://ssrn.com/abstract=2848528>.
- Liu, Y. (2007). The Long-term Impact of Loyalty Programs on Consumer Purchase Behavior and Loyalty. *Journal of Marketing*, 71(4):19–35.
- Liu, Y. and Yang, R. (2009). Competing Loyalty Programs: Impact of Market Saturation, Market Share, and Category Expandability. *Journal of Marketing*, 73(1):93–108.
- Maaten, L. v. d. and Hinton, G. (2008). Visualizing Data Using t-SNE. *Journal of Machine Learning Research*, 9(Nov):2579–2605.
- Manchanda, P., Ansari, A., and Gupta, S. (1999). The Shopping Basket: A Model for Multicategory Purchase Incidence Decisions. *Marketing Science*, 18(2):95–114.
- Manchanda, P., Dubé, J.-P., Goh, K. Y., and Chintagunta, P. K. (2006). The Effect of Banner Advertising on Internet Purchasing. *Journal of Marketing Research*, 43(1):98–108.
- Marketing Science Institute (2016). MSI Research Priorities 2016–2018. https://www.msi.org/uploads/articles/msi_rp16-18.pdf (Accessed 2019-02-20).
- McFadden, D. (1974). The Measurement of Urban Travel Demand. *Journal of Public Economics*, 3(4):303–328.

- Meyer-Waarden, L. (2007). The Effects of Loyalty Programs on Customer Lifetime Duration and Share of Wallet. *Journal of Retailing*, 83(2):223–236.
- Meyer-Waarden, L. (2015). Effects of Loyalty Program Rewards on Store Loyalty. *Journal of Retailing and Consumer Services*, 24:22–32.
- Mild, A. and Reutterer, T. (2003). An Improved Collaborative Filtering Approach for Predicting Cross-Category Purchases Based on Binary Market Basket Data. *Journal of Retailing and Consumer Services*, 10(3):123–133.
- Milkman, K. L., Rogers, T., and Bazerman, M. H. (2010). I'll Have the Ice Cream Soon and the Vegetables Later: A Study of Online Grocery Purchases and Order Lead Time. *Marketing Letters*, 21(1):17–35.
- Minnema, A., Bijmolt, T. H., and Non, M. C. (2017). The Impact of Instant Reward Programs and Bonus Premiums on Consumer Purchase Behavior. *International Journal of Research in Marketing*, 34(1):194–211.
- Montgomery, A. L. and Smith, M. D. (2009). Prospects for Personalization on the Internet. *Journal of Interactive Marketing*, 23(2):130–137.
- Murthi, B. and Sarkar, S. (2003). The Role of the Management Sciences in Research on Personalization. *Management Science*, 49(10):1344–1362.
- Naik, P., Wedel, M., Bacon, L., Bodapati, A., Bradlow, E., Kamakura, W., Kreulen, J., Lenk, P., Madigan, D. M., and Montgomery, A. (2008). Challenges and Opportunities in High-Dimensional Choice Data Analyses. *Marketing Letters*, 19(3):201.
- Narasimhan, C., Neslin, S. A., and Sen, S. K. (1996). Promotional Elasticities and Category Characteristics. *The Journal of Marketing*, 60(2):17–30.
- NCH Marketing Services (2019). 2018 Year-End Coupon Facts at a Glance. <https://www.nchmarketing.com/2018-year-end-coupon-facts-at-a-glance.aspx> (Accessed 2019-04-03).
- Neslin, S. A., Van Heerde, H. J., et al. (2009). Promotion Dynamics. *Foundations and Trends® in Marketing*, 3(4):177–268.
- O'Brien, L. and Jones, C. (1995). Do Rewards Really Create Loyalty? *Harvard Business Review*, 73(3):75–82.
- Orhan, A. E. and Pitkow, X. (2017). Skip Connections Eliminate Singularities. *arXiv Preprint arXiv:1701.09175*.
- Osuna, I., González, J., and Capizzani, M. (2016). Which Categories and Brands to Promote With Targeted Coupons to Reward and to Develop Customers in Supermarkets. *Journal of Retailing*, 92(2):236–251.

- Palmatier, R. W. and Sridhar, S. (2017). *Marketing Strategy: Based on First Principles and Data Analytics*. Macmillan International Higher Education.
- Pancras, J. and Sudhir, K. (2007). Optimal Marketing Strategies for a Customer Data Intermediary. *Journal of Marketing Research*, 44(4):560–578.
- Park, J., Naumov, M., Basu, P., Deng, S., Kalaiyah, A., Khudia, D. S., Law, J., Malani, P., Malevich, A., Satish, N., Pino, J., Schatz, M., Sidorov, A., Sivakumar, V., Tulloch, A., Wang, X., Wu, Y., Yuen, H., Diril, U., Dzhulgakov, D., Hazelwood, K. M., Jia, B., Jia, Y., Qiao, L., Rao, V., Rotem, N., Yoo, S., and Smelyanskiy, M. (2018). Deep Learning Inference in Facebook Data Centers: Characterization, performance Optimizations and Hardware Implications. *CoRR*, abs/1811.09886.
- Pathak, S. (2017). Inside Walmart’s Advertising Blitz. <https://digiday.com/marketing/inside-walmarts-advertising-blitz> (Accessed 2019-02-20).
- Peppers, D. and Rogers, M. (1997). *The One to One Future: Building Relationships One Customer at a Time*. Currency-Doubleday.
- Peppers, D., Rogers, M., and Dorf, B. (1999). Is Your Company Ready for One-to-One Marketing. *Harvard Business Review*, 77(1):151–160.
- Raju, J. S. (1992). The Effect of Price Promotions on Variability in Product Category Sales. *Marketing Science*, 11(3):207–220.
- Rossi, P. E., McCulloch, R. E., and Allenby, G. M. (1996). The Value of Purchase History Data in Target Marketing. *Marketing Science*, 15(4):321–340.
- Rowley, J. (2005). Building Brand Webs: Customer Relationship Management Through the Tesco Clubcard Loyalty Scheme. *International Journal of Retail & Distribution Management*, 33(3):194–206.
- Ruiz, F. J. R., Athey, S., and Blei, D. M. (2018). SHOPPER: A Probabilistic Model of Consumer Choice With Substitutes and Complements. *CoRR*, abs/1711.03560.
- Russell, G. J. and Petersen, A. (2000). Analysis of Cross Category Dependence in Market Basket Selection. *Journal of Retailing*, 76(3):367–392.
- Rust, R. T. and Verhoef, P. C. (2005). Optimizing the Marketing Interventions Mix in Intermediate-Term CRM. *Marketing Science*, 24(3):477–489.
- Sahni, N. S., Zou, D., and Chintagunta, P. K. (2016). Do Targeted Discount Offers Serve as Advertising? Evidence From 70 Field Experiments. *Management Science*, 63(8):2688–2705.

- Schwarzer, G., Carpenter, J. R., and Rücker, G. (2015). *Meta-Analysis With R*. Springer.
- Seetharaman, P. and Chintagunta, P. K. (2003). The Proportional Hazard Model for Purchase Timing: A Comparison of Alternative Specifications. *Journal of Business & Economic Statistics*, 21(3):368–382.
- Shaffer, G. and Zhang, Z. J. (2002). Competitive One-to-One Promotions. *Management Science*, 48(9):1143–1160.
- Shah, D., Rust, R. T., Parasuraman, A., Staelin, R., and Day, G. S. (2006). The Path to Customer Centricity. *Journal of Service Research*, 9(2):113–124.
- Simester, D., Timoshenko, A., and Zoumpoulis, S. I. (2019a). Efficiently Evaluating Targeting Policies: Improving Upon Champions vs. Challenger Experiments. *Management Science* (forthcoming).
- Simester, D., Timoshenko, A., and Zoumpoulis, S. I. (2019b). Targeting Prospective Customers: Robustness of Machine Learning Methods to Typical Data Challenges. *Management Science* (forthcoming).
- Söderlund, M. and Colliander, J. (2015). Loyalty Program Rewards and Their Impact on Perceived Justice, Customer Satisfaction, and Repatronize Intentions. *Journal of Retailing and Consumer Services*, 25:47–57.
- Sudhir, K. (2016). The Exploration-Exploitation Tradeoff and Efficiency in Knowledge Production. *Marketing Science*, 35(1):1–9.
- Sutskever, I., Vinyals, O., and Le, Q. V. (2014). Sequence to Sequence Learning With Neural Networks. In *Advances in Neural Information Processing Systems*, pages 3104–3112.
- Taigman, Y., Yang, M., Ranzato, M., and Wolf, L. (2014). Deepface: Closing the Gap to Human-Level Performance in Face Verification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1701–1708.
- Tam, K. Y. and Ho, S. Y. (2006). Understanding the Impact of Web Personalization on User Information Processing and Decision Outcomes. *MIS Quarterly*, 30:865–890.
- Taylor, G. A. and Neslin, S. A. (2005). The Current and Future Sales Impact of a Retail Frequency Reward Program. *Journal of Retailing*, 81(4):293–305.
- Timoshenko, A. and Hauser, J. R. (2019). Identifying Customer Needs From User-Generated Content. *Marketing Science*, 38(1):1–20.

- Train, K. E. (2009). *Discrete Choice Methods With Simulation*. Cambridge University Press.
- Tucker, C. E. (2014). Social Networks, Personalized Advertising, and Privacy Controls. *Journal of Marketing Research*, 51(5):546–562.
- Valassis (2016). Cashing in on CPG Coupons. <https://www.slideshare.net/valassisce/valassis-cashingincpgcoupons> (Accessed 2017-08-17).
- Venkatesan, R. and Farris, P. W. (2012). Measuring and Managing Returns From Retailer-Customized Coupon Campaigns. *Journal of Marketing*, 76(1):76–94.
- Venkatesan, R. and Kumar, V. (2004). A Customer Lifetime Value Framework for Customer Selection and Resource Allocation Strategy. *Journal of Marketing*, 68(4):106–125.
- Verhoef, P. C. (2003). Understanding the Effect of Customer Relationship Management Efforts on Customer Retention and Customer Share Development. *Journal of Marketing*, 67(4):30–45.
- Vesanen, J. (2007). What Is Personalization? A Conceptual Framework. *European Journal of Marketing*, 41(5):409–418.
- Vinyals, O., Toshev, A., Bengio, S., and Erhan, D. (2015). Show and Tell: A Neural Image Caption Generator. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2015, Boston, MA, USA, June 7-12, 2015*, pages 3156–3164.
- Walmart (2005). Our Retail Divisions. http://corporate.walmart.com/_news_/news-archive/2005/01/07/our-retail-divisions (Accessed 2016-12-30).
- Walmart (2016). Walmart.com’s History and Mission. https://help.walmart.com/app/answers/detail/a_id/6 (Accessed 2016-12-30).
- Wedel, M. and Kamakura, W. A. (2012). *Market Segmentation: Conceptual and Methodological Foundations*. Springer Science & Business Media.
- Wedel, M. and Kannan, P. (2016). Marketing Analytics for Data-Rich Environments. *Journal of Marketing*, 80(6):97–121.
- Winer, R. S. and Neslin, S. A. (2014). *The History of Marketing Science*. World Scientific.
- Xu, B., Wang, N., Chen, T., and Li, M. (2015). Empirical Evaluation of Rectified Activations in Convolutional Network. *arXiv Preprint arXiv:1505.00853*.

- Yi, Y. and Jeon, H. (2003). Effects of Loyalty Programs on Value Perception, Program Loyalty, and Brand Loyalty. *Journal of the Academy of Marketing Science*, 31(3):229–240.
- Zanutto, E. L. and Bradlow, E. T. (2006). Data Pruning in Consumer Choice Models. *Quantitative Marketing and Economics*, 4(3):267–287.
- Zhang, J. and Breugelmans, E. (2012). The Impact of an Item-Based Loyalty Program on Consumer Purchase Behavior. *Journal of Marketing Research*, 49(1):50–65.
- Zhang, J. and Wedel, M. (2009). The Effectiveness of Customized Promotions in Online and Offline Stores. *Journal of Marketing Research*, 46(2):190–206.
- Zhang, M. and Luo, L. (2018). Can User Generated Content Predict Restaurant Survival: Deep Learning of Yelp Photos and Reviews. *Available at SSRN: <https://ssrn.com/abstract=3108288>*.

Selbstständigkeitserklärung

Ich versichere, die von mir vorgelegte Dissertation selbständig und ohne unerlaubte Hilfe und Hilfsmittel angefertigt, sowie die benutzten Quellen und Daten anderen Ursprungs als solche kenntlich gemacht zu haben.

Ich bezeuge durch meine Unterschrift, dass meine Angaben über die bei der Abfassung meiner Dissertation benutzten Hilfsmittel, über die mir zuteil gewordene Hilfe sowie über frühere Begutachtungen meiner Dissertation in jeder Hinsicht der Wahrheit entsprechen.

Berlin, 26. Mai 2019

Sebastian Gabel