

THOMAS W. FRICK

The Implications of Advertising Personalization for Firms, Consumers, and Ad Platforms



THE IMPLICATIONS OF ADVERTISING
PERSONALIZATION FOR FIRMS,
CONSUMERS, AND AD PLATFORMS

The Implications of Advertising Personalization for Firms, Consumers, and Ad Platforms

De implicaties van advertentie personalisatie voor bedrijven,
consumenten en advertentieplatforms

Thesis

to obtain the degree of Doctor from the
Erasmus University Rotterdam
by command of the
rector magnificus

Prof.dr. R.C.M.E. Engels

and in accordance with the decision of the Doctorate Board

The public defense shall be held on

Friday, 21st of September 2018 at 11:30

by

THOMAS WALTER FRICK
born in Hechingen, Germany.

Erasmus University Rotterdam

The logo of Erasmus University Rotterdam, featuring the word "Erasmus" in a stylized, cursive script.

Doctoral Committee

Promotors: Prof.dr. T. Li
Prof.dr.ir. H.W.G.M. van Heck

Other members: Prof.dr. P.A. Pavlou
Dr. J.M.T. Roos
Prof.dr. R. Telang

Erasmus Research Institute of Management - ERIM

The joint research institute of the Rotterdam School of Management (RSM)
and the Erasmus School of Economics (ESE) at the Erasmus University Rotterdam
Internet: <http://www.irim.eur.nl>

ERIM Electronic Series Portal: <http://repub.eur.nl/>

ERIM PhD Series in Research in Management, 452

ERIM reference number: EPS-2018-452-LIS

ISBN 978-90-5892-522-0

©2018, Thomas Walter Frick

Design: PanArt www.panart.nl

Cover image design: Simone Lilmoes www.simonelilmoes.com

This publication (cover and interior) is printed by Tuijtel on recycled paper, BalanceSilk®.

The ink used is produced from renewable resources and alcohol free fountain solution.

Certifications for the paper and the printing production process: Recycle, EU Ecolabel, FSC®, ISO14001.

More info: www.tuijtel.com

All rights reserved. No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or by any information storage and retrieval system, without permission in writing from the author.



Acknowledgements

After pretty much six years that I spent on my Master's and PhD, my time in Rotterdam comes to an end. The loose idea of pursuing an academic career had already come up during my Bachelor's. Especially my experiences during my collaborations with Prof Tobias Kretschmer and the members of his department at LMU Munich and my exchange stay at University of Sydney, had nurtured the seed in my head that being an academic might not be that bad after all. But it was only after having completed my Master's in Business Information Management with a, back then, far too ambitious Master's thesis that I decided to go into a PhD program. Now, I have completed this trajectory and there are a lot of people I want to thank for their support along the way.

First of all, thank you Ting for being my supervisor. Thank you for giving me a lot of freedom that helped me to develop into a more independent researcher. I always felt you did not shy away from interacting with me as a colleague instead of a subordinate. Eric, thank you for being there for me whenever I needed support. I have learned a lot from you when it comes to how to survive in academia. Thank you for coaching me with my first university teaching assignments.

There are several other academics that influenced my academic journey significantly and I am glad to have some of them as members of my PhD committee. Rahul, I cannot thank you enough for hosting me in Pittsburgh in the beginning of 2017. Working with you was in retrospect probably the determining factor that made me decide to pursue an academic career after my PhD. I have learned a tremendous amount from you in such a short time and can only hope that there will be further collaborations in the future to learn even more. Paul, I was truly amazed by your style of collaboration, your extreme kindness, and your academic skills. It is a pleasure and inspiring to work with such an established academic that is so welcoming, friendly and dedicated.

I am extremely thankful to the members of my PhD committee for their efforts and support. Jason, thank you for acting as secretary of my PhD committee and providing the marketing perspective on my dissertation. Your feedback was very helpful in

improving this work. I also want to thank Prof Jan Damsgaard for being a part of my PhD committee (and for giving me a job at Copenhagen Business School). Thanks to Prof Benedict Daellert for being a member of my PhD committee.

Two people that supported me tremendously during my PhD project are Maarten and Robert. Thank you for giving me access to actual business data, being true experts in what you are doing, and maybe most importantly, for becoming friends along the way.

My gratitude also goes to a lot of colleagues at Rotterdam School of Management. Dimitris, you have been following my journey already as my Master's thesis coach and have become a great friend over time. I will always remember all the fun things we experienced together during our conference travels. I really enjoyed having somebody with a 'normal' music taste around. Rodrigo, thanks for being both the go to guy for academic questions as well as for getting a beer in the 'Smitse'. Clint, thanks for being a great office mate, having plenty of discussions about academic life, and for being my paranimf. Marcel, thanks for coming back from your US visit just in time to act as my paranimf. I hope that we will share some more great memories in the future.

There are plenty of other colleagues to thank and I am running the risk of forgetting some of them. Thank you Mark, Zike, Tobias, Yashar, Otto, Wolf, Ksenia, and Jeffrey for being great colleagues and making the department an amazing place to be around. But how much fun could the life at the Technology and Operations Management department have been without all of the great PhD colleagues. Derck, May, Camill, Timo, Mohammad, Francesco, Davide, Ivo, Micha, Christina, Wouter, Paul, Nick, Sarita, Konstantina, Markus - thank you for making the whole process so much more fun. Very special thanks go to Ingrid, Cheryl, Carmen, and Mirjam for keeping things together in the department. Also a big shout out to the many great PhD colleagues at Carnegie Mellon University that made my research visit so much more fun.

Special thanks go to the staff of the Erasmus Research Institute of Management (ERIM) for their support throughout my PhD trajectory.

It is difficult to put the gratitude I have for my parents and how they have supported me over all these years into words. Mama und Papa, danke für all die Unterstützung egal in welche Richtung mein Leben sich über die Jahre entwickelt hat. Ihr wart immer für mich da.

Without one very special person this endeavor would have probably not been possible. Thank you Simone for being my pillar, for giving me a different perspective, for sharing happy moments and for helping me through difficult times. Honestly, I could not have done it without you.

Thomas W. Frick
Rotterdam, July 2018

Contents

1	Introduction	1
1.1	Advertising Personalization	3
1.2	Advertising Performance Measurement	3
1.3	Research Motivation	3
1.4	Practical Relevance	6
1.5	Dissertation Overview	7
1.6	Declaration of Contributions	10
2	Personalization Specificity in Social Retargeting:	
	A Field Experiment	13
2.1	Introduction	13
2.2	Theory	17
2.2.1	Category-Specific Ad Personalization	18
2.2.2	Product-Specific Ad Personalization	19
2.2.3	Social Targeting	20
2.2.4	Personalization Specificity and Social Targeting	22
2.3	Field Experiment	24
2.4	Analysis and Results	26
2.4.1	Likelihood to Click	28
2.4.2	Likelihood to Purchase	30
2.4.3	Robustness Checks	32
2.5	Additional Analysis	33
2.5.1	Temporal Targeting and Preference Development	33
2.5.2	Browsing Depth and Preference Development	36
2.5.3	Consumer Demographics	38

2.6	Discussion	40
2.6.1	Theoretical Contributions	42
2.6.2	Practical Implications	43
2.6.3	Limitations and Future Research	44
Appendices		47
A2.1	Social Actions per Ad Type	47
A2.2	Economic Implications of Product-Specific Personalization and Social Targeting	48
A2.3	Data Structure Experiment	48
A2.3.1	Campaign Set-Up	48
A2.3.2	Facebook’s Breakdown Functionality	49
A2.3.3	Facebook’s Reporting Structure	49
A2.4	Consumer Response Likelihood Analysis on Consumer-Level	50
A2.5	Likelihood to Click Estimates Using Unique Clicks	51
A2.6	Alternative Estimators for Click and Purchase Probabilities	52
A2.7	Estimates for Post-View Purchase Probability	53
A2.8	Robustness Check: Average Reach of Ad Attribute Combination	53
3	Social Influence and Visual Attention in the Personalization Privacy Paradox: An Eye Tracking Study	57
3.1	Introduction	57
3.2	Related Literature	62
3.3	Theory	63
3.3.1	Advertising Personalization	63
3.3.2	Personalization Privacy Paradox	64
3.3.3	Informational Social Influence	66
3.3.4	Informational Social Influence as Mitigation of Privacy Concerns	66
3.3.5	Consumer Attention	68
3.4	Methodology	71
3.4.1	Pre-Test	72
3.4.2	Experimental Procedure	73
3.4.3	Sample	75
3.4.4	Manipulation Checks	76
3.4.5	Measures	76
3.5	Analysis	77
3.6	Discussion	82
3.6.1	Practical Implications	85
3.6.2	Limitations	85
Appendices		88
A3.1	Survey Measures Used in Pre-Test	88
A3.2	Survey Measures Used in Scenario-Based Experiment	89

A3.3 Overview of Variables Used in Analysis and Data Collection Method . . .	90
A3.4 Logit Regression for Click Probability	91
A3.5 OLS Regression for Privacy Concerns	92
A3.6 Serial Mediation Model	92
4 Pay For What You Get - Incentive Misalignments in Programmatic Advertising: Evidence From A Randomized Field Experiment	95
4.1 Introduction	95
4.2 Related Literature	100
4.2.1 Behaviorally Targeted Advertising	100
4.2.2 Ad Effectiveness	101
4.2.3 Ad Auctions	103
4.3 Theory	103
4.4 Ad Allocation Process and Infrastructure	108
4.5 Payment and Pricing for Ad Impressions	110
4.6 Experiment & Analysis	110
4.6.1 Bidding & Increase in Purchase Probability	113
4.6.2 Bidding & Profit Contribution	117
4.7 Theoretical Contributions	119
4.8 Practical Implications	121
4.9 Limitations	121
Appendices	124
A4.1 Bid Prediction with Consumer Characteristics & Behavioral Data . . .	124
A4.2 Share of Consumers Allocated to Treatment and Control Group	125
A4.3 Estimation of Purchase Probability Using Alternative Measures for Bid	126
A4.4 Prediction of Purchase Probabilities	128
A4.5 Alternative Cost Measures	130
A4.6 Revenue Regression	132
A4.7 Controlling for the Effect of Competition for Ad Impressions	135
5 Conclusions	139
5.1 Synthesis of Findings	140
5.2 Academic Contributions	145
5.3 Practical Relevance	147
5.4 Limitations	148
5.5 Future Research	149
References	153

Back Matter	165
English Summary	165
Nederlandse Samenvatting	167
About the Author	169
Author Portfolio	170
The ERIM PhD Series	173

Introduction

Consumers are confronted with a plethora of advertising in their daily lives. From more traditional advertising types such as print ads in magazines, billboard ads on their way to work, and TV ads in between their favorite show, to the relatively new forms of digital advertising. While traditional forms of advertising have decreased in their relative market share over the past years, digital advertising has increased its market share and keeps growing at a two digit rate (Liu, 2016). Since the first digital advertising campaign conducted by AT&T on the website HotWired in 1994, digital advertising has enjoyed enormous growth and is now the fastest growing type of advertising (McStay, 2010). Worldwide digital advertising spending is predicted to surpass \$200 billion in 2017 (McNair, 2017).

There are several reasons for the tremendous success of digital advertising. While it remains elusive to offer a complete list of success factors, we want to present some of the dominant aspects that paved the route of success for digital ads.

Notably, consumers spend an exceeding amount of their time online. The average US adult spends 5 hours and 53 minutes per day with internet connected devices (EMarketer, 2017b). With more than half of the planet's population predicted to use the internet in 2019 (EMarketer, 2017a), an enormous amount of consumers can be reached online. These trends in consumer behavior are being monitored by firms that want to reach consumers with their product offerings. Digital advertising represents the right media to address consumers in online environments where they spend an increasing amount of their time.

Compared to traditional advertising, e.g. large billboards, digital advertising is less bound in space (McStay, 2010). While traditional advertising largely relies on consumers' ability to recall advertising that might become relevant only later - imagine an ad for a chocolate bar that can only influence a consumer's purchase decision when entering a supermarket - digital advertising can influence consumers more directly. Hyperlinks that are connected to digital ads allow firms to steer consumers to websites on which they offer consumers more information on their products and the opportunity to purchase directly. Digital advertising allows firms to have a more immediate influence on consumers' purchase decisions.

The major operational aspects that separate digital advertising from traditional forms of advertising are the potential to use fine grained individual-level consumer data to adjust, optimize, and monitor digital advertising as well as the ability to measure its performance. The technological capability to dynamically adjust digital advertising to the preferences of consumers is called *Advertising Personalization* (Bleier and Eisenbeiss, 2015a; Lambrecht and Tucker, 2013). Advertising personalization offers firms a huge potential, as it allows them to address consumers with more relevant ads that trigger more positive consumer responses to these ads.

While firms have identified personalization as a way to improve their interaction with consumers, only few firms consider their personalization strategies to be in advanced stages (McCarthy, 2017). Consumers' privacy concerns regarding the use of their personal data for personalization purposes increases the difficulty for firms to define their personalization strategies, especially in an advertising context (Awad and Krishnan, 2006; Sutanto et al., 2013).

Next to the challenges related to defining and implementing an ad personalization strategy, firms aim to justify their investments in advertising by measuring its return. *Advertising Performance Measurement* aims to capture the impact that advertising has on firms' financial performance. Although ad platforms, that offer firms the technology to address consumers with digital ads, report ad performance measures, research has started to question whether these performance measures are a good representation of firms' return on ad spend (Johnson et al., 2017a).

The ability to personalize ads and the ability to measure ad performance have been largely enabled by developments in information and communication technologies (ICT) over the last years. With this dissertation, we aim to help firms to overcome challenges related to (1) **Advertising Personalization** and (2) **Advertising Performance Measurement**.

1.1 Advertising Personalization

With the availability of individual-level data on consumers' characteristics and behavior, firms gained the ability to predict consumers' preferences and personalize digital ads for individual consumers. Ad personalization is defined as the firm-initiated adjustment of ad content towards the preferences of consumers (Arora et al., 2008). Digital advertising allows marketers to personalize advertisements for individual consumers using, among others, information about consumers' demographics, interests, location, browsing behavior, and even social connections. Generally, several information stimuli, including offline and online organic content as well as advertising, are competing for consumers' attention simultaneously. One major way for firms to differentiate from other content providers and advertisers that compete for consumers' attention, is to increase ad relevance through ad personalization.

1.2 Advertising Performance Measurement

Next to the opportunities to increase ad relevance by adjusting ad content to the individual preferences of consumers, the ability to track consumers' reactions to digital ads allows firms to actually measure to what extent marketing campaign objectives are met. Digital advertising has been praised as being much more measurable than traditional advertising. With specific data on variables such as how many consumers have been confronted with ads, clicked on them, and conducted a purchase after having seen an ad, firms become able to assess the return on investments in advertising. Advertising performance measurement allows firms to assess their return on investments in advertising and how to allocate their marketing budgets in a smarter fashion.

1.3 Research Motivation

While technically, firms can make use of ICT to personalize their ads and measure consumers' behavioral responses to ads, there are several theoretical questions and practical challenges that remain unanswered. With this dissertation, we aim to contribute to the literature on advertising personalization and the economics of digital advertising. This dissertation is guided by the research question:

How do firms' advertising personalization strategies affect consumer responses and how can these consumer responses be assessed using ad platforms' reporting systems?

Industry reports point out that while marketers classify personalization as the most important marketing capability, they simultaneously perceive personalization as the biggest challenge within their organizations (Adobe Systems Inc., 2014). The potential to increase advertising relevance through ad personalization depends on how accurately firms can predict consumers' preferences. A non-accurate prediction of consumer preferences, where consumers are addressed with ads featuring products that they do not like, has been shown to lead to consumer resistance and annoyance (Arora et al., 2008). Previous research presents inconclusive findings with respect to the question to what extent advertising should be personalized. Some work argues generic ads outperform more personalized dynamic retargeting ads that are personalized based on consumers' browsing behavior (Lambrecht and Tucker, 2013). Only when consumers have construed preferences, meaning that they have narrowed down their preferences to a specific product and are close to making a purchase decision (Simonson, 2005), a higher level of ad personalization leads to more positive consumer responses to ads. Contradicting these findings, other work in the area of personalized advertising claims that ads that apply a higher degree of ad personalization lead to more positive consumer responses than less personalized ads (Bleier and Eisenbeiss, 2015a). For both researchers and practitioners it remains challenging to unite such contradictory findings.

Personalization of advertising hinges on the availability of consumer-level data that can be used to generate information regarding consumers' preferences. To achieve accurate preference predictions, especially for higher levels of ad personalization, firms need to have access to such data. While technically, data on consumers' characteristics and online behaviors can be recorded, consumers tend to be concerned about their data being used for advertising purposes (Sutanto et al., 2013). Although personalization of digital advertising increases ad relevance for consumers, it simultaneously triggers consumer privacy concerns. This phenomenon is described as personalization privacy paradox in the information systems literature (Awad and Krishnan, 2006; Malheiros et al., 2012; Sutanto et al., 2013). Data that is required for personalization of ads might be data that consumers are not willing to share deliberately. Recent policy changes implemented by the European Union strengthen and underline consumers' right to privacy. The General Data Protection Regulation (GDPR) establishes clear

standards that define how firms have to manage consumers' data and how firms can process this data (EUR-Lex, 2016). Consumers' perception of inappropriate use of their personal information for personalization purposes leads to so called personalization reactance (White et al., 2008). Research has shown that consumers have an interest in limiting third parties' access to their personal information (Utz and Kramer, 2009). Especially in advertising contexts, where consumers perceive personalization as less beneficial compared to personalization of other services, consumers are more privacy sensitive when it comes to the use of their personal data (Awad and Krishnan, 2006; Sutanto et al., 2013). It remains challenging for firms to adequately balance personalization gains with consumer privacy concerns. Research can help firms to develop strategies that take the personalization privacy paradox into account in an advertising personalization context.

To assess which advertising personalization strategy is most beneficial for firms and most acceptable for consumers, firms need be able to measure consumers' responses to ads accurately. Simultaneously, managers need to be able to interpret ad measurement reports correctly to draw the right conclusions for business strategy. Essentially, firms need to assess whether their investments in advertising pay off. While, arguably, in a first step we need to increase managers' understanding of how to evaluate advertising performance, the identification of the economic return on investments in advertising remains challenging (Dalessandro et al., 2012). Despite the claim that digital advertising allows firms to measure advertising performance, most digital advertising contexts do not allow firms to answer the fundamental economic question: How much additional profit did my firm generate when advertising compared to when not serving advertising to consumers? Firms struggle with implementing well-designed experiments to identify the economic effect of ads both because of the methodological challenge but also due to technical limitations. Work in the area of economics of advertising has suggested a design for an information system that would allow ad platforms, that handle the buying of ad impressions on behalf of firms, to correctly identify the return on ad investments and report this metric to firms (Johnson et al., 2017a). Until now, it remains questionable whether ad platforms are willing to carry the costs of implementing such information systems. Current ad reporting standards allow ad platforms to report inflated performance measures. Reporting the actual economic effect of advertising might decrease marketers' perception of ad performance and simultaneously their willingness to invest in advertising leading to a revenue decrease for ad platforms. Furthermore, related work has started to question whether contracts between firms and ad platforms, to which firms outsource the bidding

for ad impressions, are favorable for firms (Johnson and Lewis, 2015). The practical implications of incentives in these outsourcing contracts need to be assessed empirically, to offer a better understanding whether current contract specifications are harmful to firms.

1.4 Practical Relevance

Renting out advertising space has become the main source of revenue for Internet companies like Google, Facebook, Yahoo, and TripAdvisor. Such a revenue model requires companies to attract high volumes of traffic with their services that makes it attractive for firms to address consumers with ads on their platforms. Most social media platforms are financed by this revenue model called advertising model (Schumann, 2014). The massive rise in the number of social media users, the high demand of reliability and availability of services, as well as the urge to ever develop new applications for the platforms to keep users interested, has created financial pressure on social media platform providers. This financial pressure demands to work under a financially feasible business model in which these services, that are mostly expected to be free, can be provided for no charge. While services are provided to users free of charge, users are confronted with advertising that generates revenue for platform providers by renting out advertising space on their websites.

When implementing an advertising revenue model, platform providers face the constant need to satisfy both advertisers as well as consumers. This is the case as online platforms benefit from both positive direct and indirect network externalities. Positive direct network externalities describe that users utility for a service or product increases in the number of users on their side of the market (Katz and Shapiro, 1994). Positive indirect network externalities describe that users of a product or service on one side of the market benefit from an increase in the number of users on the other side of the market (Katz and Shapiro, 1985; Liebowitz and Margolis, 1994). Users value online platforms, e.g. social network sites, higher in case there are more users present that they can communicate with (direct network externality). At the same time, advertisers value an online platform higher if there are more consumers that they can advertise to (indirect network externality) (Clements, 2004).

Despite the positive indirect network externalities for advertisers with an increase in the number of users on an online platform, consumers tend to prefer advertising free environments and try to avoid advertising (Drèze and Hussherr, 2003; Gal-Or

et al., 2018). This conflict is amplified by high demands on the return on ad spend from advertisers, leading to platforms introducing additional and novel ways to individually address consumers with more relevant advertising. One way to accommodate advertisers' interest to serve more relevant ads is to offer opportunities for advertising personalization. At the same time, advertising personalization leads to an increase in consumers' concerns about the utilization of their personal data (Sutanto et al., 2013). Therefore, platforms are constantly struggling with balancing the interests of both users and advertisers while operating under a financially feasible business model.

Our research is of significant interest to the major stakeholders on the demand side of digital advertising: firms, consumers, and ad platforms. We shed light on which advertising personalization strategy benefits firms and how personalization of ads affects consumers' concerns regarding the use of their personal information. These insights help ad platforms to balance advertisers' and consumers' interests. Further, we investigate the economic value of digital advertising and whether ad platforms optimize the ad allocation process in the economic interest of firms or solely in their own interest.

1.5 Dissertation Overview

We conceptualize the context of advertising personalization, more specifically the demand side of advertising personalization, as the relationship between three major stakeholders (see Figure 1.1). A *firm* has an interest to serve ads to a *consumer* that can be addressed with digital advertising via an *ad platform* that is mediating the relationship between *firm* and *consumer*. In the different chapters of this dissertation, we focus on different aspects of the relationships between what we consider the three main stakeholder on the demand side of digital advertising.

This dissertation is structured as follows. In Chapter 2, we investigate, with the help of a field experiment, how specific advertising personalization should be and whether consumers should be socially targeted in personalized advertising. Chapter 3 zooms into how consumers perceive personalized digital ads, allocate their attention, and respond to personalized ads given that their personal information was used which triggers consumer privacy concerns. Chapter 4 focuses on the empirical assessment of the economic relationship between firms and ad platforms and whether this relationship is governed by a contract that leads to a beneficial outcome

for firms. Below we present the abstracts of the three main chapters in this dissertation.

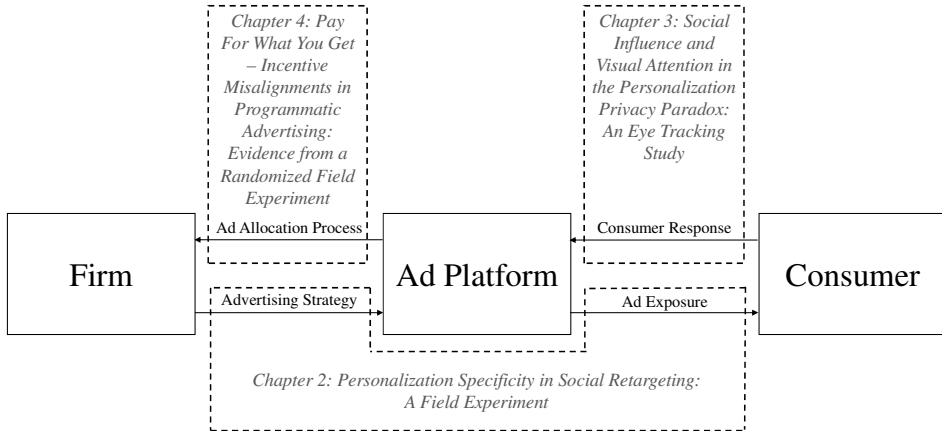


Figure 1.1: Dissertation Overview

Chapter 2 - Abstract This study investigates the effectiveness of personalization specificity and social targeting in the context of social retargeting. Social retargeting combines the features of social advertising, in which consumers are targeted based on social connections, and retargeting, for which consumers' browsing behavior is used to personalize ad content to consumers' preferences. We compare consumers' responses to product- and category-specific advertising personalization in a large-scale randomized field experiment in collaboration with a major e-retailer and assess the impact of socially targeting consumers in the context of personalized advertising. Contradicting prior empirical findings, our results indicate that product-specific ads generally outperform less personalized category-specific ads. While theory suggests a positive effect, we find that social targeting leads to less positive consumer responses to personalized ads. Further, socially targeted consumers are not more responsive to more personalized product-specific ads. We show that our results remain robust and driven by ad personalization when controlling for temporal targeting, how deep consumers browse the e-retailer's website, and consumer characteristics. Our study contributes to the IS and marketing literature related to personalization in digital advertising and provides valuable suggestions for firms' personalization strategies.

Chapter 3 - Abstract The personalization privacy paradox suggests that the personalization of advertising increases ad relevance but simultaneously triggers privacy concerns as firms make use of consumers' information. We combine a lab experiment with eye tracking and survey methodology to investigate the role of informational social influence and attention in the personalization privacy paradox for digital advertising. We find that informational social influence increases consumers' likelihood to click on ads but does not reduce consumer privacy concerns originating in personalization. Our findings contradict the presence of a negativity bias directing consumers' attention to negatively perceived stimuli. We show that privacy concerns decrease consumers' attention towards personalized ads, subsequently leading to a decrease in ad clicks and supporting a positive role of visual attention for advertising performance. By objectively measuring visual attention, we obtain a richer understanding of how consumers process information and make decisions. We show that privacy concerns, triggered by personalization, negatively influence ad performance through a decrease in attention towards ads.

Chapter 4 - Abstract In programmatic digital advertising, firms outsource the bidding for ad impressions to ad platforms. We theoretically assess the contracts governing this outsourcing relationship and find evidence for a potential incentive misalignment. Based on the contract structure, advertising platforms have an incentive to target consumers with higher inherent purchase probabilities independent of the effect of ads on consumers' purchase probabilities. Nevertheless, the implications of such an incentive structure for the firm are not straightforward and depend on both the ad platform's actual behavior and the correlation between absolute and incremental purchase probabilities. With the help of a large-scale randomized field experiment, addressing 20,918 individual consumers with ads, we empirically investigate the implications of the bidding optimization deployed by the ad platform for the firm. Our unique data set allows us to both causally assess the impact of ads on consumers' purchase probabilities and whether this impact is heterogeneous depending on the bids placed for consumers' ad impressions. In accordance with incentives specified in contracts between firm and ad platform, we find that ad platforms target consumers that are more likely to purchase independent of the effect of ads on their purchase probability. We find no significant correlation between the inherent purchase probability of consumers and the increase in purchase probabilities caused by ads. More expensive ads do not have a higher impact on consumers' purchase probabilities. Therefore, ad platforms bidding optimization does not align with the economic interest of firms.

Firms try to adjust their willingness to pay for purchases reported by the ad platform to match the platform's actual success contribution. Nevertheless, this adjustment remains without effect as it does have no influence on the incentive structure in the outsourcing contract. Advertising platforms' increasing capabilities to use large amounts of individual level data to predict consumers' inherent purchase probabilities increase the severity of this issue and emphasize the empirically confirmed incentive misalignment.

1.6 Declaration of Contributions

The majority of this work has been conducted independently by the author. More precisely, the author was responsible for defining the research questions and scope of the research, integrating the work with related literature, analyzing data, interpretation of results, and writing the chapters included in this dissertation. Nonetheless, this work benefited from discussions with co-authors that helped to trigger a process of critical thinking and improvements of the chapters.

Chapter 1: This chapter was independently written by the author.

Chapter 2: This chapter is joint work with Prof. T. Li. The majority of this work has been conducted independently by the author. While Prof. T. Li supported the author via discussions with the final definition of the research question, the author conducted most of the work for this chapter. This included the identification of related literature and theoretical relevance, convincing a partner firm to conduct the field experiment, setting up the field experiment, collection all relevant data, analyzing the data, as well as consolidating all relevant findings and writing the chapter. The co-author helped to improve the work along the way with suggestions for improvements.

Chapter 3: This chapter is joint work with Prof. T. Li, and Prof. P.A. Pavlou. For this work the co-authors provided valuable feedback regarding the design and execution of the lab experiment as well as help to improve the presentation of the contributions of the work. The majority of the chapter including the design of the experiment, the implementation of the experiment including the set up of the eye tracking device, programming of personalized websites, and design of the experimental procedure, data extraction, analysis, definition of theoretical and practical contributions, and

writing of the chapter was done by the author. The work benefited from the help of two student assistants that led participants one-by-one through the lab experiment making use of ERIM's Research Participation Pool (ERPS).

Chapter 4: This chapter is joint work with Prof. R. Telang and Dr. R. Belo. The majority of the conceptual work of this paper was conducted during the author's research visit to the Heinz College, Carnegie Mellon University Pittsburgh. During this time, Prof. R. Telang provided feedback in regular sessions that allowed the author to improve the work significantly. The majority of the work including field data collection, data analysis, identification of related literature, definition of theoretical and practical contributions, consolidation of findings, and writing the chapter was conducted by the author. Both co-authors provided valuable feedback and inputs that significantly improved the chapter. This work benefited from the financial support of the Vereniging Trustfonds Erasmus Universiteit Rotterdam that partially funded the authors research visit to Carnegie Mellon University Pittsburgh.

Chapter 5: This chapter was independently written by the author.

Personalization Specificity in Social Retargeting: A Field Experiment¹

2.1 Introduction

Worldwide digital advertising spending is predicted to surpass \$200 billion in 2017 (McNair, 2017). This number indicates that countless firms, in addition to digital content providers, are competing for consumers' attention with digital advertising online. One major way to differentiate from other firms competing for consumers' attention is to increase ad relevance through ad personalization (Arora et al., 2008). In advertising personalization firms adjust their ad content to consumers' preferences with the aim to positively influence consumer responses to ads. In a study by Adobe, marketers named marketing personalization as the most important marketing capability while being the biggest challenge within their organizations (Adobe Systems Inc., 2014). Although, most firms are in the process of implementing personalization strategies

¹Earlier versions of this study appeared in the following conference proceedings or were presented at the below mentioned conferences and workshops:

- Frick, T.W. & Li, T. (2016). Personalization in Social Retargeting: A Field Experiment. In *International Conference on Information Systems (ICIS)*, Dublin, Ireland.
- Frick, T.W. & Li, T. (2016). Social Retargeting: A Field Experiment. In *Statistical Challenges in eCommerce Research Symposium (SCECR)*, Naxos, Greece.
- Frick, T.W. & Li, T. (2016). Social Retargeting: A Field Experiment. In *The Economics of Information and Communication Technologies, ZEW Conference, Centre for European Economics Research*, Mannheim, Germany.
- Frick, T.W. & Li, T. (2015). Social Retargeting: A Randomized Field Experiment. In *37th ISMS Marketing Science Conference*, Baltimore, US.

and have recognized their potential value, only 6% of firms consider themselves in advanced stages of implementing their personalization strategy (McCarthy, 2017).

Advertising personalization relies on the availability of consumer data to personalize ad content. One online space where this consumer data, such as demographics and interests, is available is social networking sites. In 2015, the biggest social networking site, Facebook, introduced the functionality to dynamically retarget consumers by making use of their external browsing behavior to personalize advertising. The objective of this paper is to examine the effects of this new form of advertising called *Social Retargeting*, which combines the features of retargeting and social advertising.

In *retargeting*, consumers' browsing behavior is used to infer their preferences and target them with personalized ads on external websites (Bleier and Eisenbeiss, 2015a; Lambrecht and Tucker, 2013). While research generally finds that personalized communication with consumers has positive effects on consumer-firm interactions (Bleier and Eisenbeiss, 2015a; Tam and Ho, 2005; Hoffman and Novak, 1996; Komiak and Benbasat, 2006), it remains unclear how specific ad personalization should be. Arora et al. (2008) point out the costs and benefits of different levels of personalization. Less specific personalization requires less detailed information on consumers and reduces the risk of preference misclassification which leads to negative responses to ads. More specific personalization allows firms to match consumers' preferences more closely, and decrease consumers' search costs through more specific recommendations.

In *social advertising*, advertisers target consumers by using their social connections to infer their preferences and make these social connections explicit in the ad text to foster consumers' identification with the advertised product (Tucker, 2016). We define the combination of these advertising techniques, using consumers' underlying social networks to target them with ads and making these social connections explicit in the ad text, as *Social Targeting*. Prior research suggests that social targeting has a positive effect on consumer responses to ads through homophily of connected users and informational social influence (Bakshy et al., 2012). Recent work, however, points towards the tendency of consumers not to comply with informational social influence from their peers when signaling their identity on social networking sites (Sun et al., 2017). Although consumers want to identify with a favorable social group, they simultaneously strive for uniqueness (Chan et al., 2012).

Until now, it remains unclear how socially targeted consumers respond to personalized ads which create the perception of unique recommendations for consumers. Both personalization and social targeting have been shown to individually lead to positive consumer responses. However, a combination of personalization and social targeting

has not been investigated so far. In this paper, we answer the following research questions: (1) How do consumers respond to different levels of ad personalization specificity? (2) Do consumers that are socially targeted respond differently to different levels of personalization specificity?

To answer these research questions, we conducted a large-scale randomized field experiment in collaboration with a major European e-retailer on the advertising platform of Facebook. We randomly assigned 198,234 individual consumers to one of two types of ads with different levels of ad personalization specificity: (1) In the *category-specific* (low specificity) ad personalization condition, consumers see ads that advertise a product category they had previously visited. (2) In the *product-specific* (high specificity) ad personalization condition, consumers see ads that advertise a specific product matching their browsing behavior. We measure and analyze consumer responses to these ads by recording whether consumers click on an ad and/or make a purchase.

We find that consumers respond more positively to product-specific than category-specific ads regarding both click and purchase probabilities. With respect to social targeting, surprisingly, we find that socially targeted consumers, who according to previous studies would respond more positively to ads (Bakshy et al., 2012; Tucker, 2016), are in fact less likely to click on personalized ads and/or make a purchase after being exposed to personalized ads. Perhaps most interestingly, we find that social targeting does not lead to higher consumer acceptance for more specific ad personalization. On the contrary, social targeting decreases consumers' probability to click on more personalized ads, suggesting a conflict between more specific ad personalization and social targeting.

We use the uniqueness theory (Chan et al., 2012) to explain why social targeting has a negative effect on consumer responses to social retargeting. Consumers receive two different, and likely conflicting, information signals from socially retargeted ads. On the one hand, the retargeted ad is uniquely personalized for a consumer. On the other hand, however, the ad is shown with a consumer's friend endorsements in ad texts, suggesting similarities rather than uniqueness. The inclusion of social identities, friends' names, depersonalizes the ad that was meant for the individual receiver. Thus the reduction in ad effectiveness happens through a decrease in the perceived personalization. Consumers strive for uniqueness and wish to be different from their peers where the feeling of being too similar leads to emotional reactance (Berger and Heath, 2008). Our results suggest that the conflict between a uniquely personalized ad and social identities is stronger for product-specific ads (compared

to category-specific ads), further supporting our argument. The conflict between a uniquely personalized product ad and social identities is stronger for more specific ad personalization.

Our results remain robust when considering consumers' preference development, more detailed information on consumers' browsing behavior, and demographic information on consumers. The effects we observe are also economically significant. Product-specific personalization leads to an increase in click probability of 120% and a 214% increase in purchase probability compared to category-specific personalization. Socially targeted consumers have a 13% lower click probability and a 62% lower purchase probability than non-socially targeted consumers.

Our study contributes to the literature in the area of advertising personalization in several ways. First, we contribute to the discussion of adequate levels of advertising personalization by investigating the effects of personalization specificity. We focus on the question of how specific personalization, in terms of recommending a category or product, should be. Previous research found that generic brand ads outperform dynamically retargeted ads (Lambrecht and Tucker, 2013). Our results challenge this empirical finding by showing that consumers react more positively to more specific ad personalization. We attribute the difference in findings to our cleaner experimental design. In previous work, dynamically retargeted ads displayed several products, potentially confounding the experimental treatment with differences in visual attractiveness of ads and effects originating in the composition of choice sets presented to consumers. A visually less attractive ad design or consumers' difficulty in deciding which product to click might decrease ad performance of dynamically retargeted ads compared to generic brand ads. For our study, we carefully designed our experimental ad treatments to make sure we isolate the effects of category- and product-specific ad personalization and identify their effects on consumer responses in a cleaner fashion.

In addition, our results provide evidence that the effectiveness of more specific ad personalization decreases slower as the time between a consumer's website visit and her ad exposure increases. This finding contradicts prior findings that suggest the opposite effect (Bleier and Eisenbeiss, 2015a). Such difference might originate in the fact that we focus on advertising search good related products, consumer electronics, where consumers face lower consumption uncertainty. Previous studies investigated ad personalization effectiveness in the context of experience goods, holiday services (Lambrecht and Tucker, 2013) and sports fashion (Bleier and Eisenbeiss, 2015a).

Second, to our knowledge, this is the first study that investigates the effect of social targeting in the context of personalized ads. While prior research found positive effects of social targeting (Bakshy et al., 2012; Tucker, 2016), our results, surprisingly, show that socially targeted consumers react less positive to personalized ads. This negative effect is enhanced for more specific ad personalization. We use the uniqueness theory and distinguish between social identities and personal identity to explain the reactance behavior of consumers.

The increasing popularity of using social networks to reach consumers underlines the importance of our study to business practice. Our findings suggest that firms can benefit more when they readdress consumers with highly personalized product-specific ads, especially as soon as possible after consumers' website visits. Firms need to be very cautious about the current practice of socially targeting consumers by default in social advertising as it likely leads to negative consumer responses to personalized ads.

The rest of the paper will be organized as follows. First, we will provide the underlying theoretical foundation of our study. We develop hypotheses for the effects of personalization specificity, social targeting, and their interaction. Next we present our method, empirical model, and results. To conclude, we discuss our findings, present theoretical and practical implications, and point out limitations and potential for future research.

2.2 Theory

Advertising personalization is defined as firm-initiated adjustment of advertising content towards the preferences of consumers with the goal to improve consumer responses to ads (Arora et al., 2008). Personalized communication with consumers has been found to increase customer loyalty and consumers' attention towards marketing communication (Ansari and Mela, 2003). In the information systems literature, advertising personalization is categorized as decision personalization, supporting consumers to more easily identify and choose products that match their preferences (Thirumalai and Sinha, 2013). Matching consumers' preferences that change dynamically with advertising content remains challenging for firms. Consumers (re)construct their preferences, utilizing accumulated and relevant experiences and gathering additional information which ultimately leads to stabilized preferences (Hoeffler and Ariely, 1999).

Prior literature demonstrated positive effects of personalized advertising based on consumers' past browsing behavior, commonly called retargeting (Bleier and Eisenbeiss,

2015a; Lambrecht and Tucker, 2013). For this type of advertising firms make use of information on consumers' browsing behavior on their website to readdress consumers with ads matching this behavior on external websites. Generally, browsing behavior, especially which products consumers browse, has been pointed out as a good indicator of consumers' preferences. Nevertheless, recent studies investigating the optimal level of personalization in advertising come to inconsistent conclusions. Related work finds that less personalized generic brand ads outperform dynamically personalized ads, which only work better for consumers that have narrowly defined preferences (Lambrecht and Tucker, 2013). In contrast, Bleier and Eisenbeiss (2015a) claim that a high degree of content personalization in ads leads to more positive consumer responses than less personalized ads.

Although personalization has been shown to positively affect consumers' reactions to advertising, firms are struggling with how specific advertising personalization should be. While less specific personalized advertising uses consumers' inferred preferences to recommend a product category, highly specific personalized advertising recommends a specific product. Firms need to decide which level of personalization specificity yields the highest returns for them. This decision is difficult as, based on theory, there are arguments for the superiority of both category-specific and product-specific advertising personalization. To demonstrate the conflict in theoretical reasoning we develop competing hypotheses in the following sections.

2.2.1 Category-Specific Ad Personalization

Although more specific advertising personalization offers the chance to increase advertising relevance for consumers, its success is highly dependent on the preference prediction accuracy underlying the personalization. Misclassification of consumer preferences, for example presenting a consumer with a product that she dislikes, can lead to consumer resistance and annoyance (Arora et al., 2008). Theory suggests that category preferences are more likely to be classified accurately as consumer preferences for product categories are more stable than preferences for specific products (Simonson, 2005; Tam and Ho, 2006). Product-specific preferences are constructed up until the moment of the product purchase. Previous studies found that, on average, generic brand ads outperform ads for specific products (Lambrecht and Tucker, 2013). Therefore, less specific personalization can be more favorable than highly specific personalization as it decreases the risk of misclassifying consumer preferences triggering consumer annoyance.

Triggered by consumer privacy concerns, consumers might react negatively to ads when they have concerns that too much of their personal information is being used to personalize ads. The dilemma of personalization leading to an increase in ad relevance but simultaneously increasing consumer privacy concerns is called personalization-privacy paradox in the information systems literature (Awad and Krishnan, 2006; Lee et al., 2011; Sutanto et al., 2013). Research indicates that consumers are more concerned about personalization when their awareness of personalization is increased, for example through the inclusion of their names in promotional e-Mails (Wattal et al., 2012). When confronted with highly specific personalization, consumers have higher privacy concerns stemming from the use of their personal information. These arguments lead us to the following hypothesis:

Hypothesis 1a (H1a Competing): *Category-Specific Personalization leads to a more positive consumer response than Product-Specific Personalization in Social Retargeting.*

2.2.2 Product-Specific Ad Personalization

When consumers browse particular products, advertisers can infer that consumers may be interested in these or similar products. Showing ads with a specific product, allows advertisers to be closer to the actual preferences of a consumer. Using more details on consumers' browsing behavior and advertising specific products that match consumers' preferences allows advertisers to achieve higher ad relevance (Bleier and Eisenbeiss, 2015a; Tam and Ho, 2005). More relevant advertising content is processed with more cognitive effort and therefore more likely to influence consumers' preference construction (Ho and Bodoff, 2014). Advertising content that is in line with consumers' preferences is more likely to be considered via the central route of persuasion (Tam and Ho, 2005). A consumer that has looked at a particular product is more likely to have invested time and effort in product evaluation to narrow down her choice set. In this case, consumers may perceive less specific ads, that advertise a product category, as less relevant, as these ads refer to a step in their purchase process that they have already taken.

Moreover, consumers consider product-specific ads more relevant as they are more likely to recognize that these ads are personalized to their preferences. Previous research found that perceived personalization increases consumers' intention to adopt recommendations (Komiak and Benbasat, 2006). An increase in perceived personalization was also found to decrease consumers' ad avoidance (Baek and Morimoto, 2012).

In the context of social networks, perceived personalization has been found to increase consumers' perceived ad relevance as well as their intention to click on ads (Keyzer et al., 2015).

Consumers use a specificity heuristic when assessing the quality of recommendations (Palmeira and Spassova, 2012): More specific recommendations are evaluated more favorably. More extreme advertising claims for reputable advertisers have been shown to positively influence ad credibility (Goldberg and Jon Hartwick, 1990). Another advantage for consumers in more specific personalization is that they receive customized offers that allow them to more easily make decisions (Xiao and Benbasat, 2007). While less specific personalization requires consumers to choose between different product options, product-specific recommendations can reduce choice overload effects and minimize search costs (Ansari and Mela, 2003). This allows consumers to make purchase decisions more efficiently. Assisting consumers in making their choices can help consumers to overcome the confusion originating especially in large product assortments (Thirumalai and Sinha, 2013). These arguments dispute the favorability of less specific personalization in advertising and lead us to derive the following competing hypothesis:

Hypothesis 1b (H1b Competing): *Product-Specific Personalization leads to a more positive consumer response than Category-Specific Personalization in Social Retargeting.*

2.2.3 Social Targeting

Social network platforms have extensive information about their users. This information includes demographics, preferences and interests, as well as social connections. Recent research has focused on what can be inferred from consumers' social connections and how this information can be leveraged, e.g. for the purpose of personalization (Aral and Walker, 2011; Muchnik et al., 2014). One way to leverage social connections for marketing purposes is social targeting. In social targeting, firms use consumers' social connections to infer their preferences and subsequently address consumers whose preferences match with the firm's product offerings. Next to that, the social connections underlying the targeting are made explicit in the ad text as social endorsements with the aim to increase consumers' trust in the advertiser and the perceived relevance of the ad. These types of ads are then called social advertising, where "ads are targeted based on underlying social networks and highlight when a friend has 'liked' a product or organization" (Tucker, 2016, p. 1). We define social targeting as the combination

of using consumers' underlying social networks to target them and making these connections explicit in the ad text as social endorsements. While most research finds that leveraging social connections to socially target consumers has positive implications on ad performance (Bakshy et al., 2012), more recently, there are examples in which consumers do not generally react positively to socially targeted ads (Tucker, 2016). Again, we reveal the conflicts in theories used to explain the effects of social targeting by developing competing hypotheses.

Positive Effect of Social Targeting

Prior research has shown that users that are connected in social networks are likely to share similar preferences, which is referred to as *homophily of connected users* (Aral et al., 2009). This preference similarity of connected users can be used to infer consumers' preferences. Knowing the preferences of a consumer's friends, firms can target and personalize advertising content based on these social connections. Prior studies found that social network friends of consumers with a high affinity for a brand, are likely to have an affinity for this brand as well (Provost et al., 2009). Furthermore, consumers are usually influenced by their peers' actions when forming their preferences (Tucker, 2016). In social advertising, the social connections underlying the targeting are made explicit. Names of users that are fans of the advertising brand as well as friends with the targeted consumer are displayed in advertisements. This so-called *social endorsement* is supposed to increase ad effectiveness by exploiting a user's social network via *social influence*. The use of social endorsements provides a positive influence on how individuals perceive advertising on social media (Bakshy et al., 2012). This type of influence resulting from socially endorsed advertising is called informational social influence (Kwahk and Ge, 2012). Informational social influence helps individuals to accept externally received information to be true (Deutsch and Gerard, 1955). In social advertising this means a socially endorsed ad is viewed as being more credible. Consumers perceive the information that their friends are connected to the advertising brand as evidence for the quality of the ad content. Prior research found evidence of a persuasive effect (informational social influence) of social endorsements in social advertising being present in addition to a targeting effect as users with similar interests tend to be connected (homophily of connected users) (Bakshy et al., 2012). These arguments lead to the following hypothesis:

Hypothesis 2a (H2a Competing): *Social Targeting leads to a more positive consumer response in Social Retargeting.*

Negative Effect of Social Targeting

There are also theoretical arguments that point towards a negative effect of social targeting. Recent work shows that consumers tend to not conform with their friends' actions in the context of social networks when they want to express their personality (Sun et al., 2017). Despite the fact that informational social influence has been shown to trigger conformity (Deutsch and Gerard, 1955) consumers are simultaneously striving for uniqueness (Chan et al., 2012). Uniqueness theory describes consumers' drive to be different from others where "too much similarity leads to a negative emotional reaction" (Berger and Heath, 2008, p. 594). Uniqueness theory combines the urge of individuals to identify themselves with others (social identities) as well as the need to differentiate themselves to define their personal identity (Snyder and Fromkin, 1980). Individuals tend to adhere to favorable social identities while simultaneously defining their personal identity through differentiation (Brewer, 1991). While the personal identity is unique, social identities are related to common characteristics that are popular in a certain social group and adopted by individuals.

In the context of personalized advertising, consumers are confronted with a conflict when being socially targeted. Advertisers address them with ads that are personalized to their preferences giving consumers the impression that recommendations are made uniquely for them. This should allow consumers to identify with the personal offers that matches their preferences. However, by using social connections to target consumers and making these social connections explicit in the ads the presented recommendations are being depersonalized. Social identities "depersonalize the self-concept" (Brewer, 1991, p. 476). The fact that consumers see ads that recommend products specifically for them does conceptually not match with the social endorsement of friends which results in a decrease in perceived personalization. We hypothesize:

Hypothesis 2b (H2b Competing): *Social Targeting leads to a more negative consumer response in Social Retargeting.*

2.2.4 Personalization Specificity and Social Targeting

We showed that there are theoretical arguments for both a positive and negative effect of social targeting on consumer responses to social retargeting ads. The investigation of the moderating role of social targeting on personalization specificity can give us deeper insights into the theoretical explanation for this effect.

If socially targeted consumers react more positively to personalized ads, social targeting should also positively moderate the relationship between personalization specificity and consumer responses to ads. To increase the accuracy of personalization in advertising, firms can leverage consumers' social networks. Consumers that are connected with friends that like a product on social networks are more likely to have preferences that favor this product as well (Bakshy et al., 2012; Tucker, 2016). The fact that connected consumers share similar preferences (homophily of connected users) allows advertisers to gain additional information on consumers' preferences. By using information on consumers' social connections, firms can achieve higher accuracy in the prediction of consumer preferences leading to an increase in ad relevance and more positive consumer responses (Arora et al., 2008). This increase in accuracy allows firms to make more specific product recommendations to consumers.

Further, social endorsements that are included in socially targeted ads, allow consumers to understand that their friends are connected to the advertiser, leading to an increase in trust in the advertiser (Bakshy et al., 2012). Trust has been shown to decrease consumers' reactance and privacy concerns towards personalized recommendations (Bleier and Eisenbeiss, 2015b; Komiak and Benbasat, 2006). Therefore, in the presence of a positive direct effect of social targeting on consumer responses to personalized ads, we expect social targeting to positively moderate product-specific ad personalization.

Hypothesis 3a (H3a Competing): *Social Targeting positively moderates the effect of Product-Specific Ad Personalization on consumer responses to Social Retargeting ads.*

On the contrary, the theoretical arguments for a negative effect of social targeting on consumer responses to social retargeting point towards a negative moderating effect of social targeting on product-specific ad personalization. As argued above, the inclusion of friends' names in the advertising text that endorse an ad depersonalizes the ad which conflicts with the personalized recommendation made by the advertiser (Brewer, 1991). This conflict is stronger when the ad personalization is more specific, as consumers perceive such a product recommendation as more unique and therefore as conflicting more strongly with the inclusion of social identities in the ad. When a product-specific ad triggers a higher degree of perceived personalization with consumers, the presence of social identities, through the inclusion of friends' names in the ad text, depersonalizes the ad more strongly. A friend endorsement for a product category still allows a consumer to differentiate from friends by choosing a product within the advertised

category. But a friend endorsement for a specific product leaves a consumer with a limited ability to differentiate and make a unique product choice signalling her personal identity. Therefore, in the case of a negative direct effect of social targeting, we expect that social targeting is negatively moderating the effect of highly specific ad personalization on consumer responses to personalized ads.

Hypothesis 3b (H3b Competing): *Social Targeting negatively moderates the effect of Product-Specific ad Personalization on consumer responses to Social Retargeting ads.*

2.3 Field Experiment

We conducted a large-scale randomized field experiment in collaboration with a major European e-commerce company to investigate the effectiveness of different levels of personalization specificity in social retargeting. Our partner company sells a wide range of products with a focus on consumer electronics. For our study, we focus on the product categories of laptops, cameras, tablet computers, smart phones, and televisions. For the experiment, we solely advertise to consumers in the newsfeed area of Facebook as the newsfeed is generally the focal area for consumers and captures most of their attention (Wishpond, 2014).

Consumers that browsed the partner company’s website, viewed at least a category-level page, and were active users of Facebook, were eligible to participate in our experiment. Using their browsing behavior, we randomly assigned either category- or product-specific personalized social retargeting ads to these consumers. The random assignment to the two personalization treatments took place on our partner company’s website by assigning one of two conditions to consumers’ Facebook pixels (a cookie stored on consumers’ computers that can be read by Facebook). Consumers that then visited Facebook were treated with ads matching this assignment. This assignment method offers an advantage over conventional cookie targeting. Once consumers reach Facebook’s website without deleting their cookie, they are allocated to their assigned treatment group. Facebook stores this assignment linked to a consumer’s user account. This way, consumers remain in a treatment group even if they delete their cookies after reaching Facebook. If consumers delete their cookie before reaching Facebook, they are not addressed with advertising and remain eligible to participate in the experiment in case they re-visit our partner company’s website and receive a new, independent assignment to a treatment group. Additionally, we address the hypothetical case that

consumers are assigned to both treatment groups, i.e. because of technical issues, by excluding consumers with a double assignment from the experiment. This way, we can guarantee a clean between-subject design for our personalization treatments.² Throughout the experiment, consumers remained in their respective treatment groups.³

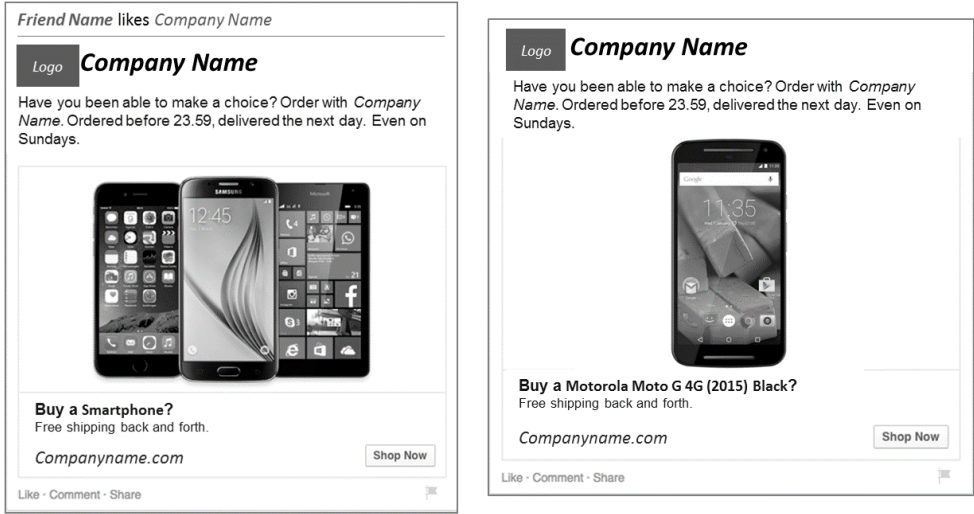
We operationalized personalization specificity by displaying ads that were related to either the last visited product category (category-specific) or the last visited product (product-specific). We made sure that the two types of ads were exactly the same besides the product and category attributes as shown in Figure 2.1. In contrast to former studies that, based on the chosen personalization algorithm, displayed several products to consumers simultaneously, we only advertised a single category or product per ad. This way, we aim to isolate the effect originating from category- and product-specific personalization and rule out alternative explanations originating from the difference in visual appeal or confounding factors originating from the composition of choice sets that are presented to consumers in ads with several products. Category ads displayed the three most popular products within a product category (in terms of sales) in a single ad image.

By default, ads were socially targeted when consumers were via one or several friends connected to the Facebook page of our partner company. Socially targeted ads displayed the name(s) of the friend(s) that liked our partner company’s Facebook page by stating “[Friend Name] likes [Company Name]” at the top of the ad (see Figure 2.1a). Generally, Facebook’s advertising algorithm displays friend connections to the advertising firm whenever possible. This means that friend connections need to be present and the friend that is supposed to appear as an endorser in the ad has not withdrawn Facebook’s right to use her name for advertising purposes in her account settings (Tucker, 2016). Notably, our social targeting operationalization does not represent an experimental treatment variable but rather a consumer characteristic, i.e. being connected to the advertiser’s Facebook page, that is used by the advertiser to target consumers and which is made explicit in the ad text.

We ran our field experiment for 28 consecutive days in May 2015. Overall, our experiment generated 3,476,626 impressions for 198,234 individual consumers. Consumers

²We further discuss the limited potential of contamination through social interactions with ads in Appendix A2.1.

³We run all ad campaigns with the same budget restrictions, the firm’s willingness to pay per 1,000 impressions (CPM). This way the ad platform, Facebook, has no incentive to select different types of consumers into the ad treatment groups when being paid per impression. This differs from campaigns that are optimized based on consumers’ propensity to respond as common in cost per click (CPC) or cost per acquisition (CPA) optimized campaigns. We do not detect a systematic difference in the costs per impression for consumers in the two ad treatment groups.



a) Category-Specific Social Retargeting Ad

b) Product-Specific Social Retargeting Ad

Figure 2.1: Experiment Conditions

were shown a maximum of two ads on a daily basis. We measure the ad effectiveness using both clicks and purchases. Clicks measure how many times consumers have clicked on a social retargeting ad. Purchases indicate how many times consumers have purchased from our partner company within 28 days after clicking an ad. The Facebook ad reporting tool does attribute purchases to a consumer's last clicked ad impression before the purchase. The ads generated 25,577 clicks, leading to an overall average click-through rate of 0.736%, and 1,070 purchases (within 28 days after having clicked on an ad), resulting in an average click-to-conversion rate of 4.183%. Consumers were excluded from the experiment after conducting a purchase with our partner firm to avoid serving consumers ads of products they had already purchased.

2.4 Analysis and Results

Figure 2.2 presents model-free evidence of the average click-through and purchase rates of consumers that are confronted with either category-specific personalized or product-specific personalized ads. We compare average click-through rates and purchase rates (from impression to sale). For both measures highly personalized product-specific ads outperform less personalized category-specific ads. These model-free results offer an

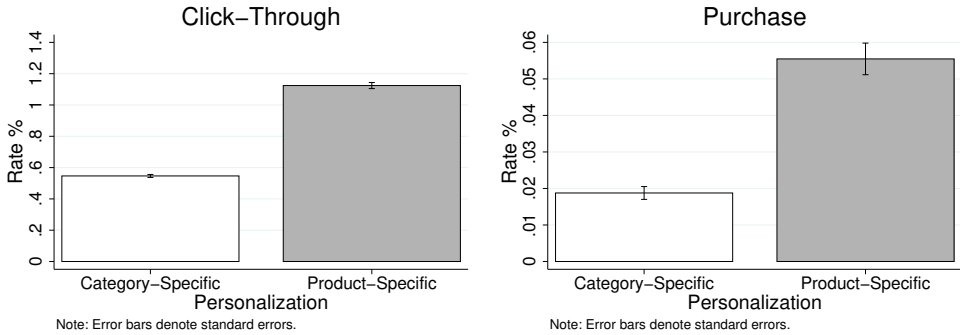


Figure 2.2: Consumer Response for Category-Specific and Product-Specific Ad Personalization.

initial hint on the more positive consumer response to more specific ad personalization, however, social targeting, seasonality effects of different sales dates, and heterogeneous popularity of different product categories might influence consumer responses to ads. To control for the impact of these factors on consumer responses, we move on to estimate logistic regression models. These logistic regression models take the binary nature of our consumer response variables, i.e. click and purchase, into account.

We analyze the data from our field experiment on an ad impression-level. An ad impression represents the event of displaying an ad to a consumer. Our model estimates the unobserved probability of a positive consumer response to an ad, i.e. a click or a purchase,⁴ for an ad impression i . We denote the probability of a positive consumer response as $Pr(\text{Consumer Response}_i = 1)$ and model the latent probability, denoted by U_i , by using a logit function of personalization specificity and social targeting as well as additional control variables. We assume an independent and identically distributed extreme value distribution of the error term.

$$Pr(\text{Consumer Response}_i = 1) = \frac{\exp(U_i)}{1 + \exp(U_i)}$$

$$U_i = \beta_0 + \beta_1 \text{product-specific}_i + \beta_2 \text{social targeting}_i + \beta_3 \text{product-specific}_i \times \text{social targeting}_i + \gamma X_i + \epsilon_i$$

⁴We define consumer response as a consumer's binary response decision. For our different analyses we exchange the binary dependent variables with a focus on click and purchase behavior.

Where β_0 is the constant term. $product-specific_i$ is equal to 1 when an impression features a product-specific ad, 0 when featuring a category-specific ad. The binary variable $social\ targeting_i$ is equal to 1 when an ad is targeted to consumers that are via friends connected to our partner company’s Facebook page and includes a social endorsement, otherwise 0. The coefficient for the interaction term, $product-specific_i \times social\ targeting_i$, allows us to assess whether socially targeted consumers react differently to product-specific ad personalization. X_i represents a vector of ad controls including day, product category, country, and device (mobile and desktop) fixed effects. As our ad treatments are randomized, one purpose for including the control variables is to increase efficiency in our estimations. At the same time, our social targeting variable is not a randomized treatment. Therefore the control variables help us to rule out biases originating in heterogeneity such as seasonality in sales dates, differences in product category attractiveness, cultural differences between consumers in different countries, as well as different responsiveness across devices. ϵ_i represents the idiosyncratic error term. Table 2.1 gives the summary statistics for our main variables.

Table 2.1: Descriptive Statistics (n = 3,476,626)

VARIABLES	Mean	Std. Dev.	Min	Max
click	0.0074	0.0855	0	1
purchase	0.0003	0.0175	0	1
product-specific	0.3272	0.4692	0	1
social targeting	0.8221	0.3824	0	1

2.4.1 Likelihood to Click

The results of our logistic regressions estimating the click probabilities are summarized in Table 2.2. Column (1) presents our model with only the control variables for dates, product categories, countries, and devices. In column (2) we include the binary treatment variable $product-specific_i$ to estimate the effect of highly personalized product-specific ads compared to less personalized category-specific ads on click probabilities. We find that high levels of personalization specificity significantly increase the likelihood of an ad impression leading to a click ($\beta_{product-specific} = .742$, $p < .001$). This confirms the findings from our model-free comparison between category- and product-specific ads. Furthermore, this result confirms the positive effects of more specific advertising personalization and supports *hypothesis H1B*. In column (3) we enter the

main effect of *social targeting_i*. This variable essentially estimates whether consumers that are connected to the firm and see (a) social endorsement(s) in an ad react differently to social retargeting ads compared to consumers that are not socially targeted. Contrary to prior findings in the literature, we find that socially targeted ad impressions lead to lower click-through probabilities ($\beta_{social\ targeting} = -.176, p < .001$) thus supporting *hypothesis H2B*. Next, we include the interaction between *product-specific_i* and *social targeting_i* into our model to investigate whether consumers that are socially targeted react differently to product-specific ad personalization. We find that consumers that are connected to the advertiser are less likely to click on highly personalized ads ($\beta_{product-specific \times social\ targeting} = -.070, p = .025$). This means that highly personalized ads are in fact less effective for consumers that are connected with the advertiser and see a social endorsement. This contradicts the notion that connected consumers are more likely to accept higher levels of personalization. In contrast to the literature, our finding points out that there is a conflict between personalization and social targeting, in which social connections are made explicit. This conflict seems to be stronger when ad personalization is more specific, supporting *hypothesis H3B*.

Table 2.2: Logistic Regressions for Click Probabilities

VARIABLES	(1) Click	(2) Click	(3) Click	(4) Click
product-specific		0.742*** (0.013)	0.741*** (0.013)	0.796*** (0.028)
social targeting			-0.176*** (0.016)	-0.141*** (0.022)
product-specific \times social targeting				-0.070** (0.031)
date controls	Yes	Yes	Yes	Yes
product category controls	Yes	Yes	Yes	Yes
country controls	Yes	Yes	Yes	Yes
device controls	Yes	Yes	Yes	Yes
Constant	-4.052*** (0.028)	-4.320*** (0.028)	-4.177*** (0.031)	-4.205*** (0.034)
Observations	3,476,626	3,476,626	3,476,626	3,476,626
Chi2	2,647.067	6,046.068	6,169.709	6,174.713
-2 Log Likelihood	299,593.098	296,194.097	296,070.456	296,065.451
AIC	299,661.098	296,264.097	296,142.456	296,139.451
BIC	300,105.191	296,721.252	296,612.672	296,622.730

Standard errors in parentheses, *** p<0.01, ** p<0.05, * p<0.1

As the interpretation of interaction coefficients in logistic regression is not straightforward, we calculate the marginal effects of our logit model estimates (Forman, 2005; Luo et al., 2013). We find, consistent with the coefficients in our logit model, that ads with product-specific personalization that are not socially targeted perform better. The marginal effect for the interaction between $product-specific_i \times social\ targeting_i$ is $-.0005$ ($p < .05$). We further assess the economic implications of our findings (see Appendix A2.2, Table A2.2). We see that product-specific ad personalization leads to an absolute increase of 0.61% in click probability and a relative increase of 120.38% compared to category-specific personalization. Product-specific ads lead on average to around 6 additional clicks per 1,000 impressions compared to category-specific ads.⁵ For our experiment with around 3.5 million impressions, treating consumers with solely product-specific ads instead of category-specific ads would result in an increase of 21,425.25 clicks.⁶ Similarly we see that social targeting decreases the probability of a click by 0.10% in absolute terms and 13.03% relative to non-socially targeted ads. When product-specific ads are socially targeted, this decrease is even stronger with additional 0.05% decrease in absolute click probability.

2.4.2 Likelihood to Purchase

We repeat the estimation process using the dependent variable $purchase_i$. As for clicks, we model the latent probability of a purchase by using a logit function of personalization specificity and social targeting.⁷

In line with our findings for click probabilities, we see that product-specific social retargeting ads increase purchase probabilities ($\beta_{product-specific} = 1.090$, $p < .001$) (Table 2.3 column(2)). We find that the effect of product-specific personalization remains positive and significant also when entering the main effect of social targeting as well as their interaction effect. Again, we confirm that product-specific personalization leads to a more positive consumer response than category-specific personalization, supporting *hypothesis H1B*. We also find a significant negative effect for social targeting on the probability to purchase ($\beta_{social\ targeting} = -1.026$, $p < .001$), supporting *hypothesis H2B*. In line with the analysis for click probabilities, we see that more specific ad personalization is less favorable for socially targeted ads. However, the result is not

⁵We multiply the absolute increase in clicks by 1,000 impressions to get to this figure.

⁶We multiply the absolute increase in clicks by 3.5 million impressions to get to this figure.

⁷In our model for purchase probabilities, we measure $social\ targeting_i$ as the percentage of socially endorsed impressions for a specific ad attribute combination since the advertising platform does not provide exact data that allow linking a socially targeted impression with a purchase. For more information on the structure of our data please consult Appendix A2.3.

significant and does therefore not provide full support for *hypothesis H3A*. As in our model for click probabilities, the marginal effect of *product-specific_i × social targeting_i* is, with -0.00002 ($p = .871$), negative. Nevertheless, this effect is not significant for purchase probabilities.

Table 2.3: Logistic Regressions for Purchase Probabilities

VARIABLES	(1) Purchase	(2) Purchase	(3) Purchase	(4) Purchase
product-specific		1.090*** (0.062)	1.084*** (0.062)	1.144*** (0.378)
social targeting			-1.026*** (0.215)	-0.973** (0.387)
product-specific × social targeting				-0.075 (0.463)
date controls	Yes	Yes	Yes	Yes
product category controls	Yes	Yes	Yes	Yes
country controls	Yes	Yes	Yes	Yes
device controls	Yes	Yes	Yes	Yes
Constant	-7.231*** (0.139)	-7.677*** (0.143)	-6.836*** (0.226)	-6.878*** (0.346)
Observations	3,476,626	3,476,626	3,476,626	3,476,626
Chi2	107.319	418.820	439.780	439.806
-2 Log Likelihood	19,320.559	19,009.058	18,988.098	18,988.072
AIC	19,388.559	19,079.058	19,060.098	19,062.072
BIC	19,832.652	19,536.213	19,530.315	19,545.350

Standard errors in parentheses, *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Again, we provide a table in the appendix that presents the economic implications of our findings (see Appendix A2.2, Table A2.3). Product-specific ad personalization leads to an absolute increase of 0.04% in purchase probability and a relative increase of 213.92% compared to category-specific personalization. Product-specific ads lead on average to around 0.37 purchases more per 1,000 impressions.⁸ For our experiment with around 3.5 million impressions this would lead to an increase of 1,301.30 purchases when solely serving product-specific ads instead of category-specific ads.⁹ Social targeting decreases the probability of a purchase by 0.03% in absolute terms and 62.22% relative to non-socially targeted ads. When product-specific ads are socially targeted, this decrease is slightly stronger, but not significant.

⁸We multiply the absolute increase in purchases by 1,000 impressions to get to this figure.

⁹We multiply the absolute increase in purchases by 3.5 million impressions to get to this figure.

2.4.3 Robustness Checks

To assess the robustness of the results in our main analysis we analyze our data on different levels of aggregation and re-estimate our models with alternative dependent variables and different estimators. First, we analyze our data on a consumer level.¹⁰ We find results consistent with our main analyses (see Appendix A2.4, Table A2.4).

Second, we estimate our models for click probabilities replacing our dependent variable clicks with unique clicks, measuring only one click per consumer in case a consumer clicks on an ad several times. We find that our results are consistent with the results for clicks (see Appendix A2.5, Table A2.5).

Third, for our estimates for both click and purchase probabilities, our results remain consistent when estimating probit and linear probability models with the same model specifications (see Appendix A2.6, Table A2.6). In fact, we find that social targeting negatively moderates the positive impact of product-specific personalization on purchase probabilities in our linear probability model estimating purchase probability, supporting *hypothesis H3B*.

Fourth, we analyze the effect of our focus variables on purchase probability by operationalizing a purchase as a buying event within 7 days after seeing an ad instead of 28 days after having clicked on an ad (see Appendix A2.7, Table A2.7). Our results remain largely consistent.

Last, due to the structure of our dataset, we are analyzing our data on an impression-level and are not able to control for consumer-specific variables. One consumer-specific variable that might potentially bias our results is the frequency of ad impressions per consumer. This might be the case if the frequency of ad impressions per consumer differs between consumers in our two personalization treatment groups as well as for socially targeted and non-socially targeted consumers. A higher average frequency of impressions per consumer for any of these focus ad attributes is likely to decrease the effectiveness of a single impression. To control for the frequency of impressions we generate a variable measuring the number of uniquely addressed consumers per ad attribute combination and control for this variable in our analysis (see Appendix A2.8). Our results remain consistent.

¹⁰Strictly speaking, we analyze our data for unique consumer, campaign, day combinations. Our data allows us to infer how many individual consumers were addressed per campaign and how often these consumers clicked on an ad or conducted a purchase on a specific day of the campaign. Facebook aggregates data on campaign level and does not allow us to follow consumers throughout the duration of the campaigns. Therefore, we decide to conduct our main analysis on an ad impression level.

2.5 Additional Analysis

Our results provide important insights into advertising personalization in social retargeting. To make sure that it is personalization that is driving our results and that we have operationalized our personalization specificity treatment adequately, we run additional analyses. We investigate to what extent our results remain stable when controlling for (1) temporal targeting, as well as the (2) browsing behavior of consumers. To mitigate the concern that our social targeting variable is not an experimental treatment and socially targeted consumers might react differently to personalized ads due to unobserved consumer characteristics we re-run our analysis controlling for (3) consumer demographics.

2.5.1 Temporal Targeting and Preference Development

Temporal targeting aims at addressing consumers at the point in time when they are most receptive towards marketing messages (Luo et al., 2013). Timing is an essential aspect of addressing consumers with personalized advertisements as consumers develop their preferences over time through consideration processes and gathering of further experiences (Hoeffler and Ariely, 1999; Simonson, 2005). Timing is especially crucial for personalized advertising techniques that base their preference predictions on consumers' browsing behavior. An increase in the time between observing consumers' behavior and addressing consumers with personalized ads based on this behavior is likely to be correlated with a change in consumers' preferences. The theory of constructive preferences argues that preferences evolve over time as consumers generate more experiences that influence their product preferences (Simonson, 2005). Therefore, advertising that is personalized based on consumer behavior loses its effectiveness with a decrease in recency of the related consumer behavior. Our personalized ads, both category- and product-specific, should generally decrease in performance over time in case they are being perceived as personalized.

We introduce the variable *temporal targeting_i* into our model, which is a continuous variable indicating the time between a consumer's website visit and an ad impression in days. A lower value for *temporal targeting_i* indicates a higher degree of temporal targeting - fewer days between the consumer's website visit and the ad impression. We run our models including *temporal targeting_i* for both click and purchase probabilities, but focus on the temporal distance between website visit and ad impression from one to seven days. The data for our additional analysis consists of 1,457,527 impressions

for 148,588 individual consumers.¹¹ Reducing the temporal distance window allows us to further assess the robustness of our results when focusing on a smaller time window compared to our main analysis in which consumers are addressed with ads up to 4 weeks after their website visit.

It is likely that more active consumers, i.e. who visit Facebook more frequently, are more likely to be addressed with a smaller temporal distance to their website visit. These active consumers are also more likely to click on ads and purchase products online, which has been coined activity bias in previous work (Lewis et al., 2011). This activity bias might amplify the effect of our temporal targeting coefficient as the variable measures not only the effect originating in the increase in time between a website visit and the confrontation with an ad, but also the effect of a decrease in consumers' activeness. To control for this issue, we simultaneously introduce the variable *consumer activeness_i* which measures the average number of reactions to a unique ad attribute combination per impression.¹² This variable gives a good indication of how active consumers respond to an ad. Intuitively, a higher rate of actions towards ads indicates more active consumers, which should also be more likely to click and purchase.

Table 2.4 shows the results of our temporal targeting models. We first focus on the inclusion of *temporal targeting_i* and its interplay with product-specific personalization (Column 1 and 3). Then we investigate the impact on the whole model including social targeting (Column 2 and 4). Generally, we find that product-specific personalization consistently outperforms less personalized category-specific ads in terms of click probability. Furthermore, as expected, we find that consumers respond less positive to ads with a decrease in temporal targeting. This points towards an adequate operationalization of personalization in our ads. With a decrease in the recency of consumers' behavior used for our personalization the performance of both ads decreases as consumers preferences evolve over time. In addition, we find a significant and positive coefficient for the interaction between *product-specific_i* and *temporal targeting_i* on click probabilities. This suggests that highly personalized ads decrease in performance slower, which contradicts earlier findings (Bleier and Eisenbeiss, 2015a; Simonson, 2005).

¹¹The reason for focusing on ad impressions served within the first 7 days after consumers' website visits is that for larger temporal distances the granularity of our data moves from a daily- to a weekly-level. The focus on impressions within the first 7 days leads to a decrease in the sample size used for this analysis.

¹²The variable measures the average number of reactions per impression on a unique ad attribute combination. These reactions include things such as likes, comments, and shares as response to an ad. More details on the structure of our data can be found in Appendix A2.3.

Table 2.4: Logistic Regressions for Click and Purchase Probabilities Controlling for Temporal Targeting

VARIABLES	(1) Click	(2) Click	(3) Purchase	(4) Purchase
product-specific	0.278*** (0.031)	0.286*** (0.046)	0.666*** (0.142)	0.432 (0.413)
social targeting		-0.088*** (0.031)		-0.310 (0.397)
product-specific \times social targeting		-0.011 (0.043)		0.283 (0.473)
temporal targeting	-0.190*** (0.007)	-0.190*** (0.007)	-0.241*** (0.035)	-0.242*** (0.035)
product-specific \times temporal targeting	0.021** (0.009)	0.021** (0.009)	0.036 (0.045)	0.037 (0.045)
consumer activeness	13.884*** (0.107)	13.873*** (0.107)	6.762*** (0.149)	6.758*** (0.150)
date controls	Yes	Yes	Yes	Yes
product category controls	Yes	Yes	Yes	Yes
country controls	Yes	Yes	Yes	Yes
device controls	Yes	Yes	Yes	Yes
Constant	-4.127*** (0.044)	-4.053*** (0.050)	-7.234*** (0.209)	-6.975*** (0.391)
Observations	1,457,527	1,457,527	1,457,527	1,457,527
Chi2	25,102.828	25,121.694	1,555.911	1,556.517
-2 Log Likelihood	146,893.266	146,874.400	10,541.392	10,540.787
AIC	146,969.266	146,954.400	10,617.392	10,620.787
BIC	147,432.572	147,442.090	11,080.698	11,108.477

Standard errors in parentheses, *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

In column (2) (Table 2.4) we enter $social\ targeting_i$, as well as the interaction with $product-specific_i$ into our model. Our results remain mostly consistent. While product-specific personalization increases click probabilities, social targeting decreases the likelihood of a click. Again, temporal targeting shows a significant and negative effect pointing to a decrease in performance of social retargeting ads over time. Although the coefficient remains negative, we do not find a significant effect for the interaction between product-specific personalization and social targeting.

We replicate our analysis but exchange $click_i$ with $purchase_i$ as binary dependent variable. We estimate logistic regression to assess the impact of temporal targeting on purchase probabilities (see Table 2.4, column (3-4)). Again, we see a positive and significant effect of product-specific personalization. This effect does not remain significant when including social targeting in the model. The coefficient for temporal

targeting confirms, for both models estimating purchase probabilities, that our ads decrease in performance with an increase in temporal distance between website visit and ad confrontation. This once more supports our operationalization of ad personalization. As consumers develop their preferences over time, the performance of our personalized ads decreases with an increase in the temporal distance to a consumer’s website visit.

2.5.2 Browsing Depth and Preference Development

Consumers’ browsing depth, i.e. how deep consumers browse a firm’s website, can give a good indication about how well consumers have defined their preferences (Bleier and Eisenbeiss, 2015a; Lambrecht and Tucker, 2013; Moe, 2003). To confirm that the higher performance of highly personalized, product-specific ads is coming from a higher level of personalization specificity and not from confounding factors like ad attractiveness, we assess to what extent personalization that is closer to consumers’ browsing behavior leads to more positive consumer responses. We assume that product-specific ads perform better when they resemble consumers’ browsing behavior more closely, since they predict consumers’ preferences more accurately. To assess this relationship, we distinguish between consumers that have solely browsed category pages and consumers that have browsed specific product pages. We introduce a binary variable indicating if a consumer browsed product pages (*product browsing_i* = 1) or only category-level pages (*product browsing_i* = 0) before ad exposure.¹³ We expect to find a positive and significant effect for the interaction of *product-specific_i* and *product browsing_i*. Since consumers are more likely to recognize personalization that happens immediately after their website visit, we restrict the time window of our analysis, as before, to include impressions served within the first to seventh day after a consumer’s website visit.

Table 2.5 shows the results of our models including *product browsing_i*. We find a significant and positive effect of product-specific personalization on click probability (Column (1)). We also find that the interaction between product-specific personalization and product browsing is positive and significant for click probabilities. This supports the argument that the increase in ad performance is caused by an increase in the specificity of ad personalization. More specific ad personalization works better when matching consumers’ browsing behavior more closely.

¹³During our experiment, consumers that browsed only category-level pages and were assigned to the product-specific personalization treatment were presented with one of the top three products within the browsed category. These consumers should perceive the personalization less relevant, compared to the consumers that are addressed with products that they actually viewed during product page browsing.

Table 2.5: Logistic Regressions for Click and Purchase Probabilities Controlling for Browsing Depth

VARIABLES	(1) Click	(2) Click	(3) Purchase	(4) Purchase
product-specific	0.646*** (0.026)	0.693*** (0.041)	0.960*** (0.118)	1.175*** (0.421)
social targeting		-0.101*** (0.030)		-0.562 (0.420)
product-specific × social targeting		-0.062 (0.041)		-0.290 (0.504)
product browsing	-0.032 (0.024)	-0.032 (0.024)	-0.106 (0.119)	-0.107 (0.119)
product-specific × product browsing	0.126*** (0.033)	0.129*** (0.033)	0.068 (0.157)	0.089 (0.157)
date controls	Yes	Yes	Yes	Yes
product category controls	Yes	Yes	Yes	Yes
country controls	Yes	Yes	Yes	Yes
device controls	Yes	Yes	Yes	Yes
Constant	-4.161*** (0.039)	-4.077*** (0.046)	-7.480*** (0.190)	-7.008*** (0.393)
Observations	1,457,527	1,457,527	1,457,527	1,457,527
Chi2	3,270.856	3,316.398	234.898	245.180
-2 Log Likelihood	168,824.907	168,779.364	11,862.406	11,852.123
AIC	168,898.907	168,857.364	11,936.406	11,930.123
BIC	169,350.020	169,332.862	12,387.519	12,405.621

Standard errors in parentheses, *** p<0.01, ** p<0.05, * p<0.1

When focusing on the influence of product browsing on purchase probabilities we see that product-specific personalization does lead to more positive consumer responses, consistent with our main results (see Table 2.5 column (3)). The interaction coefficient of product-specific personalization and product browsing is positive but not significant when estimating our models for purchase probabilities.

These models can also serve as a falsification test for whether simply a *reminder effect* instead of personalization of social retargeting is driving our results. While theoretically, personalization aims at matching a consumer's preferences as close as possible with an ad, a lot of retargeting algorithms simply advertise the consumer's last visited product. As the last visited product is the most recent memory of a consumer's browsing journey there might be a reminder effect of such an ad that encourages the consumer to pick up their shopping process where they left. In case such a reminder effect is driving our results, we would expect that advertising that matches consumers'

browsing behavior more closely, would always lead to more positive consumer responses. In contrast, we find that product-specific personalization also performs better for consumers that only browse category pages, as visible in the positive coefficient of the direct effect of product-specific personalization. This is a good indication that it is the specificity of personalization that leads to more positive consumer responses not a reminder effect.

2.5.3 Consumer Demographics

As our social targeting operationalization does not represent an experimental treatment variable, we run the risk that the effect that we find for social targeting is confounded by other factors, omitted from our analysis. A major concern is that consumers self-select into being connected to our partner company on Facebook. These consumers are by default socially targeted and see socially endorsed ads. Next to that, these consumers are also likely to be different from consumers that are not connected with our partner company on Facebook. This difference in unobserved consumer characteristics might drive the negative effect of social targeting on consumer responses to social retargeting ads. To investigate that, we make use of the fact that the Facebook advertising tool gives us access to additional demographic information on the consumers that we address with ads. More precisely, we make use of age and gender information of consumers in our analysis to see whether demographic factors, that are likely to be correlated with other unobserved consumer characteristics, change the coefficient of our social targeting variable. We re-run our main models for both click and purchase probabilities and include gender and age information as control variables.¹⁴

Table 2.6 gives the results of our analysis.¹⁵ For the model for click probabilities, we find that the results are consistent with our main models. We find that when controlling for age categories and gender, although decreasing in magnitude, the negative effect of social targeting and the interaction between product-specific personalization and social targeting are still negative and significant.

¹⁴For a limited number of consumers, Facebook does not have the gender or age information and returns the value unknown. We include these categories as baselines in our analysis, also because we are not interested in what the actual effect of age or gender is but whether there are other confounding factors that explain the effect of social targeting.

¹⁵For both logistic models the number of observations deviates slightly from the number of observations in our main models. This is the case as some age categories do perfectly predict that no click or purchase occurs. Observations for these categories are therefore excluded from the models.

Table 2.6: Logistic Regressions for Click and Purchase Probabilities Controlling for Consumer Characteristics

VARIABLES	(1) Click	(2) Purchase
product-specific	0.795*** (0.028)	1.330*** (0.397)
social targeting	-0.108*** (0.022)	-0.365 (0.415)
product-specific × social targeting	-0.067** (0.031)	-0.303 (0.487)
age 13-17	1.069 (1.034)	<i>omitted</i>
age 18-24	-0.506*** (0.167)	0.870 (1.086)
age 25-34	-0.539*** (0.166)	1.050 (1.084)
age 35-44	-0.246 (0.166)	1.219 (1.084)
age 45-54	0.036 (0.167)	1.341 (1.085)
age 55-64	0.290 (0.177)	1.342 (1.140)
female	0.035 (0.101)	-0.397 (0.413)
male	0.195* (0.101)	-0.015 (0.411)
date controls	Yes	Yes
product category controls	Yes	Yes
country controls	Yes	Yes
device controls	Yes	Yes
Constant	-3.931*** (0.137)	-8.237*** (1.070)
Observations	3,476,619	3,476,585
Chi2	7,367.108	502.428
-2 Log Likelihood	294,951.426	18,925.425
AIC	295,041.426	19,013.425
BIC	295,629.197	19,588.133

Standard errors in parentheses, *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

This strengthens our argument that the negative effect is actually driven by the conflicting effect of making the social targeting explicit in personalized advertising. When estimating the model for purchase probabilities, we find that the effects for social targeting and the interaction of social targeting with product-specific personalization are negative but non-significant. Although, this does not allow us to confirm our hypothesis that social targeting does negatively influence purchase probabilities, we do also find no support for a positive impact of social targeting on consumer responses to personalized ads as suggested in previous research (Bakshy et al., 2012; Tucker, 2016).

2.6 Discussion

With the growing availability of detailed online data on consumers and their online behaviors, opportunities in advertising personalization are constantly increasing. This research is among the first studies to investigate the effects of advertising personalization in social retargeting. More specifically, we investigate how consumers respond to personalization specificity in ads and whether their responses change when being socially targeted. We distinguish between two different levels of personalization specificity: category-specific and product-specific personalization. In collaboration with our partner company, we ran a field experiment to readdress consumers on Facebook after their website visits with these two types of personalized ads, using consumers' browsing behavior observed on the company's e-commerce website to infer consumer preferences. To assess consumers' reactions to ads we measure clicks and/or purchases.

Our results offer several important insights into how consumers respond to ad personalization. First, consumers generally respond more positively to product-specific than category-specific personalization. This finding is in line with the notion that highly personalized digital advertising matches consumer preferences more closely and is therefore perceived as more relevant by consumers (Arora et al., 2008). Our finding contradicts prior research that showed that less personalized generic brand ads perform better than more personalized ads, unless consumers have narrowed down their preferences and are close to making their purchase decision (Lambrecht and Tucker, 2013). We attribute the difference in findings to our cleaner experimental design that allows us to isolate the effects of personalization specificity unconfounded of other mechanisms. We make sure that our two ad treatments differ solely in the level of personalization specificity. Previous research, that showed that dynamically retargeted

ads perform worse than generic brand ads, potentially suffers from effects originating in the differences in ad designs. Less visually attractive dynamic retargeting ads or the fact that dynamic retargeting ads display several products making it more difficult for consumers to decide which product to click might cause dynamic retargeting ads to perform worse than generic brand ads.

To make sure it is the level of personalization specificity that drives our results, we rule out alternative explanations and further investigate the underlying mechanisms. In case a reminder effect was driving the finding that product-specific ads lead to more positive consumer responses, we would expect such ads to work better for only consumers that browsed actual product pages. These consumers would then be reminded of their last visited product and could pick up their purchase process where they left. However, we find product-specific ads also work better for consumers that browsed only the category pages. Furthermore, we find that both personalized ad treatments decrease in performance with an increase in the time between a consumers' website visit and an ad impression. This is because consumers develop their preferences over time rendering personalization based on previously observed consumer behavior less relevant (Simonson, 2005). We also find that highly personalized ads work better for consumers with more developed preferences, indicated by their browsing behavior. Both of these findings suggest that our ads are being perceived as personalized by consumers leaving us confident to have adequately operationalized ad personalization.

Second, we investigate the effects of social targeting on consumer responses to personalized ads. We find that, surprisingly, socially targeting consumers with personalized ads has a negative effect on consumer responses. This finding contradicts previous findings that suggested a positive impact of social targeting because of homophily of connected users and informational social influence (Bakshy et al., 2012; Tucker, 2016). These studies point out that (1) advertisers should be able to predict consumers' preferences more accurately as consumers' preferences should be similar to the preferences of their connected peers (homophily); and (2) displaying friends' names in an ad leads to more trust in the advertiser and a higher perceived ad relevance (social influence). However, our results show that in the context of personalized ads, social targeting in fact *negatively* influences consumer responses.

We explain this result using the uniqueness theory (Chan et al., 2012), which reconciles consumers' need to identify themselves with others (social identities) and their urge to be different with the goal to define their personal identity (Brewer, 1991; Snyder and Fromkin, 1980). The distinction we draw between social identity and personal identity in uniqueness theory is important in our context, which distinguishes

our work from earlier research in social advertising. The focus of our study lies on personalized advertising, rather than advertising in general. Research has pointed out the conflict between social identities and individuals self-concepts where the emphasis of social identities can lead to a depersonalization process (Brewer, 1991). We argue that through the inclusion of friends' names in the ad text, consumers are confronted with social identities that depersonalize ads presented to them. Seeing ads that recommend products specifically for consumers trigger an increase in perceived personalization. However, the inclusion of social identities via friend endorsements does conflict with personalized recommendations. Therefore, social endorsements decrease consumers' feeling that an ad has been uniquely personalized to their preferences leading to a lower propensity to positively respond to personalized ads.

If social targeting interferes with the personalization of advertising as it conflicts with consumers' perception of uniqueness, this conflict should be enhanced for more specific personalization. Thus, we also tested if social targeting negatively moderates personalization specificity. We empirically find that ads that are more specific in their personalization, and are therefore perceived as more unique, are more severely harmed by social targeting. Our results are robust when controlling for consumer demographics (age and gender).

2.6.1 Theoretical Contributions

Despite the increasing popularity of social targeting and ad personalization, studies investigating personalized advertising on social networks are scarce. While prior research has considered the respective effect of ad personalization (Bleier and Eisenbeiss, 2015a; Lambrecht and Tucker, 2013) and social targeting (Bakshy et al., 2012; Tucker, 2014, 2016) in isolation, we offer a first consideration of the combination of ad personalization and social targeting.

First, we advance the discussion on adequate levels of advertising personalization (Arora et al., 2008). By analyzing consumer responses to personalization specificity in digital ads, we challenge prior findings that claimed that less personalized ads lead to more positive consumer responses in most cases (Lambrecht and Tucker, 2013). We affirm that more specific ad personalization, that offers consumers a more specific product recommendation in personalized ads, leads to more positive consumer responses.

Second, we are the first to jointly investigate advertising personalization and social targeting. By investigating these advertising techniques jointly, we are able to identify

performance limitations caused by a combining these advertising techniques. This way, we are able to promote a better understanding of ad personalization and its limitations. We find that socially targeted consumers are less responsive to personalized ads, contradicting prior findings that assert the positive effect of social targeting (Bakshy et al., 2012; Tucker, 2016). Although social targeting has been considered in the advertising literature, no prior work to our knowledge has explored the combined effect of social targeting and ad personalization. We show that the use of friends' names to personalize ads has a negative effect on consumer responses to personalized ads. We also find that social targeting negatively moderates the effect of personalization specificity. That is, consumers react less positive to ads that recommend a specific product to them based on their preferences when they see that their friends endorse the ad, leading to the depersonalization of the ad.

2.6.2 Practical Implications

Advertising personalization enjoys an increasing popularity in the digital advertising industry with most marketers praising its higher response and engagement rates (EMarketer, 2015). However, advertisers struggle to find the optimal specifications for their personalized ads. Our results shed light on their question of how specific advertising personalization should be and inform ad platforms and policy makers.

First our research can help marketers to define their personalization strategy. Our research suggests that more specific ad personalization leads to more positive consumer responses. Furthermore, personalized advertising decreases in performance with an increase in the time between a consumers' website visit and the confrontation with a personalized ad. We recommend firms to retarget consumers with highly specific ad personalization as soon as possible after their website visits.

From a technical perspective, highly specific advertising personalization requires firms to implement more complex systems that are more costly to maintain compared to systems for less specific ad personalization (Zhang and Wedel, 2009). Although we cannot directly address the question of whether the implementation of such systems is financially beneficial for a firm, we can provide evidence that more specific ad personalization leads to more positive consumer responses, which could justify the financial investment for such systems.

Next, our findings challenge Facebook's default strategy to socially target consumers whenever possible. We find that socially targeting consumers decreases, rather than increases, advertising performance in the context of personalized ads. While ads that

are more specific and closer to consumers' preferences lead to more positive consumer responses, exploiting and using consumers' social connections seems to come at a cost. Currently, advertisers do not have the option to not include social endorsements in social ads. Our empirical evidence suggests that social networking ad platforms should reevaluate their policies and assess whether socially targeted ads do underperform untargeted ads contingent on the type of advertising, e.g. personalized advertising.

In light of recent events social networking sites, such as Facebook, should consider limiting their use of consumers' information for commercial purposes. Our results show that for personalized advertising, social targeting actually decreases ad performance, suggesting removing social endorsements from ads might be beneficial for both users and advertisers.¹⁶

2.6.3 Limitations and Future Research

Although our study provides valuable new insights into the personalization of digital advertising, some limitations need to be considered when interpreting our findings.

First, our study focuses on the comparison of different levels of advertising personalization specificity. Therefore, our results offer insights into which level of personalization specificity leads to more positive consumer responses. Future research can examine the overall effectiveness of personalized ads by comparing treated consumers with non-treated consumers or to consumers that are addressed with non-personalized ads.¹⁷ Such approach would represent the assessment of ad effectiveness more adequately as promoted by recent research (Johnson et al., 2017a,b; Lewis and Rao, 2014).

Second, we address consumers with ads that are personalized based on their browsing behavior. Therefore, only consumers that have indicated their interest in our partner firm's products, by visiting the firm's website, are eligible to be participants of our experiment. Although such a selection might raise concerns regarding to what extent our findings can be generalized out of sample, such a selection is common in the practice of advertising personalization. Previous work discussed the issue that consumers need

¹⁶In March 2018, the New York Times published an article on information from 87 million Facebook profiles being harvested for commercial purpose by data analytics company Cambridge Analytica (Meredith, 2018). Next to the media attention and concerns of social network users this led to an investigation by the Federal Trade Commission into whether Facebook had violated user privacy regulations.

¹⁷A major reason for not using non-personalized ads as the baseline in our experiment is that our partner company knew they had worse performance compared to personalized ads. Therefore, our partner company was not willing to invest marketing budget in such an ad treatment.

to be profiled before they can be addressed with personalized advertising (Thirumalai and Sinha, 2013).

Furthermore, such a selection of consumers has raised concerns regarding the presence of an activity bias that might confound ad effectiveness estimates (Lewis et al., 2011). The term activity bias describes that more active consumers are more likely to be addressed with ads. Such a selection leads to an upwards bias in ad effectiveness estimates as these more active consumers are more likely to respond positively to ads than the 'average' consumer. With respect to our study, we are likely to generally select more active consumers, consumers that visit our partner firm's website, into our sample. Furthermore, we are more likely to address more active consumers (more frequently) with ads as they are more likely to log into their Facebook accounts allowing us to serve ads to them. In our case, such an activity bias should influence both treatment groups, category- and product-specific ads, symmetrically, therefore not biasing our estimates in a way that limits how we can compare category- with product-specific ads. Nevertheless, we acknowledge that this activity bias might influence our out of sample generalizability as less active consumers are less likely to be part of our sample.

Third, it is worth noting that our social targeting operationalization does not represent a randomized treatment variable in our field experiment. As such we run the risk that unobserved consumer characteristics might bias our results. More specifically, consumers that are connected to the advertiser via a friend might be different from consumers that are not connected. Recent work has pointed out that unobserved consumer characteristics can heavily bias ad effectiveness estimates that are not derived from randomized experiments (Gordon et al., 2018). We try to remedy this issue by controlling for consumer demographics (i.e., gender and age). We find that results remain consistent with our main models, mitigating the concern of omitted variable bias. Future research could consider the type of relationship between a consumer and the 'friends' displayed as endorsers, and investigate the moderating role of tie strength between consumers and the friend(s) endorsing the ad more closely.

Fourth, it would be useful to understand the generalizability of the observed effects for different product categories. Our partner company is specialized in selling consumer electronics. Such search goods have lower consumption uncertainty compared to experience goods, making it easier for consumers to assess the quality of the product within an ad, and allowing firms to personalize ads more accurately. Previous research that studied personalized ads focuses on ads for experience goods such as sports fashion (Bleier and Eisenbeiss, 2015a) and holiday services (Lambrecht and Tucker,

2013). Future research could investigate to what extent the type of advertised product influences consumer responses.

Appendices

A2.1 Social Actions per Ad Type

We provide some descriptive statistics regarding the frequency of social interactions (shares, comments, likes) with the ads in Table A2.1 to give some insights on potential contamination of our experimental treatments. More precisely, by conducting a social interaction with an ad, this ad might potentially appear in the newsfeed of a consumer’s friend displaying this social interaction together with the ad. We have no numbers on how often advertising was displayed to individual consumers because of social actions. In line with our main analyses, in which we find product-specific ads to outperform category-specific ads, we also find that consumers more frequently interact with product-specific ads. Based on the numbers presented in the table we consider the problem of contamination as limited for the following reasons: (1) We find that the absolute numbers of social actions are low compared to our sample size of 3,476,626 impressions. (2) A social action on an ad is not sufficient to determine that the ad a consumer interacted with will be displayed to friends. This selection depends on the Facebook algorithm that determines which content to serve to consumers on the platform. (3) Given that an ad appears on a friend’s newsfeed, this friend needs to be part of the experiment for the ad to represent an actual contamination.

Table A2.1: Social Actions per Ad Type

	CS Ad	PS Ad	CS Ad [%]	PS Ad [%]	Total	Total [%]
Post Shares	28	38	0.0012%	0.0033%	66	0.0019%
Post Comments	44	104	0.0019%	0.0091%	148	0.0043%
Post Likes	904	1062	0.0386%	0.0934%	1966	0.0565%

CS = category-specific; PS = product-specific

A2.2 Economic Implications of Product-Specific Personalization and Social Targeting

Table A2.2: Economic Implications for Clicks

	Marginal Effect (%)	Change (%)	Absolute (pM)	Absolute (Exp.)
product-specific	0.61%	120.38%	6.12	21,425.25
social targeting	-0.10%	-13.03%	-0.96	-3,370.85
product-specific \times social targeting	-0.05%		-0.46	-1,592.50

Change (%) = the change in percentage compared to the baseline;

Absolute (pM) = absolute effect per 1,000 impressions;

Absolute (Exp.) = absolute effect for 3.5 million impressions in experiment

Table A2.3: Economic Implications for Purchases

	Marginal Effect (%)	Change (%)	Absolute (pM)	Absolute (Exp.)
product-specific	0.04%	213.92%	0.37	1,301.30
social targeting	-0.03%	-62.22%	-0.35	-1,224.65
product-specific \times social targeting	-0.00%		-0.02	-66.50

Change (%) = the change in percentage compared to the baseline;

Absolute (pM) = absolute effect per 1,000 impressions;

Absolute (Exp.) = absolute effect for 3.5 million impressions in experiment

A2.3 Data Structure Experiment

To be able to analyze the data from our experiment on an ad impression level, we need to structure our campaigns in a way that allows us to disentangle the different ad attributes in our analysis. In this section, we present a short summary of our data structure that will allow the reader to understand how we received information on such a granular level from the Facebook ad reporting tool. We are able to disentangle our aggregated data with the help of three features: (1) campaign set-up, (2) Facebook’s breakdown functionality, (3) Facebook’s reporting structure.

A2.3.1 Campaign Set-Up

We structured our campaigns in a way that allows us to distinguish between unique combinations of ad characteristics. More specifically we distinguish between eight features with a different number of attributes:

$$[Personalization\ Specificity] \times [Product\ Category] \times [Browsing\ Behavior] \times [Device] \\ \times [Temporal\ Distance] \times [Date] \times [Country]$$

This data structure leaves us with 22,400 unique combinations of ad features for which we know how many impressions were shown, how many clicks generated, and how many purchases conducted.

A2.3.2 Facebook's Breakdown Functionality

With the help of Facebook's breakdown functionality in the reporting tool we are able to get insights into aggregated consumer characteristics for our unique ad attribute combinations.

$$[Age\ Group] \times [Gender]$$

Facebook distinguishes between 8 age categories (including unknown) and 3 genders (including unknown). This leaves us, taking our unique ad attribute combinations into account, with a total of 537,600 unique attribute combinations.

A2.3.3 Facebook's Reporting Structure

In Facebook's reporting tool we can distinguish between social impressions and non-social impressions. This way we know how many impressions for a unique ad attribute combination were socially targeted of the overall number of impressions served for a specific ad attribute combination. Further, the tool gives us the number of overall clicks as well as the number of clicks on socially targeted ads. This way, we can distinguish between socially targeted and non-socially targeted ads as well as the respective number of clicks resulting from the unique ad attribute combination. For clicks we have therefore 1,075,200 potential unique combinations of ad attributes and their respective number of clicks. For purchases, the reporting tool does not distinguish between purchases resulting from socially targeted ads and purchases resulting from non-socially targeted ads. Therefore, we make use of the variation in the share of socially targeted ads per each of the 537,600 unique ad attribute combinations to assess the influence of social targeting on consumers' purchase probabilities.

A2.4 Consumer Response Likelihood Analysis on Consumer-Level

We analyze our data on a consumer- instead of impression-level to assess the robustness of our results. More precisely, our data allows us to infer unique consumer, campaign, day combinations. This means that we know how many unique consumers for each of our specific campaigns was addressed on a specific day of the experiment. The same consumers can be then addressed the following day, where we are not able to link these consumers to the consumers addressed the previous day. Therefore, our sample size for this analysis (2,794,878 consumer, campaign, day observations) is significantly larger than the number of consumers addressed in the experiment (198,234 individual consumers). Consumers are addressed with ads on several days of the experiment. Nevertheless, such analysis allows us to control for the fact that consumers might have seen several ads on a specific day of the campaign.

Table A2.4: Consumer-Level Logistic Regressions for Click and Purchase Probabilities

VARIABLES	(1) Click	(2) Purchase
product-specific	0.751*** (0.029)	1.142*** (0.381)
social targeting	-0.126*** (0.023)	-0.821** (0.390)
product-specific × social targeting	-0.067** (0.032)	-0.129 (0.466)
date controls	Yes	Yes
product category controls	Yes	Yes
country controls	Yes	Yes
device controls	Yes	Yes
Constant	-4.001*** (0.035)	-6.718*** (0.348)
Observations	2,794,878	2,794,878
Chi2	5515.568	415.684
-2 Log Likelihood	269905.689	18545.444
AIC	269979.689	18619.444
BIC	270454.891	19094.646
Standard errors in parentheses, *** p<0.01, ** p<0.05, * p<0.1		

A2.5 Likelihood to Click Estimates Using Unique Clicks

To assess the robustness of our main model for clicks, we replace our dependent variable for clicks with a different operationalization. We analyze the impact of our focal variables on *unique clicks*. Unique clicks do not count repeated clicks of consumers that click on the same ad several times but instead count such repeated clicks as single click instance.

Table A2.5: Logistic Regressions for Click Probabilities Using Unique Clicks

VARIABLES	(1) Click	(2) Click	(3) Click	(4) Click
product-specific		0.749*** (0.013)	0.749*** (0.013)	0.810*** (0.029)
social targeting			-0.171*** (0.016)	-0.132*** (0.023)
product-specific × social targeting				-0.077** (0.033)
date controls	Yes	Yes	Yes	Yes
product category controls	Yes	Yes	Yes	Yes
country controls	Yes	Yes	Yes	Yes
device controls	Yes	Yes	Yes	Yes
Constant	-4.162*** (0.029)	-4.434*** (0.030)	-4.294*** (0.033)	-4.326*** (0.035)
Observations	3,476,626	3,476,626	3,476,626	3,476,626
Chi2	2,253.963	5,393.747	5,499.897	5,505.379
-2 Log Likelihood	275,488.388	272,348.603	272,242.453	272,236.972
AIC	275,556.388	272,418.603	272,314.453	272,310.972
BIC	276,000.481	272,875.759	272,784.670	272,794.250

Standard errors in parentheses, *** p<0.01, ** p<0.05, * p<0.1

A2.6 Alternative Estimators for Click and Purchase Probabilities

Table A2.6: Alternative Estimators for Click and Purchase Probabilities

VARIABLES	(1) Probit Click	(2) LPM Click	(3) Probit Purchase	(4) LPM Purchase
product-specific	0.2930*** (0.0102)	0.0073*** (0.0003)	0.3283*** (0.1048)	0.0008*** (0.0002)
social targeting	-0.0492*** (0.0078)	-0.0008*** (0.0001)	-0.2630** (0.1043)	-0.0002** (0.0001)
product-specific × social targeting	-0.0315*** (0.0115)	-0.0017*** (0.0003)	-0.0396 (0.1282)	-0.0005** (0.0002)
date controls	Yes	Yes	Yes	Yes
product category controls	Yes	Yes	Yes	Yes
country controls	Yes	Yes	Yes	Yes
device controls	Yes	Yes	Yes	Yes
Constant	-2.1754*** (0.0126)	0.0160*** (0.0004)	-3.0891*** (0.0942)	0.0007*** (0.0001)
Observations	3,476,626	3,476,626	3,476,626	3,476,626
R-squared		0.0020		0.0001
Chi2	6,131.8350		438.7122	
-2 Log Likelihood	296,108.3296		18,989.1659	
AIC	296,182.3296		19,063.1659	
BIC	296,665.6078		19,546.4441	

Standard errors in parentheses, *** p<0.01, ** p<0.05, * p<0.1

LPM = linear probability model

A2.7 Estimates for Post-View Purchase Probability

As robustness check, we replace our dependent variable for purchases with a different operationalization. In our main analysis, we operationalize purchases as purchases contingent on having clicked on an ad that are conducted within 28 days after the click. In the robustness check below, purchases represent purchases that occur within 7 days after having been exposed to an ad.

Table A2.7: Logistic Regressions for Purchase Probabilities Using Post-View Purchases

VARIABLES	(1) Purchase	(2) Purchase	(3) Purchase	(4) Purchase
product-specific		0.251*** (0.029)	0.249*** (0.029)	0.167 (0.177)
social targeting			-0.459*** (0.108)	-0.509*** (0.150)
product-specific × social targeting				0.101 (0.215)
date controls	Yes	Yes	Yes	Yes
product category controls	Yes	Yes	Yes	Yes
country controls	Yes	Yes	Yes	Yes
device controls	Yes	Yes	Yes	Yes
Constant	-6.411*** (0.081)	-6.490*** (0.081)	-6.108*** (0.121)	-6.068*** (0.148)
Observations	3,476,626	3,476,626	3,476,626	3,476,626
Chi2	635.752	710.709	728.232	728.453
-2 Log Likelihood	76,642.014	76,567.058	76,549.535	76,549.314
AIC	76,710.014	76,637.058	76,621.535	76,623.314
BIC	77,154.107	77,094.213	77,091.751	77,106.592

Standard errors in parentheses, *** p<0.01, ** p<0.05, * p<0.1

A2.8 Robustness Check: Average Reach of Ad Attribute Combination

As we analyze our data on impression-level and do only have access to user information in an aggregated manner we are not able to directly control for how many impressions are served for a specific consumer. Nevertheless, the frequency of impressions is an important control measure as a higher frequency of impressions per user for any of the ad features we are focusing on might bias our results. A higher frequency per

consumer for a certain attribute is likely to lead to a decrease in the effectiveness per ad impression for this respective attribute. Therefore, we generate a variable that measures the average reach for each of our unique ad attribute combinations. This variable indicates the number of consumers reached divided by the impressions served for each unique ad attribute combination. In case each impression for an ad attribute combination reaches a unique consumer the variable takes the value 1. We include this variable into our analysis to control for potential issues due to an unequal average impression frequency for our personalization-specificity treatment as well as our social targeting variable. Our analysis yields results consistent with our main models (see Table A2.8 and Table A2.9).

Table A2.8: Logistic Regressions for Click Probabilities Controlling for Average Reach of Ad Attribute Combination

VARIABLES	(1) Click	(2) Click	(3) Click
product-specific	0.646*** (0.013)	0.649*** (0.013)	0.704*** (0.028)
social targeting		-0.116*** (0.016)	-0.081*** (0.022)
product-specific × social targeting			-0.070** (0.031)
reach	2.334*** (0.062)	2.264*** (0.063)	2.265*** (0.063)
date controls	Yes	Yes	Yes
product category controls	Yes	Yes	Yes
country controls	Yes	Yes	Yes
device controls	Yes	Yes	Yes
Constant	-6.092*** (0.056)	-5.943*** (0.059)	-5.972*** (0.061)
Observations	3,476,626	3,476,626	3,476,626
Chi2	7,458.736	7,512.531	7,517.682
-2 Log Likelihood	294,859.901	294,806.107	294,800.956
AIC	294,931.901	294,880.107	294,876.956
BIC	295,402.118	295,363.385	295,373.295

Standard errors in parentheses, *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table A2.9: Logistic Regressions for Purchase Probabilities Controlling for Average Reach of Ad Attribute Combination

VARIABLES	(1) Purchase	(2) Purchase	(3) Purchase
product-specific	0.938*** (0.064)	0.941*** (0.064)	1.118*** (0.376)
social targeting		-0.663*** (0.215)	-0.510 (0.387)
product-specific × social targeting			-0.221 (0.460)
reach	3.873*** (0.348)	3.727*** (0.352)	3.732*** (0.352)
date controls	Yes	Yes	Yes
product category controls	Yes	Yes	Yes
country controls	Yes	Yes	Yes
device controls	Yes	Yes	Yes
Constant	-10.661*** (0.310)	-10.002*** (0.379)	-10.131*** (0.465)
Observations	3,476,626	3,476,626	3,476,626
Chi2	545.218	554.250	554.481
-2 Log Likelihood	18882.660	18873.628	18873.397
AIC	18954.660	18947.628	18949.397
BIC	19424.877	19430.906	19445.737

Standard errors in parentheses, *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Social Influence and Visual Attention in the Personalization Privacy Paradox: An Eye Tracking Study¹

3.1 Introduction

Personalization of marketing communication has been praised by advertising platforms for its positive implications for advertising performance (Adobe Systems Inc., 2014). By increasing availability and use of individual-level consumer data, advertising can be personalized to match individual consumer preferences. Providers of personalization technology point towards consumers' positive reactions to advertising personalization, such as increases in click and purchase probability, caused by an increase in ad relevance. However, research shows that consumers are concerned about the use of their information to personalize advertising content (Bleier and Eisenbeiss, 2015b; Sutanto et al., 2013; Tucker, 2014). Consumers experience information privacy concerns, which negatively influence consumer responses to ads, as third parties use their information

¹Earlier versions of this study appeared in the following conference proceedings or were presented at the below mentioned conferences and workshops:

- Frick, T.W., Li, T. & Pavlou, P. (2017). The Role of Social Influence and Attention in the Personalization Privacy Paradox: An Eye Tracking Study. In *Statistical Challenges in eCommerce Research Symposium (SCECR)*, Hanoi, Vietnam.
- Frick, T.W., Li, T. & Pavlou, P. (2016). Investigating the Impact of Social Influence on the Personalization-Privacy Paradox: An Eye Tracking Study. In *INFORMS Annual Meeting*, Nashville, US.
- Frick, T.W. & Li, T. (2015). Understanding Information Privacy Concerns in Social Advertising: An Eye Tracking Study. In *37th ISMS Marketing Science Conference*, Baltimore, US.

for personalization purposes without being in control of their own data. “Information privacy refers to the desire of individuals to control or have some influence over data about themselves” (Bélanger and Crossler, 2011, p. 1017). Consumer privacy concerns describe that consumers are concerned that they are not in control over the use of their personal data.

Consumer privacy concerns tend to be especially present in contexts in which consumers question the value that they gain from personalization and where they perceive that third parties aim to monetize on their information, which is common in advertising (Sutanto et al., 2013). This leads to the challenge for firms to adequately balance an increase in ad relevance with the simultaneous increase in consumer privacy concerns caused by ad personalization to avoid negative reactions to ads from consumers. The simultaneous increase in ad relevance and consumer privacy concerns through personalization has been coined personalization privacy paradox (Awad and Krishnan, 2006; Sutanto et al., 2013). While consumers’ valuation of marketing content increases as they are exposed to more relevant ads, they also experience higher privacy concerns as the personalization comes at the cost of their personal data.

In contexts that are rich of individual-level data and therefore attractive for personalization purposes, such as social networking sites, consumers are likely to be addressed with personalized ads that trigger privacy concerns (Tucker, 2014). Social advertising allows firms to personalize ads based on demographics, interests, browsing behavior, and, peculiar to social advertising, social connections. In social ads, the social connections used to personalize ads are made explicit as social endorsements. By including social endorsements in social ads, which are information cues on consumers’ friends that like the advertising firm, firms try to trigger informational social influence. Informational social influence describes that consumers are influenced in their decision making by information about the actions of their peers (Kwahk and Ge, 2012). For example, when a consumer sees that one of her friends likes a movie on social media, this information will influence the consumer’s decision to watch this movie. The information that the friend likes the movie is used to infer its quality. Further, informational social influence leads individuals to perceive externally received information to be true (Deutsch and Gerard, 1955). Research has pointed out that through informational social influence consumers perceive ads as more relevant and credible (Bakshy et al., 2012; Tucker, 2016).

Although firms can make use of more detailed consumer profiles in social advertising that improve personalization of ads, these ads are, due to incorporation of detailed consumer data, also more likely to induce consumer privacy concerns. Despite

the increasing popularity of social advertising, little is known about how it affects consumers' privacy concerns.

While the personalization privacy paradox has been investigated in the advertising context before (Aguirre et al., 2015; Bleier and Eisenbeiss, 2015b; Lee et al., 2011; Sutanto et al., 2013), we seek to contribute to the literature by investigating whether consumers' information privacy concerns can be mitigated through informational social influence in social advertising. As informational social influence creates trust in the advertising firm, consumers might display a higher acceptance for the use of their personal information for personalization purposes.

To get a deeper understanding of consumers' cognitive processes when evaluating ads, we assess the role of visual attention in the personalization privacy paradox by recording consumers' eye movements. By measuring visual attention, which has been identified as key coordination mechanism for information processing (LaBerge, 1995), we can identify which information is being processed by an individual when evaluating an ad (Wedel and Pieters, 2007).

While attention has been mostly viewed as an enabler for ad performance (Lee and Ahn, 2012; Wedel and Pieters, 2007), research in the area of psychology points towards the presence of a negativity bias directing consumers' attention (Fiske, 1980; Kanouse, 1984; Smith et al., 2003). The negativity bias in visual attention describes that individuals are more likely to dedicate attention to stimuli that cause negative emotions. By measuring both consumer privacy concerns, which have been shown to cause negative emotions such as vulnerability (Aguirre et al., 2015), and visual attention, we are able to get unique insights into the presence of such a negativity bias for personalized advertising. This investigation is especially important as it challenges the generally positive role of visual attention in advertising. The presence of such a negativity bias directing consumers' attention when exposed to personalized ads would lead to higher attention towards ads that trigger higher consumer privacy concerns. Visual attention would then not necessarily have positive implications for consumer responses to ads but might negatively influence ad performance.

Our study aims to answer two research questions related to the personalization privacy paradox in the context of advertising: (1) Can informational social influence mitigate the increase in consumer privacy concerns caused by the personalization of ads? (2) How does consumers' attention, which is required to cognitively and affectively process information in an advertisement, influence ad performance in the context of the personalization privacy paradox?

To answer these research questions, we conducted a multi-method study using a set of lab experiments, eye tracking technology, and questionnaires, to investigate how social advertising affects consumers' privacy concerns by shaping their affective and cognitive reactions. Consumers' attention patterns give an excellent view on which information of a marketing stimuli consumers process and present a basis to predict consumers' reactions (Wedel and Pieters, 2007).

For the experiments, we created different social advertising conditions in a lab environment resulting in a $2 \times 2 \times 2$ between-subject experimental design. We randomly assigned participants to being confronted with (1) personalized or non-personalized ads, (2) which were socially endorsed or not, and (3) advertised two different types of products. We constructed personalized advertising based on experimental scenarios and dynamically include participants' actual social connections to socially endorse ads. We measure consumer responses to the different ad conditions by investigating whether participants in the experiment click on the ad presented to them.

We find that informational social influence leads to more positive consumer responses to ads, measured as clicks. Nevertheless, informational social influence does not mitigate privacy concerns triggered by the personalization of advertisements. Further, our study disentangles the attentional processes within the personalization privacy paradox. We find that personalization positively affects consumer responses to ads, partially mediated by an increase in attention towards the ad. At the same time, our analysis of serial mediation reveals that there is a significant negative indirect effect as personalization increases privacy concerns which lead to a decrease in attention, eventually decreasing advertising performance. This finding supports the notion of a generally positive effect of attention on advertising performance and refutes the idea that firms can leverage a negativity bias in attention to attract consumers' eye balls.

Our results offer valuable insights for marketing practitioners: First, while informational social influence cannot mitigate privacy concerns triggered by personalized ads, it directly affects ad performance. While former studies claimed that personalization of ads does not increase attention compared to non-personalized ads for consumers with low cognitive demand tasks (Bang and Wojdyski, 2016), we show that personalized ads significantly increase attention, even for low cognitive demand tasks, such as browsing on social networking sites. Notably, we show that intrusive ads, which cause privacy concerns, will eventually decrease ad performance via a reduction in attention.

By taking a multi-method approach, we contribute to the IS literature shedding light on the personalization privacy paradox in an advertising context. Our study differs from former work on the personalization privacy paradox as we observe actual privacy

concerns of consumers originating in confrontation with personalized ads. Most prior studies focus on consumers' willingness to be profiled for the personalization of services prior to an actual confrontation with a personalized service (Awad and Krishnan, 2006). We acknowledge a clear difference between the assessment of consumers' privacy concerns *ex ante* and *ex post* the confrontation with a personalized service.

While several factors that mitigate privacy concerns in a personalization context such as justification of personalization (White et al., 2008), trust (Bleier and Eisenbeiss, 2015b; Chellappa and Sin, 2005), control over personal information (Song et al., 2016; Tucker, 2014), and communication media (Aguirre et al., 2016) have been investigated, we are the first to assess the role of informational social influence in mitigating privacy concerns caused by personalization. Although theory points towards a presence of such a mitigation, we do not find empirical support for a moderating effect. Informational social influence does not help marketers to overcome consumer privacy concerns in personalized ads. This might be the case as the inclusion of friend connections in ads represents a use of personal information in itself, triggering privacy concerns that counterbalance the moderating effect of informational social influence.

To the best of our knowledge, we present the first study to empirically test the simultaneous influence of personalization and privacy concerns on consumers' attention and the resulting effects on consumer responses to ads. We theoretically derive an explanation for both a potential positive and negative effect of consumer privacy concerns on attention towards a personalized ad. While research in psychology points towards the presence of a negativity bias directing consumers' attention, suggesting an increase in attention with an increase in privacy concerns, we find that privacy concerns do negatively influence consumers' attention. This finding is in line with the notion of attention being an enabler for ad performance and underlines the positive role of visual attention in advertising. While consumers' attention is limited and different information stimuli are competing for consumers' attention simultaneously (Desimone and Duncan, 1995), we find that more intrusive ads cannot help firms to gain more of consumers' attention. Although more intrusive ads that increase consumers' privacy concerns might initially generate more attention towards an ad, the effect of privacy concerns on attention overall is negative, leading to a decrease in ad effectiveness for more intrusive ads.

3.2 Related Literature

The investigation of the performance of personalized advertising dependent on advertising characteristics and consumers' perceptions of these characteristics has a long history in marketing research (Arora et al., 2008). Related to that, several studies focus on the cognitive processes within consumers that determine their reaction towards ads (Pieters et al., 1999, 2010; Pieters and Wedel, 2012; Rosbergen et al., 1997; Wedel and Pieters, 2000). Eye tracking methodology allows researchers to get an indication of which information consumers process cognitively, as attention towards an information element represents a prerequisite for cognitive processing (Wedel and Pieters, 2007).

Early research in the area of visual marketing investigated the predictive power of eye fixation data on consumers' ability to distinguish different print ads and purchase probabilities (Treisman and Gregg, 1979). Leven (1991) found that consumers tend to scan print ads before investigating them in detail. Rosbergen et al. (1997) point out the importance of consumer heterogeneity when investigating consumer attention. Consumers differ in the attention patterns they show towards ads. This finding is especially crucial, as it shows that not only bottom-up, ad design features, but also top-down factors, consumer characteristics, determine attention patterns. Related studies showed that consumers' goals significantly influence consumers' attention patterns (Pieters and Wedel, 2007).

Visual attention has also been found to represent an important mechanism when trying to understand ad effectiveness and wear-out, where ads perform less well with an increase in the number of a consumers' exposures to an ad (Pieters et al., 1999). While fixations on an advertisement's brand and image increase brand memory, advertising text has been found to not influence brand memory (Wedel and Pieters, 2000).

Another focal topic within the area of visual marketing is banner blindness. Banner blindness describes consumers' habit of training themselves to ignore digital banner advertisements, therefore reducing ad banners' ability to distract consumers from their focal task (Drèze and Hussherr, 2003). Consumers' attention patterns when confronted with banner ads challenge the commonly used click-through success measure as their attention leads to an increase in brand awareness independent of a click-through. Researchers found that most consumers do actually fixate on banner ads at least once and memorize ads that match website content better (Hervet and Gue, 2011). Other research finds that animations in banners do decrease consumers' attention as well as the positive effect of attention on memory (Lee and Ahn, 2012).

3.3 Theory

Because of the increase in the popularity of ad personalization technologies and the intertwined requirement for granular consumer data, the topic of consumer privacy has moved into the focus of both researchers and policy makers. Concurrent to the increase in ad relevance through personalization, consumers are concerned that they lose control over their personal information that is being used to adjust ads to their preferences. Although advertising effectiveness seems to be positively influenced by ad personalization, privacy concerns decrease advertising effectiveness creating the necessity for firms to balance these two effects. In this section, we will provide the underlying theoretical concepts for our study and develop the conceptual model that guided our research. We will start by discussing advertising personalization and the personalization privacy paradox. Next, we will introduce the concept of informational social influence and how informational social influence might mitigate the effect of ad personalization on consumer privacy concerns. Lastly, we will explain the role of consumer attention related to cognitive processes that determine consumers' reactions towards ads. Firms are increasingly focusing on measuring consumers' reactions to ads, which we label *consumer response* in our study, to assess advertising performance.²

3.3.1 Advertising Personalization

Advertising personalization describes the firm-initiated adjustment of advertising content towards the preferences of consumers (Arora et al., 2008) with the ultimate goal to positively influence consumers' perceptions and make them conduct business with the advertiser (Ansari and Mela, 2003). Advertising content is adjusted to match the preferences of either a consumer segment (one-to-n personalization) or an individual consumer (one-to-one). Advertising personalization has been studied in varying contexts such as e-Mail marketing (Sahni et al., 2018; Wattal et al., 2012; White et al., 2008), banner ads (Bleier and Eisenbeiss, 2015a; Lambrecht and Tucker, 2013), social media advertising (Bakshy et al., 2012; Tucker, 2014, 2016), and mobile advertising (Kim and Han, 2014). Prior research mostly points out the positive implications of advertising personalization on ad performance. In the context of e-Mail marketing, research finds that individually personalizing e-Mails by including consumers' names in the subject line positively influences open rates, sales leads, purchases, and consumer retention while decreasing the number of consumers unsubscribing from an e-Mail

²In our analysis, we operationalize *consumer response* as consumers' propensity to click on an ad.

list (Sahni et al., 2018). Bleier and Eisenbeiss (2015a) show that highly personalized banner ads do outperform less personalized ads. In the mobile advertising context, personalization is found to have a positive relationship with ad informativeness, credibility, and the entertaining aspects of an advertising message (Kim and Han, 2014). In line with former findings we hypothesize:

Hypothesis 1 (H1): *Personalization of ads leads to more positive consumer responses to ads.*

Although personalization has been found to have a positive impact on advertising effectiveness through an increase in advertising relevance, research has also revealed the negative implications of ad personalization through an increase in information privacy concerns. “Information privacy refers to the desire of individuals to control or have some influence over data about themselves” (Bélanger and Crossler, 2011, p. 1017). To successfully personalize advertisements, advertisers require not only consumer data to infer preferences but also consumers’ willingness to be profiled and to adopt services that are personalized to their preferences (Chellappa and Sin, 2005). Highly personalized advertisements require detailed information on consumers to adjust ad content adequately to consumers’ preferences (Arora et al., 2008). Consumers’ feelings of losing control over their private information that is being used by a third-party causes information privacy concerns (Sutanto et al., 2013). Aguirre et al. (2015) show that consumers feel vulnerable when realizing that advertisers collected their information to personalize ads without their prior consent. This feeling of vulnerability leads to negative consumer responses for personalized ads. The conceptual conflict of personalization inducing both higher ad relevance and consumer privacy concerns is coined personalization privacy paradox (Awad and Krishnan, 2006; Sutanto et al., 2013).

3.3.2 Personalization Privacy Paradox

Although personalized content in digital advertising receives higher levels of attention it also increases privacy concerns of users (Malheiros et al., 2012). Personalized digital advertising faces the constant trade-off between relevance gain and privacy cost in the form of personalization reactance - “that is, psychological resistance in response to subjectively inappropriate personalization” (White et al., 2008, p. 40). Personalization technologies rely on consumers’ personal information to match content to their preferences (Sutanto et al., 2013). Nevertheless, by choice consumers tend to

restrict third parties' access to their personal information (Utz and Kramer, 2009). The utilization of personal information by third parties leads to consumer privacy concerns. Consumers have been found to be especially privacy sensitive about their personal information in an advertising context (Awad and Krishnan, 2006). This can be explained by the fact that personalized advertising is being perceived as less beneficial compared to other personalized services, e.g. personalized music play lists. In the context of e-Mail marketing, research found that consumers show personalization reactance when confronted with highly personalized marketing communication which is not justified by the advertiser (White et al., 2008) and that increasing the level of personalization leads to higher perceived privacy risks of consumers (Song et al., 2016). We hypothesize:

Hypothesis 2 (H2): *Personalization of ads increases consumer privacy concerns.*

When looking at the impact of consumer privacy concerns, research found that privacy concerns decrease consumers' willingness to adopt personalized services (Sheng et al., 2008; Song et al., 2016). When comparing users' willingness to save personalized offers from a mobile application, Sutanto et al. (2013) find that consumers are more willing to do so in a privacy-safe version of the application. Privacy calculus theory describes, that when deciding on whether to release personal information consumers compare the costs in form of privacy concerns with the benefits gained from accessing a service (Xu et al., 2011). When privacy concerns are too high they inhibit consumers to adopt a service. Therefore, privacy concerns represent a cost that is factored into consumers' decision whether to adopt personalized services (Chellappa and Sin, 2005). When privacy costs are higher than the benefits gained from a personalized advertisement consumers will react negatively. In the context of e-Mail marketing, personalization reactance was found to decrease consumers' willingness to accept offers (White et al., 2008). Next to that, privacy concerns of consumers have been found to increase ad avoidance (Baek and Morimoto, 2012) and to decrease consumers' intentions to click on banner ads (Bleier and Eisenbeiss, 2015b). In line with former findings we hypothesize:

Hypothesis 3 (H3): *Consumer privacy concerns lead to negative consumer responses to ads.*

3.3.3 Informational Social Influence

Social Influence has been found to influence both consumers' product evaluations and decisions (Yadav et al., 2013). Normative social influence describes what is commonly known as peer pressure, where individuals feel social pressure to follow the decisions of their peers to belong to and not conflict with a certain social group. On the other hand, informational social influence describes the informational signal that consumers receive from their peers that is taken into account when evaluating their environment (Burnkrant and Cousineau, 1975). Informational social influence helps consumers to evaluate aspects of their environment, as information inferred from their peer's association with these aspects reduces uncertainty (Park and Lessig, 1977).

In this paper, we focus on peers' association with a commercial product advertised to consumers and how this association is affecting consumers' responses to ads. Consumers have the tendency to believe that their peers made product decisions based on better information and are therefore inclined to follow their peers' decisions (Bonabeau, 2004). Informational social influence leads individuals to evaluate information received from third-parties as trustworthy (Deutsch and Gerard, 1955). With increasing uncertainty, the influence of peers increases (Liu and Sutanto, 2012). It is assumed that informational social influence plays a role for a wide range of product categories (Yadav et al., 2013). In social advertising, advertisers try to tap into the effect of informational social influence by making the association of consumers' peers with advertisers explicit to consumers. This allows consumers to decrease uncertainty regarding the advertised product as they learn that their connected peers are in a favorable relationship with the advertising brand. In social advertising, informational social influence, operationalized as social endorsements, has been found to positively influence consumers' responses to ads (Bakshy et al., 2012; Tucker, 2014, 2016). We follow this argumentation and hypothesize:

Hypothesis 4 (H4): *Informational social influence leads to more positive consumer responses to ads.*

3.3.4 Informational Social Influence as Mitigation of Privacy Concerns

Research in the area of the personalization privacy paradox has investigated potential moderators that influence the strength of the relationship between personalization and consumer privacy concerns. According to White et al. (2008), personalization

reactance can be mitigated by justification of personalization and perceived utility of the personalized service. Essentially, when consumers perceive the use of their personal data as justified as it increases the quality of the service deploying the personalization, they react less negative to the use of their data. Next to that, trust in the advertiser has been shown to alleviate the effect of personalization on privacy concerns (Bleier and Eisenbeiss, 2015b; Chellappa and Sin, 2005). Furthermore, when consumers' perception of control over their personal information increases, personalization induces lower levels of privacy concerns (Song et al., 2016; Tucker, 2014). Moreover, the communication medium through which the personalized communication is transmitted influences the strength of the relationship between personalization and privacy concerns (Aguirre et al., 2016).

In our study, we investigate whether informational social influence can mitigate the relationship between advertising personalization and consumer privacy concerns. In line with previous research, we argue that informational social influence leads to consumers evaluating advertising more positively (Bakshy et al., 2012; Tucker, 2016). Consumers evaluate information received from third-parties as more trustworthy through informational social influence (Deutsch and Gerard, 1955). This increase in trust is likely decreasing consumers' perceived privacy risks (Lee and Rha, 2016) and feeling of vulnerability due to an invasion of their privacy (Aguirre et al., 2015). Both perceived privacy risks (Lee and Rha, 2016) and trust in the advertiser (Bleier and Eisenbeiss, 2015b) have been found to moderate the effect of personalization on consumer privacy concerns.

Next to that, consumers are particularly affected by informational social influence if it increases their understanding or helps them to deal with a problematic situation or conflict (Burnkrant and Cousineau, 1975). This means that consumers are especially prone to being influenced by others if they are facing a conflict for which they need to find a solution. We argue that consumers are confronted with such a conflict when facing the personalization privacy paradox induced by a personalized ad. Although, consumers perceive the presented ad as more relevant they are concerned about their privacy and that their personal information is being used without their consent. With the aim to resolve this conflict, consumers are likely to process the information that their friends seem to be in a favorable relationship with the advertiser. Such information should decrease the effect of personalization on consumer privacy concerns. We hypothesize:

Hypothesis 5 (H5): *Informational social influence mitigates the effect of personalization on consumer privacy concerns.*

3.3.5 Consumer Attention

Eye movements give a good indication of what information consumers are processing as information acquisition requires consumers to direct the fovea, the central part of the eye, towards the information that is being processed (Wedel and Pieters, 2007). By identifying eye fixations, during which an individual's eye is moving only very little, researchers can learn which information is being processed by an individual. This is important as the acquired information influences consumers' preference development and eventually choice. Attention has been pointed out as key coordination mechanism for information processing (LaBerge, 1995). Attention for information can either be captured bottom-up by stimuli or directed top-down guided by an objective of the individual (Greenberg, 2012).

On a social networking site, consumers tend to freely browse their newsfeed in a non-goal-directed but rather exploratory browsing mode, also described as low cognitive demand task (Bang and Wojdyski, 2016), leading to a smaller amount of top-down directed attention. At the same time, different posts in the newsfeed area of a users' social network feed, including news postings, friend postings, and advertisements, are difficult to be distinguished. This means that bottom-up bias in attention direction is minimal (Desimone and Duncan, 1995). This approach makes sense for the social network site acting as ad platform, as consumers have been found to train themselves to ignore easy to distinguish advertisements that inhibit stronger bottom-up attention attraction, referred to as banner blindness (Lee and Ahn, 2012). This leads to consumers evaluating on the spot to which content area they want to dedicate their attention. Personalization acts here as a top-down guidance mechanism for consumers' attention allocation by biasing their attention towards the content perceived as more relevant (Desimone and Duncan, 1995).

Research has shown, personalized content is more likely to catch consumers' attention and allows individuals to remember the advertising content as it entails personal relevance and is easier to process for individuals (Köster et al., 2015). This is also the case if personalization, as top down factor, is interfering with another top down factor, such as a search task. Especially, when consumers are conducting high demanding cognitive tasks, personalized ads have been shown to be particularly more effective than non-personalized ads in catching consumers' attention (Bang and Wojdyski,

2016). Consumers have been found to spend twice as much time on ads that are personalized, including their picture, than on non-personalized ads (Malheiros et al., 2012).

Furthermore, web personalization, in the context of recommendation systems, has been found to increase consumer engagement and attention which is required to evaluate the presented content (Tam and Ho, 2005). When including name and geographic information in ads in order to personalize them, researchers found a significant increase in attention towards these ads compared to non-personalized ads (Bang and Wojdyski, 2016). In line with former findings we hypothesize:

Hypothesis 6 (H6): *Personalization increases attention towards ads.*

We know from eye tracking research that consumers' attention directed to an advertisement is necessary to cognitively process the information in an ad (Lee and Ahn, 2012). While personalization acts as a top-down enabler, increasing the likelihood that consumers spend attention on an ad, attention in itself is the basis for consumers to form their opinions about an ad. Attracting consumers' attention is especially crucial in situations in which several stimuli are competing for consumers' attention simultaneously as described in the biased competition theory (Desimone and Duncan, 1995) and common in the digital advertising space. When consumers decide to dedicate attention to a specific stimulus, this means that less attention can be spent on other stimuli simultaneously as consumers' attention is naturally limited.

This argumentation underlines the importance of attention in the information processing stage, enabling the behavior stage (e.g. clicking, purchasing, etc.) (Köster et al., 2015). Former research has shown that attention measures captured from eye tracking are predictive of future sales (Wedel and Pieters, 2007). Further, attention has been found to increase consumers' elaboration of advertising content (Tam and Ho, 2005). We hypothesize that attention positively influences consumers' responses to ads:

Hypothesis 7 (H7): *Attention has a positive impact on consumers' responses to ads.*

While theory suggests that personalization is increasing consumer attention to ads, the relationship between consumer privacy concerns and consumers' attention directed towards an advertisement is less trivial. Although personalization increases the degree to which consumers evaluate ad content by increasing their attention to the ad, this increase in cognitive processing does not necessarily lead to a positive response

to an ad (Tam and Ho, 2005). Related research finds that when personalizing a holiday booking task in the lab and tracking participants' eye movements, participants' attention towards ads increases while consumers have simultaneously higher privacy concerns (Köster et al., 2015). It remains unclear what the relationship between consumer privacy concerns and attention is. Both a positive and negative effect of privacy concerns on attention can be theoretically explained. Research in the area of psychology has shown that individuals' moods influence their attention patterns (Wadlinger and Isaacowitz, 2006).

On the one hand, privacy concerns might increase consumers' attention towards an ad as consumers want to figure out who is using their personal information and for what reason. In the case of personalized advertising, this would lead to two conflicting effects of attention. Privacy concerns would increase consumers' attention inducing cognitive processing that fosters negative feelings related to a loss in privacy. At the same time, personalization would increase the attention that fosters positive processing of the ad content that is perceived as relevant. In such a case, the overall effect of attention on consumers' responses to ads depends on whether privacy concerns or personalization represent the stronger top-down factor driving attention.

Malheiros et al. (2012) find, although not testing the actual effect of privacy concerns on consumer attention, that personalized ads increase the attention significantly but also increase consumer privacy concerns, triggering a more negative perception of ads. This points towards a potential positive effect of privacy concerns on consumer attention, where consumer attention might negatively influence consumers' responses to ads.

Research in the area of psychology has pointed towards the presence of what is called negativity bias, also negativity effect, in information integration, which is the process of forming one's overall judgment regarding an object of interest (Kanouse 1984). Individuals tend to suffer from selective attention mechanisms that favor negative information over positive information (Fiske, 1980). The negativity bias is basically describing that "our attention is automatically drawn to negative information more strongly than it is automatically drawn to positive information" (Smith et al., 2003, p. 171). Recent research showed that in a second screen setting, consumers dedicate more attention to social media messages when these messages are negative than when confronted with positive messages, supporting the presence of a negativity bias (Kätsyri et al., 2016). In line with the presence of a negativity bias we hypothesize:

Hypothesis 8a (H8a Competing): *Privacy concerns increase the attention towards an ad.*

Assuming that consumer attention increases with privacy concerns, challenges the positive implication of attention on consumers' responses to ads. When privacy concerns, that reduce ad effectiveness (Aguirre et al., 2015; Bleier and Eisenbeiss, 2015b), increase consumer attention to ads, this points towards the presence of a negative type of attention. Former research has shown that attention is required to cognitively process ad content and evaluate advertising, whether positively or negatively (Tam and Ho, 2005; Lee and Ahn, 2012).

At the same time, research has pointed out that attention is a positive enabler for ad performance (Wedel and Pieters, 2007). Next to that, consumers have been found to train themselves to avoid paying attention to banner ads that they evaluate unfavorably, which points towards consumers trying to avoid negatively perceived ad information (Drèze and Hussherr, 2003). These findings support the notion of a negative effect of privacy concerns on attention, i.e. consumers' negative feelings such as perceived vulnerability (Aguirre et al., 2015) would actually decrease consumers' attention towards an ad. The assumption of a negative effect of privacy concerns on attention is in line with the notion of a positive effect of higher levels of attention on advertising performance. While the negative emotions caused by privacy concerns draw attention away from an ad, personalization, which leads to more relevant ad content, increases attention supporting the notion of a positive effect of attention on consumer responses to ads. We hypothesize:

Hypothesis 8b (H8b Competing): *Privacy concerns decrease the attention towards an ad.*

Figure 3.1 shows our conceptual model that guides our study and summarizes the hypotheses established in the previous sections.

3.4 Methodology

To test our hypotheses, we conducted a scenario-based lab experiment in which we made use of eye tracking technology to record participants' eye movements. We supplemented the experiment with a questionnaire that participants needed to fill in after the experiment. During the experiment, we randomly allocated participants to being exposed to different types of social ads. We manipulated whether the ads

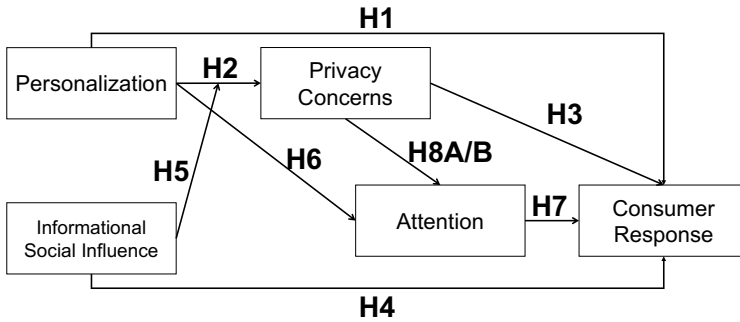


Figure 3.1: Conceptual Model

were (1) personalized or not, (2) whether they included a social endorsement, and (3) distinguished between two types of products (sexually transmitted disease (STD) treatment and gambling application).

3.4.1 Pre-Test

To identify suitable product categories for our experimental ads we ran a pre-test on Amazon’s Mechanical Turk. The aim of this pre-test was to make sure that our ads would, despite the experimental lab setting, induce privacy concerns and allow us to establish the personalization privacy paradox as the baseline model for our study.

We ran our pre-test using a scenario-based experiment in which we randomly allocated participants to being confronted with either personalized or non-personalized ads as well as one out of 5 different product categories: (1) weight loss consultation, (2) gambling application, (3) money management program, (4) STD treatment, (5) alcohol delivery service. Product categories had been pre-selected based on their privacy sensitive nature and former research (Bansal et al., 2010; Zhang et al., 2014). This set up led to a 2×5 between-subject experimental design. Appendix A3.1 presents an overview of the measures we use in the questionnaire presented to participants during the scenario-based experiment.

For the 304 participants in the pre-test, we find that our personalized treatment is being perceived as significantly more personalized ($\Delta M = 1.947$, $t = 11.032$, $p < .001$). We identify the gambling application and the STD treatment as being the most privacy-sensitive categories. Additionally, we are able to establish that privacy concerns are largely driven by our personalization treatment when comparing the personalized

with the non-personalized treatment group ($\Delta M = .375$, $t = 2.099$, $p = .036$). Also, personalized ads are being perceived as significantly more relevant by participants ($\Delta M = 2.266$, $t = 12.322$, $p < .001$).

These findings support the notion of the personalization privacy paradox. To make sure that our results are not solely driven by user characteristics of Amazon's Mechanical Turk, we replicate the study with a convenience sample from a Master's class at our university. Our findings are consistent. None of the pre-test participants was eligible to participate in the main experiment of this study.

3.4.2 Experimental Procedure

After participants were recruited from our university's participation pool, they scheduled individual appointments during which they visited the research lab. Students in the participation pool are incentivized with additional course credits to participate in the experiment. Before the start of the experiment, participants were introduced to the experimental procedure and seated.

During the experiment, all instructions were presented on the screen to create a standardized experimental experience for all participants and not introduce experimental biases caused by the experimenter. Before starting the experiment, the eye tracking device was calibrated to participants' eye movements. We made use of a 9-point calibration and validation and required participants to not deviate more than 1.0 degrees in calibration accuracy. For participants that did not meet this requirement, the calibration was repeated. In case the required accuracy was not met after several trials, participant data was discarded from the analysis, as the values recorded for their eye movements are not accurate enough to measure their attention towards the ad.

After successful calibration, participants were confronted with one of three scenarios. In case they had been randomly allocated to the personalized treatment, they were facing a scenario text that explained that they had been searching for either an (1) online poker platform to earn additional money (Gambling), or a (2) sexually transmitted disease (STD) treatment to treat a potentially contracted STD (STD treatment). In case participants had been allocated to the non-personalized treatment, they faced a scenario text explaining that they had been looking for a new laptop. All the scenario texts explained to participants that they had decided not to purchase the respective product.

Next, the instruction text pointed out that some days later the participant decided to log into her Facebook account. Participants were then required to log into their actual Facebook accounts. In a next step, they were redirected to a personalized Facebook mock up page. These Facebook mock up pages included a post on a news article and an advertisement. In case participants were in the personalized treatment group, they were confronted with an ad that matched their search scenario. In case participants had been randomly allocated to the non-personalized treatment, they were confronted to the same ads with the difference that these ads were not matching their search scenario (searching for a laptop).

We decided to operationalize the experimental manipulation for personalization this way, as participants are questioned on their perceptions of the advertising and not of the scenario. Presenting a different ad to participants in the non-personalized treatment group would have meant that effects such as attractiveness of the advertised product category or the advertising visual might have confounded our personalization treatment. Through our operationalization, we made sure that observed differences in participants' behavior and responses regarding participants' perception were driven by perceived personalization.

Next to the personalization treatment, we randomly allocated participants to either seeing ads including social endorsements or not. These social endorsements were dynamically generated for participants by using their actual Facebook information. By displaying profile pictures of actual friends, we made sure that the endorsements were being perceived as realistic. While being on the Facebook mock-up page, we allowed participants to browse freely for as long as they wanted. Both the news text and the ad were clickable. After clicking on the news article, the ad, or closing the browser window, the experiment ended.

After leaving the Facebook mock up page, we asked participants to fill in a questionnaire that would help us to explain the mechanisms underlying their behavior. Appendix A3.2 presents the questions we used in the questionnaire presented to participants after the scenario-based experiment. As a last step, participants were debriefed. We explained to them that all products advertised were not real and that we had shown them friends' endorsements that were non-existent. We present a summary of all variables used in our analyses, how they were operationalized, and the collection method in Appendix A3.3.

3.4.3 Sample

We conducted our final experiment in the research lab with students from a European university. 290 students subscribed to our experiment. Data could not be recorded for some students for varying reasons: (1) 29 students did not come or come too late to their lab appointments, (2) 4 students were not able to log in to their Facebook account as they could not reproduce their passwords, and (3) for 5 students, technical issues led to their data not being recorded.

The 252 remaining students conducted the experiment. To guarantee the data quality for our analysis, we needed to additionally discard data of several other students. (4) For 3 students, it was obvious that they did not conduct the experiment with the required attention and caution. Their records were excluded from the dataset. (5) 17 participants did not meet the attention check in which they were required to select a specific item (“Please select somewhat disagree for this statement”), (6) 27 participants did not meet the required eye tracking accuracy of a maximum deviation of 1.0 degree, and (7) 18 participants were not able to recall the product category for their ad between 5 options correctly³.

We exclude these 18 participants from our main analysis as we are assessing participants’ perceptions of the ad in the questionnaire. Especially, when trying to investigate a relationship between privacy concerns and attention, participants need to process the ad information to some extent to inform their reactions. Attention represents a prerequisite for cognitive processing (Wedel and Pieters, 2007). We assume that testing whether participants recall the advertised product category gives a good indication of whether participants processed the information in the ad. We also find that participants that do not recall the ad spend significantly less attention on the ad.

Our final sample consists of 187 participants ($M_{age} = 20.44$, 42% *male*). Participants had been randomly allocated to the different experimental treatments and product categories eventually leading to a distribution of participants to treatments as depicted in Table 3.1.

³We re-run our main analyses including the 18 participants that could not recall the product categories advertised to them. Although these participants seem to not pay full attention, it might be the case that consumers would not recall ad categories in real life as well, simply because they do not pay attention towards advertisements or try to avoid them. Results for these analyses remain mostly consistent and can be found in the appendix.

Table 3.1: Distribution of Participants in Experimental Cells

	Non-Personalized		Personalized		Total
	Gambling	STD Treatment	Gambling	STD Treatment	
Non-Endorsed	26	21	23	26	96
Endorsed	20	22	23	26	91
Total	46	43	46	52	187

3.4.4 Manipulation Checks

We make use of post-experimental manipulation checks to confirm that our personalized treatment was perceived as more personalized by participants. On a seven-point scale from 1 (“Strongly disagree”) to 7 (“Strongly agree”) (Kalyanaraman and Sundar, 2006), we find that participants perceive the ads in our personalized treatment as significantly more personalized ($\Delta M = 2.513$, $t = 13.240$, $p < .001$). Further, when asking our participants whether their ad included a social endorsement (“Yes”) or not (“No”), a significantly higher share of participants in the social endorsement treatment notice the social endorsement (“The displayed advertisement on Facebook showed my friends who like the advertiser”) ($\Delta M = .609$, $t = 10.953$, $p < .001$).

3.4.5 Measures

We measured privacy concerns with 4 items adjusted from (Sheng et al., 2008) who based their scale on (Dinev and Hart, 2004; Smith et al., 1996): “It bothers me that the advertiser is able to track information about me”, “I am concerned that the advertiser has too much information about me”, “It bothers me that the advertiser is able to access information about me”, and “I am concerned that my information could be used in ways I could not foresee”. Items were measured on a 7-point Likert scale from 1 (“Strongly disagree”) to 7 (“Strongly agree”) with an alpha reliability of 0.91. We measure participants’ attention with the help of an infrared corneal reflection eye tracker from SMI. This eye tracker “measures the distance and angle of the reflection of infrared light from the center of the pupil to determine the point of fixation of the person, after calibration” (Wedel and Pieters, 2007, p. 124). By measuring attention with the help of an eye tracking device we gain access to a behavioral measure of attention. Consistent with related research (Lee and Ahn, 2012; Pieters and Wedel, 2004) we operationalize participants’ attention as fixation duration on pre-defined areas of interest (AOIs). Fixation duration accounts for the time that

individuals are focusing on the area of the advertising and processing its information. We operationalized responses to ads as participants' actual clicks on the presented ad.

3.5 Analysis

We present the descriptive statistics and intercorrelations of our variables in Table 3.2.

Table 3.2: Descriptive Statistics and Intercorrelations

	n	Mean	SD	alpha	Pearson Correlations							
					(1)	(2)	(3)	(4)	(5)	(6)	(7)	
(1) male	187	0.42	0.50									
(2) age	187	20.44	1.33		0.13							
(3) gambling ⁺	187	0.49	0.50		-0.08	0.03						
(4) personalization ⁺	187	0.52	0.50		0.06	-0.04	-0.05					
(5) informational social influence ⁺	187	0.49	0.50		0.10	-0.05	-0.04	0.03				
(6) privacy concerns	187	4.94	1.38	0.91	0.01	0.17*	-0.08	0.26**	0.03			
(7) attention	187	7.89	2.14		0.04	-0.10	-0.06	0.18*	0.02	-0.06		
(8) click	187	0.33	0.47		0.13	-0.06	-0.10	0.58**	0.16*	0.08	0.27**	

**p<.01; *p<.05; ⁺Binary treatment variables

We conduct a step-wise analysis to test the hypotheses that we established in our theory section. We start by analyzing the impact of personalization and informational social influence, our experimental treatments, on consumers' responses to ads which we operationalize as their likelihood to click on an ad. When looking at the average click-through rate for the ads presented to participants ($M_{click} = 0.33$), we notice that this measure is unusually high in comparison to common click-through rates digital ads achieve in the field.⁴ This difference originates in our controlled lab setting in which participants can click on the news post, the ad, or leave the Facebook mock-up page, compared to a real-life scenario in which consumers face many more options to continue their browsing journey. As we are not interested in the actual height of the probability to click but rather the differences between the different treatment conditions, we are confident that we can conduct our analyses with this measure.

Our model estimates the probability of a consumer i clicking on an ad j . The probability of a click is denoted as $Pr(click_{ij} = 1)$. We model the latent probability of a click, denoted by U_{ij} , using a logit function of personalization, informational social influence and additional control variables.

$$Pr(click_{ij} = 1) = \frac{\exp(U_{ij})}{1 + \exp(U_{ij})}$$

⁴Commonly click-through rates for social advertising are below 2% (Irvine, 2018).

$$U_{ij} = \alpha_{ij} + \beta_1 \text{personalization}_{ij} + \beta_2 \text{informational social influence}_{ij} \\ + \theta X_i + \Upsilon \text{ad category}_j + \epsilon_{ij}$$

*personalization*_{ij} is a binary variable equal to 1 when an ad is personalized otherwise zero. *informational social influence*_{ij} equals 1 if the ad contains a participants' friends dynamically extracted from Facebook to display them as endorsers of the presented advertisement, otherwise zero. X_i represents a vector of participant controls including age and gender. The variable *ad category*_j controls for the advertised product category. ϵ_{ij} represents the idiosyncratic error term. Table 3.3 presents the results from our analysis.

We find that *personalization* is significantly increasing the probability of participants to click on ads ($\beta_{\text{personalization}} = 3.512$, $p < 0.001$) supporting *hypothesis 1*. At the same time, we find that *informational social influence* leads to an increased click propensity ($\beta_{\text{informational social influence}} = 0.920$, $p = 0.023$) supporting *hypothesis 4*.⁵

Table 3.3: Logit Regression - Click Probability

	(1)	(2)	(3)
DV: Click	Logit	Logit	Logit
personalization		3.512***	3.647***
		(0.566)	(0.586)
informational social influence			0.920**
			(0.405)
gender control	Yes	Yes	Yes
age control	Yes	Yes	Yes
product category control	Yes	Yes	Yes
constant	2.138	0.633	0.047
	(2.636)	(3.297)	(3.309)
Observations	187	187	187
Chi2	6.187	78.482	83.830
-2 Log Likelihood	231.405	159.109	153.762
AIC	239.405	169.109	165.762
BIC	252.329	185.265	185.148
standard errors in parentheses, *** p<0.01, ** p<0.05, * p<0.1			

⁵We replicate the analysis including the 18 participants that could not recall the product category advertised to them. Results are mostly consistent and can be found in Appendix A3.4.

We move on to analyze the relationship between our treatment variables, *personalization* and *informational social influence*, and participants' *privacy concerns*. We conduct an ordinary least squares (OLS) regression in which we assess how *privacy concerns* of participant i that is confronted with ad j are influenced by *personalization* and *informational social influence*. We operationalize *privacy concerns* as the average of our 4 respective survey items. We estimate the following model:

$$\begin{aligned} \text{privacy concerns}_{ij} = & \beta_0 + \beta_1 \text{personalization}_{ij} \\ & + \beta_2 \text{informational social influence}_{ij} \\ & + \beta_3 \text{personalization}_{ij} \times \text{informational social influence}_{ij} \\ & + \theta X_i + \Upsilon \text{ad category}_j + \epsilon_i \end{aligned}$$

In this model, β_0 represents the constant term. β_1 and β_2 give the conditional effects of *personalization* and *informational social influence*, which are both binary treatment variables. With β_3 we measure the impact of the interaction between *personalization* and *informational social influence*. Again, X_i represents a vector of demographic control variables including gender and age and *ad category* $_j$ controls for the advertised product category. Table 3.4 shows our results.

Table 3.4: OLS Regression - Privacy Concerns

	(1)	(2)	(3)	(4)
DV: Privacy Concerns	OLS	OLS	OLS	OLS
personalization		0.728*** (0.196)	0.726*** (0.196)	0.759*** (0.267)
informational social influence			0.092 (0.195)	0.127 (0.308)
personalization \times informational social influence				-0.067 (0.401)
gender control	Yes	Yes	Yes	Yes
age control	Yes	Yes	Yes	Yes
product category control	Yes	Yes	Yes	Yes
constant	1.312 (1.434)	0.696 (1.377)	0.610 (1.381)	0.579 (1.413)
Observations	187	187	187	187
R-squared	0.037	0.107	0.108	0.108

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

We find that, as hypothesized (*hypothesis 2*), *personalization* increases *privacy concerns* ($\beta_{\text{personalization}} = 0.728$, $p < 0.001$). We do not find a significant coefficient

for the effect of *informational social influence* or the interaction between *personalization* and *informational social influence*. Therefore, our analysis does not support *hypothesis 5*.⁶

We move on to estimate the impact of our treatment variables, *consumer privacy concerns* and *attention*, on consumers' responses to ads. We do not estimate the effect of *informational social influence* on *privacy concerns* in this model, as the previous model has shown that there is no significant relationship between the variables. Instead, we include *informational social influence* to control for its effect on *consumer response*. To do so, we run a serial mediation model using the SPSS PROCESS plugin (Hayes, 2013). Our results are depicted in Figure 3.2.

We find that *personalization* does significantly increase *privacy concerns* ($B = 0.73$, $SE = .19$, $p < .001$) but *privacy concerns* do not directly influence *consumer response*, therefore not supporting *hypothesis 3*. *Personalization* does significantly increase consumers' *attention* towards an ad ($B = 0.86$, $SE = .32$, $p = .009$), supporting *hypothesis 6*. We find that *attention* is positively influencing *consumer response* ($B = 1.02$, $SE = .35$, $p = .003$), supporting *hypothesis 7*. Moreover, we find that *privacy concerns* do not directly influence *attention*, therefore not supporting *hypothesis 8A* or *8B*.

To test for mediation, we make use of bias-corrected bootstrapping and generate 95% confidence intervals around the indirect effects of *privacy concerns* and *attention*. Next to that, we generate the 95% confidence interval for the indirect effect of the serial mediation through both *privacy concerns* and *attention*. In case mediation is present in our models the generated confidence intervals do not include zero (Hayes, 2013).

Our analysis, making use of 10,000 bootstraps and bias corrected confidence intervals, shows that there is no significant indirect effect for *privacy concerns* ($ab = -.18$, $SE = .15$; 95% $LLCI = -.52$, 95% $ULCI = .08$). We do find a significant indirect effect for *attention* ($ab = .87$, $SE = .50$; 95% $LLCI = .22$, 95% $ULCI = 2.2$). The indirect effect for the serial mediation was found to be significant ($ab = -.13$, $SE = .10$; 95% $LLCI = -.45$, 95% $ULCI = -.01$). *Personalization* increases *privacy concerns*, which subsequently decreases *attention*, eventually leading to a decrease in *consumer response*. This finding is in line with the notion of the personalization privacy paradox, which claims negative implications of privacy concerns on advertising performance. We show in our model that through a decrease in *attention*, *privacy concerns* do

⁶We replicate the analysis including the 18 participants that could not recall the product category advertised to them. Results are mostly consistent and can be found in Appendix A3.5.

negatively influence advertising performance. Further, we find that *informational social influence* does positively influence *consumer response* ($B = .92$, $SE = .41$, $p = 0.02$). Table 3.5 gives an overview of which of our hypotheses was supported by our analysis.




Figure 3.2: Results Serial Mediation Model

When replicating the serial mediation analysis including the 18 participants that could not recall the product category advertised to them, we find that results remain mostly consistent (see Appendix A3.6). Interestingly, we do not find support for the serial mediation effect of *personalization* through *privacy concerns* and *attention* on *consumer response*. Recalling the ad category seems to be correlated with the extent to which participants process the ad content cognitively. We see that the effect of *personalization* on *privacy concerns* seems to be stronger in our main model (excluding the participants that were not able to recall the advertised product category). The remaining effects in the chain from *personalization* through *privacy concerns* and *attention* to *consumer response* remain very close to the effects in our main model. Essentially, the serial mediation effect becomes insignificant as we include observations from participants in the sample that spend less attention on the ad. A t-test reveals that consumers that remember the advertised product category in the survey spend significantly more attention (fixation duration) on the ad ($\Delta M = 2254.87$, $t = 2.05$, $p = .04$). Consumers that do not recall the advertised product category are likely to process the ad information less thoroughly, decreasing the impact of personalization on their privacy concerns.

Table 3.5: Overview Hypotheses

Hypothesis	Hypothesized Effect	Support*
Hypothesis 1	personalization → consumer response (positive)	S
Hypothesis 2	personalization → privacy concerns (positive)	S
Hypothesis 3	privacy concerns → consumer response (negative)	nS
Hypothesis 4	informational social influence → consumer response (positive)	S
Hypothesis 5	informational social influence <i>mitigating</i> (personalization → privacy concerns)	nS
Hypothesis 6	personalization → attention (positive)	S
Hypothesis 7	attention → consumer response (positive)	S
Hypothesis 8A	privacy concerns → attention (positive)	nS
Hypothesis 8B	privacy concerns → attention (negative)	nS

*nS = not supported, S = supported

3.6 Discussion

Our multi-method approach, combining a scenario-based experiment with eye tracking and survey methodology, allows us to gain deeper insights into consumers' evaluation of personalized advertising. The focus of our research is twofold. First, we want to investigate whether informational social influence can mitigate the effect of personalization on consumer privacy concerns. Second, we want to shed light on what is the mediating role of attention on consumers' responses to ads when confronted with the personalization privacy paradox. More specifically, we focus on the question whether privacy concerns do actually increase consumer attention that subsequently negatively influences consumers' responses to ads, or if consumer privacy concerns do decrease attention towards an ad, supporting a generally positive notion of the influence of attention on consumers' responses to ads.

In line with research in the area of ad personalization (Arora et al., 2008; Bleier and Eisenbeiss, 2015b; Lambrecht and Tucker, 2013) and supporting our *hypothesis 1*, we find that personalization leads to more positive consumer responses to ads. Ads that match consumers' preferences are more likely to induce a favorable reaction from consumers.

Supporting *hypothesis 2* and the personalization privacy paradox, we find that personalization increases consumer privacy concerns. Although, personalized ads are perceived as more relevant by participants in our experiment, participants also experience privacy concerns as their individual browsing information is being used to personalize ads. These privacy concerns are likely to stem from the perceived lack of control over personal information that the advertiser is using without consumers' consent.

We do find that the effect of privacy concerns on consumers' responses to ads is negative but not significant, therefore not supporting *hypothesis 3* and partially contradicting previous findings (Aguirre et al., 2015; Bleier and Eisenbeiss, 2015a). The reason for this might be, that privacy concerns that are induced do not impact participants' actions in our lab setting very strongly as they have the feeling of conducting the experiment in a safe environment. Participants are not using a personal device but a device provided by us.

Furthermore, our analysis shows that, in line with previous studies investigating informational social influence in the ad context (Bakshy et al., 2012; Tucker, 2016), informational social influence does positively influence consumers' responses to ads, supporting *hypothesis 4*.

We do not find informational social influence to be a moderator in the relationship between personalization and consumer privacy, therefore not supporting *hypothesis 5*. As the inclusion of social endorsements in ads represents a use of personal information regarding social connections in itself, social endorsements might potentially trigger privacy concerns. The privacy concerns originating in social endorsements might counterbalance the moderating effect of informational social influence. Triggering informational social influence with social endorsements does not seem to be a viable option for advertisers to mitigate privacy concerns caused by the personalization of ads. Still, informational social influence seems to remain an effective way to increase ad performance through more positive consumer responses to ads.

Our full model gives insights into the role of attention as cognitive enabler in processing advertising content. We find that personalization is, next to leading to more positive consumer responses to ads, increasing consumers' attention towards ads, supporting *hypothesis 6*. This finding is in line with the notion that personalization is a top-down factor steering consumers' attention towards more relevant content (Bang and Wojdyski, 2016). When consumers do cognitively process advertising content and recognize its relevance, their attention is biased towards the ad which is perceived as relevant.

We also find a significant direct relationship between attention and consumer response, supporting *hypothesis 7*. This represents evidence for the positive role of attention. Attention is necessary to process information in an ad (Lee and Ahn, 2012; Tam and Ho, 2006), acting as enabler for the behavior a consumer decides to take subsequently (Köster et al., 2015; Wedel and Pieters, 2007).

Using conditional process analysis, we find that the relationship between personalization and consumer response is partially mediated by attention. This means that part

of the positive effect of personalization on consumers' responses to ads is operating indirectly via an increase in consumers' attention towards ads.

Additionally, we find that privacy concerns seem to not have a significant direct effect on consumers' attention towards an ad, therefore supporting neither *hypothesis 8A* nor *8B*. While the direction of the effect of privacy concerns on attention is negative, standing in favor of a generally positive effect of attention on consumer response, we do not find a significant relationship.

Despite the non-significant direct relationship, we do find evidence for a serial mediation effect from personalization through privacy concerns and attention to consumer response. Having non-significant direct effects within a mediation path does not mean that the indirect effects in conditional process models are non-significant as they should be tested separately with a single test for significance (see Hayes (2013)). We find that personalizing an advertisement increases consumers' privacy concerns, which subsequently leads to a decrease in attention towards the ad, eventually leading to a negative influence on consumers' responses to ads. This finding is in line with a positive role of attention in consumers' assessment process of advertisements. Higher levels of attention are correlated with more positive consumer responses to ads. While personalization can be seen as a top down factor that positively influences the attention level of consumers by directing attention to more relevant content, privacy concerns represent a top down factor that is directing attention away from negatively perceived ad content. Although we do not find that privacy concerns do negatively influence consumer response directly, they still play a role as mediator in our model hampering consumer responses to ads.

Interestingly, this serial mediation effect becomes insignificant when re-running the serial mediation model and including observations of participants that were not able to recall the product category advertised to them (see Appendix A3.6). We find that consumers that recall the product category spent significantly more attention on the ad. As discussed earlier, attention is a prerequisite for cognitive processing of information (Wedel and Pieters, 2007). This finding points towards privacy concerns not only influencing attention, but attention also influencing to what extent consumers process information potentially resulting in privacy concerns. For consumers that do not cognitively process ad information, we are not able to find evidence for the personalization privacy paradox. This is the case as these consumers are not able to evaluate the personalized ad as being intrusive, invading their privacy, when not dedicating a minimum amount of attention towards the ad.

3.6.1 Practical Implications

Our research gives practitioners valuable insights into both the role of informational social influence and consumers' attentional mechanisms in the context of personalized advertising.

We find that the inclusion of social endorsements, as typical in social advertising, is justified through a direct positive impact on consumers' responses to ads. Consumers are significantly more likely to click socially endorsed ads than unendorsed ads. Nevertheless, we also show that triggering informational social influence with social endorsements cannot be considered a fruitful strategy to mitigate consumer privacy concerns. Consumers' acceptance of their personal information being used to personalize ads seems to not increase when endorsing personalized ads.

Next to that, we disentangle the attention mechanisms within the personalization privacy paradox. While the fact that privacy concerns have a negative influence on advertising performance is vastly established (Aguirre et al., 2015; Bleier and Eisenbeiss, 2015a; Tucker, 2014) we assess how consumers' attention processes are related to the conflict between personalization and privacy concerns. Consistent with the notion of the positive impact of attention on consumers' responses to ads, we find that consumers reduce the level of attention towards an ad if they experience privacy concerns. Practically, this means that more intrusive ads might initially catch consumers' attention but, as they decrease the overall level of attention that consumers dedicate to an ad, consumers' responses to these ads are more negative. Therefore, we suggest to advertisers to abstain from using highly intrusive ads to compete for consumers' attention, as these ads decrease consumers' attention, leading to a decrease in ad performance.

3.6.2 Limitations

We conduct our study within the research lab of a university using university students as subjects. This means that the external validity of our study might be limited. This also becomes apparent in the inflated click-through probability in our scenario-based experiment compared to commonly observed click-through probabilities for digital ads in the field. In the field, social advertising click-through rates are below 2% (Irvine, 2018). At the same time, we are forced to choose this environment for our study to support our multi-method approach, especially the assessment of privacy concerns with the help of a questionnaire and the collection of attention measures employing eye tracking.

Furthermore, we confront participants in our experiment with very privacy sensitive product advertisements for a gambling application or a STD treatment. We make this choice as we want to induce privacy concerns with participants in a lab environment that is likely to be perceived as safe. Potentially, participants might find it difficult to relate to scenarios describing their interest in a gambling application or medical support for a STD treatment. A t-test regarding the perceived realism of the situation (*'I found the described situation realistic.'*, 7-point scale) shows that there is no significant difference in participants' perception between the personalized and non-personalized treatment group ($\Delta M = .051$, $t = .314$, $p = .753$). Less privacy sensitive products would likely not induce privacy concerns with participants and therefore not allow us to assess the personalization privacy paradox.

We choose to manipulate the scenario texts to distinguish between the personalization and the non-personalization treatment groups instead of manipulation the ad copies. We make this choice as we are assessing participants' reactions to the ads and not to the scenario texts. Therefore, we want to avoid introducing confounding effects with different advertising visuals and messages.

The operationalization of our attention variable does not only capture the attention driven by our treatment variables and consumer privacy concerns. The main reason for this is that, as previously mentioned, attention represents a prerequisite for consumers to process information (Wedel and Pieters, 2007). Therefore, consumers first need to dedicate their attention to the personalized ad before they can recognize that the ad is personalized for them and that they might be concerned about their privacy because of how the firm used their personal information. Our additional analysis in Appendix A3.6 supports this notion. Experiment participants that do not recall the advertised product category spent less attention on the ad, and most likely did process the ad information less thoroughly. Furthermore, the effect of privacy concerns on attention is smaller for these consumers, presumably because they did not process the ad information thoroughly.

While it is likely that privacy concerns negatively influence consumers' responses to ads through a decrease in attention, one could argue that a certain level of attention is necessary to actually cognitively process an ad, leading to privacy concerns in the first place. Attentional processes enable cognitive processes that are necessary for individuals to construct their opinions. To counteract this issue, we operationalize attention as the overall attention spent on an ad up to the first click on a website element. This way, we make sure that we assess attention related to the overall cognitive processing of an ad. We follow the assumption that consumers need to spend

a certain amount of attention on advertising to cognitively process the information and then decide whether they want to investigate the ad further, i.e. dedicate more attention towards the ad. The dynamic relationship between attention and cognitive processes presents fruitful grounds for future research. By observing which ad elements consumers focus on, researchers can get additional information on which information stimuli influence consumers' decision to continue paying attention to an ad or allocate their attention somewhere else.

Appendices

A3.1 Survey Measures Used in Pre-Test

Table A3.1: Survey Measures Used in Pre-Test

Construct	Items	Scale	Source
Perceived Personalization	(1) The content and information featured in the advertisement targeted me as a unique individual. (2) This advertisement was 'personalized' according to my interests.	7-point	(Kalyanaraman and Sundar, 2006)
Privacy Concerns	(1) It bothers me that the advertiser is able to track information about me. (2) I am concerned that the advertiser has too much information about me. (3) It bothers me that the advertiser is able to access information about me. (4) I am concerned that my information could be used in ways I could not foresee.	7-point	(Sheng et al., 2008) based on (Dinev and Hart, 2004; Smith et al., 1996)
Personal Relevance	(1) This advertisement was useless to me. (2) This advertisement was irrelevant to me. (3) This advertisement was not personally important to me.	7-point	(Campbell and Wright, 2008)
Intention to Click	(1) I would like to click-through the advertisement to acquire further information.	7-point	(Yoo, 2007)

Constructs listed in same chronological order as asked in pre-test questionnaire.

A3.2 Survey Measures Used in Scenario-Based Experiment

Table A3.2: Survey Measures Used in Scenario-Based Experiment

Construct	Items	Scale	Source
Perceived Personalization (<i>Manipulation Check</i>)	(1) The content and information featured in the advertisement targeted me as a unique individual. (2) This advertisement was 'personalized' according to my interests.	7-point	(Kalyanaram and Sundar, 2006)
Social Endorsement (<i>Manipulation Check</i>)	(1) The displayed advertisement on Facebook showed my friends who like the advertiser.	Yes/No	-
Personal Relevance	(1) This advertisement was useless to me. (2) This advertisement was irrelevant to me. (3) This advertisement was not personally important to me.	7-point	(Campbell and Wright, 2008)
Privacy Concerns	(1) It bothers me that the advertiser is able to track information about me. (2) I am concerned that the advertiser has too much information about me. (3) It bothers me that the advertiser is able to access information about me. (4) I am concerned that my information could be used in ways I could not foresee.	7-point	(Sheng et al., 2008) based on (Dinev and Hart, 2004; Smith et al., 1996)
Attention Check 1	(1) Please select <i>somewhat disagree</i> for this statement.	7-point	-
Perceived Realism	(1) I found the described situation realistic.	7-point	-
Attention Check 2	(1) The advertisement I saw introduced a ...	5 product categories*	-
Gender	(1) What is your gender?	male/female	-
Age	(1) What is your age?	18-99	-

Constructs listed in same chronological order as asked in the questionnaire in the scenario-based experiment.

*Weight-loss product/Gambling product/Financial product/STD treatment/Alcoholic beverage

A3.3 Overview of Variables Used in Analysis and Data Collection Method

Table A3.3: Variable Overview

Variable	Operationalization	Collection Method
Personalization	Experimental treatment: (1) Personalized if scenario text matches advertised product category, (2) Non-personalized if scenario text does not match with advertised product category	Experiment
Informational Social Influence	Experimental treatment: (1) Inclusion of friends' profile pictures and number of friends that like the advertised product to trigger informational social influence, (2) Friends' profile pictures not included	Experiment
Privacy Concerns	Measured with survey items	Questionnaire
Attention	Fixation duration on advertisement in milliseconds	Eye tracker
Consumer Response	Binary variable indicating whether a participant clicked on the ad or not	<i>BeGaze</i> eye tracking software
Product Category	Experimental treatment: (1) Gambling application, (2) STD treatment	Experiment
Gender	Survey question	Questionnaire
Age	Survey question	Questionnaire

A3.4 Logit Regression for Click Probability

We re-run our analysis from Table 3.3 including also experiment participants that did not recall the product category that was advertised to them. Table A3.4 presents our results which remain consistent with the main analysis in which these participants are excluded.

Table A3.4: Logit Regression - Click Probability

	(1)	(2)	(3)
DV: Click	Logit	Logit	Logit
personalization		3.653***	3.750***
		(0.565)	(0.580)
informational social influence			0.848**
			(0.394)
gender control	Yes	Yes	Yes
age control	Yes	Yes	Yes
product category control	Yes	Yes	Yes
constant	2.224	1.630	1.323
	(2.569)	(3.224)	(3.220)
Observations	205	205	205
Chi2	6.128	87.887	92.663
-2 Log Likelihood	248.419	166.660	161.884
AIC	256.419	176.660	173.884
BIC	269.711	193.275	193.822
standard errors in parentheses, *** p<0.01, ** p<0.05, * p<0.1			

A3.5 OLS Regression for Privacy Concerns

We re-run our analysis from Table 3.4 including also experiment participants that did not recall the product category that was advertised to them. Table A3.5 presents our results which remain consistent with the main analysis in which these participants are excluded.

Table A3.5: OLS Regression - Privacy Concerns

DV: Privacy Concerns	(1) OLS	(2) OLS	(3) OLS	(4) OLS
personalization		0.622*** (0.185)	0.620*** (0.186)	0.675*** (0.258)
informational social influence			0.039 (0.187)	0.095 (0.288)
personalization × informational social influence				-0.111 (0.382)
gender control	Yes	Yes	Yes	Yes
age control	Yes	Yes	Yes	Yes
product category control	Yes	Yes	Yes	Yes
constant	1.960 (1.350)	1.570 (1.305)	1.541 (1.311)	1.489 (1.345)
Observations	205	205	205	205
R-squared	0.027	0.079	0.079	0.079

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

A3.6 Serial Mediation Model

We re-run our serial mediation analysis including observations of participants that were not able to recall the advertised product category correctly. As our sample size includes the observations for these 18 participants, we now analyze the model with 205 observations. Figure A3.1 presents our results.

We find, consistent with our main model, that *personalization* does significantly increase *privacy concerns* ($B = 0.62$, $SE = .19$, $p = .001$) but *privacy concerns* do not directly influence participants' propensity to click the ad, *consumer response*, therefore not supporting *hypothesis 3*. Similarly, *personalization* does significantly increase consumers' *attention* towards an ad ($B = 0.85$, $SE = .33$, $p = .01$), sup-

porting *hypothesis 6*. Also *attention* is positively influencing *consumer response* ($B = 1.02$, $SE = .34$, $p = .003$), supporting *hypothesis 7*. *Privacy concerns* do not directly influence *attention*, therefore not supporting *hypothesis 8A* or *B*.

We use bias-corrected bootstrapping and generate 95% confidence intervals around the indirect effects of *privacy concerns* and *attention* to test for mediation. Further, we generate the 95% confidence interval for the indirect effect of the serial mediation through both *privacy concerns* and *attention*. In case mediation is present in our models the generated confidence intervals do not include zero (Hayes, 2013).

The bias corrected confidence intervals generated from 10,000 bootstraps reveal that there is no significant indirect effect for *privacy concerns* ($ab = -.13$, $SE = .12$; 95% $LLCI = -.41$, 95% $ULCI = .08$). Consistent with the main analysis, we find a significant indirect effect for *attention* ($ab = .87$, $SE = .49$; 95% $LLCI = .23$, 95% $ULCI = 2.1$).

Interestingly, the indirect effect for the serial mediation was found to be not significant when including the observations from participants that did not recall the advertised ad category in the analysis ($ab = -.10$, $SE = .09$; 95% $LLCI = -.36$, 95% $ULCI = .01$).




Figure A3.1: Results Serial Mediation Model - Full Sample

Recalling the ad category seems to be correlated with the extent to which consumers process the ad content cognitively. We find that the effect of *personalization* on *privacy concerns* is stronger in our main model. The remaining effects in the chain from *personalization* through *privacy concerns* and *attention* to *consumer response* remain in magnitude very close to the effects in our main model. Essentially, the serial mediation effect becomes insignificant as we include observations from participants in the sample

that spend less attention on the ad. A t-test reveals that consumers that remember the advertised product category in the survey spent significantly more attention (fixation duration) on the ad ($mean_{\text{attention}} = 5885.50$) ($\Delta M = 2254.87$, $t = 2.05$, $p = .04$).

The comparison of these analyses leads back to our discussion regarding the dynamic relationship between attention, personalization, and privacy concerns. Consumers need to spend a sufficient amount of attention on the ad to retrieve and process the information from the ad cognitively. This way, consumers will recognize that an ad is personalized for them and might become concerned about their privacy as the advertiser is making use of their personal information to personalize the presented ad. When consumers do not process information in the ad, it is unlikely that privacy concerns will arise. We elaborate further on this dynamic relationship in the discussion section.

Pay For What You Get - Incentive Misalignments in Programmatic Advertising: Evidence From A Randomized Field Experiment¹

4.1 Introduction

Firms are serving millions of digital ad impressions to consumers on a monthly basis. This makes it prohibitively expensive for firms to manually determine their willingness to pay for serving a specific ad to an individual consumer. Ad platforms offer a solution to this problem by providing ad allocation tools that firms can use to automate the

¹Earlier versions of this study appeared in the following conference proceedings or were presented at the below mentioned conferences and workshops:

- Frick, T.W., Telang, R. & Belo, R. (2018). Pay For What You Get - Incentive Misalignments in Programmatic Advertising: Evidence From A Randomized Field Experiment. In *The Economics of Information and Communication Technologies, ZEW Conference, Centre for European Economics Research*, Mannheim, Germany.
- Frick, T.W., Telang, R. & Belo, R. (2018). Pay For What You Get - Incentive Misalignments in Programmatic Advertising: Evidence From A Randomized Field Experiment. In *Statistical Challenges in eCommerce Research Symposium 2018 (SCECR)*, Rotterdam, The Netherlands.
- Frick, T.W., Telang, R. & Belo, R. (2017). Pay For What You Get - Incentive Misalignments in Programmatic Advertising: Evidence From A Randomized Field Experiment. In *Workshop on Information Systems and Economics 2017 (WISE)*, Seoul, South Korea.
- Frick, T.W. & Telang, R. (2017). Pay For What You Get - Incentive Misalignments in Programmatic Advertising: Evidence From A Randomized Field Experiment. In *Conference on Digital Experimentation 2017 (CODE)*, Boston, US.
- Frick, T.W. & Telang, R. (2017). Pay For What You Get - Incentive Misalignments in Programmatic Advertising: Evidence From A Randomized Field Experiment. In *NBER Summer Institute 2017, Economics of Information Technology and Digitization Workshop*, Boston, US.

purchasing of ad impressions called programmatic advertising. Simply put, these tools use machine learning algorithms that determine the willingness to pay for a specific ad impression on behalf of the firm. More precisely, as most opportunities to display an ad are auctioned off, the ad platform makes the decision how much to bid for an ad impression on behalf of the firm. This decision is based on large amounts of data on individual consumers' characteristics and online behavior. In programmatic advertising, firms select an algorithmic optimization rule in line with their campaign objective (e.g. generate traffic, opt-ins for newsletters, purchases, etc.). The optimization algorithms implemented by the ad platform do then adjust the bid for each impression with the goal to maximize the chosen campaign objective. Recent industry reports show that in 2016, 50% of US firms spent more than half of their marketing budget on programmatic advertising to optimize their bidding for ads (AdRoll, 2017).

In general, firms want to increase profit with advertising. In *performance marketing* firms advertise with the objective to increase purchases conducted by consumers and quantify the performance of advertising. Although usually, advertising performance is operationalized as purchases, it can also be measured with metrics that have a less immediate impact on the bottom line, e.g. newsletter opt-ins. With the objective to increase profit in mind, it becomes clear what the actual value of advertising for firms is: The difference in profit when serving ads to consumers to the profit when not serving ads to consumers. Therefore, when firms outsource the task to bid for ad impressions to ad platforms they would like ad platforms to target consumers with ad impressions whose profit contribution increases because of these ad impressions.

Historically, firms have strived to incentivize ad platforms for actual success contribution. Therefore, advertisers in the digital ad industry have moved from simply paying for ad impressions (CPM), to paying only for when consumers click on ads to reach their websites (CPC), to eventually paying ad platforms only when consumers addressed with ads conduct a pre-defined target action (CPA), usually a purchase². Even though firms have now commonly implemented CPA-based contracts, in which they set a *CPA* value as their willingness to pay for a purchase reported by the ad platform, they start to recognize that they should reward ad platforms only for the increase in purchases caused by ads and not for all purchases from consumers that

² *CPM* stands for cost per mille (a thousand impressions), *CPC* for cost per click, and *CPA* for cost-per-acquisition and cost-per-action which are used interchangeably in the ad industry. Besides the wish of firms to incentivize ad platforms only for actual success contribution this evolution of ad contract structures was also driven by technological developments, most notably cookie technologies, that allow ad platforms to track clicks and target actions on firms' websites.

have been addressed with ads. This is the case as some consumers that are addressed with ads would have purchased independent of seeing advertising. While conceptually the value of advertising is clear to firms, it is often very difficult to assess the profit contribution of advertising. The identification of the increase in profit caused by advertising hinges on an experimental set-up where treatment and control group are well identified. Ad platforms do not run such experiments which would allow them to report actual ad performance, but instead report measures such as the absolute purchase probability of or absolute revenue generated by consumers that have been addressed with ads.

To deal with the fact that CPA-based contract structures reward ad platforms based on the absolute purchase probability instead of the increase in purchase probability caused by ads, firms deflate their willingness to pay for a purchase reported by an ad platform. Deflating the willingness to pay for a reported purchase, i.e. deflating the *CPA* value, has the aim to more closely match advertising's profit contribution. Nevertheless, to optimally deflate their willingness to pay for a purchase, firms need to possess information on the increase in purchase probability caused by ads. In the best case, such information is available from experiments in the past.

Additionally, despite such a deflation procedure, ad platforms are still being rewarded based on the number of purchases conducted by consumers that were addressed with ads. This means that ad platforms have an incentive to target consumers that are likely to purchase independent of their reaction to ads. The impact of this incentive structure for the firm has so far not been empirically investigated. In case, the ad platform acts in accordance with incentives defined in the contract with firms and targets high purchase probability consumers, the firm is facing a potential incentive misalignment. Whether such targeting behavior is harmful for the firm depends on whether consumers with high absolute purchase probabilities are more or less receptive to ads. The aim of this study is to empirically investigate the potential incentive misalignments in CPA-based contract structures in programmatic advertising by assessing the impact of these contract structures on firms' return on ad investments.

To do so, we run a large scale randomized field experiment in collaboration with a European e-retailer. During the experiment, we randomly allocate consumers to seeing either retargeting ads of the e-retailer or orthogonal charity ads, also called public service announcements (PSA). This way, we are able to identify the causal impact of ads on consumers' purchase probability. To investigate whether ad platforms optimize the bidding for ad impressions in the interest of the firm we let the ad platform in our

experiment optimize the bidding for ads based on a CPA-based contract. This way, we gain access to a data set that is unique with respect to several characteristics:

- (i) We have detailed information on individual consumers' browsing behavior and the ad platform's bids for individual impressions, allowing us to infer which consumers the ad platform is targeting to which degree.
- (ii) Through our experiment, we have a symmetric control and treatment group across different consumer segments. We are able to employ a CPA bidding optimization without introducing a selection bias into the control group, which has been pointed out as a flaw in experimental designs for ad effectiveness measurement in previous research (Johnson et al., 2017a). Therefore, with the help of our experiment, we can precisely identify the change in purchase probability across heterogeneous consumer segments.
- (iii) Finally, our data set includes detailed information on consumers' responses to ad impressions such as website visits, ad clicks, and purchases.

We empirically investigate, the nature and magnitude of this potential incentive misalignment between firm and ad platform. To do this, we measure the true increase in probability of purchase for consumers who are more or less extensively targeted by the ad platform. First, we find that ad platforms act in accordance with the incentives specified in CPA-based contracts and systematically target users who are more likely to purchase independent of the effect of ads. We find evidence that ad platforms are using data on consumers' browsing history and characteristics to identify high purchase probability consumers. While advertising does generally increase consumers' purchase probability, we do not find any evidence that consumers that are targeted by the ad platform are more receptive to ads. There is no significant correlation between a consumers' inherent absolute purchase probability and the increase in purchase probability caused by ads. This renders the ad platform's targeting strategy sub-optimal for firms. Furthermore, the ad platform's optimization algorithm is also not taking consumers' heterogeneous profit contributions into account when optimizing the firm's bidding strategy, confirming that the employed bidding strategy is sub-optimal.

This work contributes to the literature on economics of advertising and expands the understanding of how to assess advertising effectiveness as well as the implications of incentives specified in contracts between firms and ad platforms. While previous work has pointed out that incentives in contracts between firms and ad platforms are not specified in the true interest of firms (Dalessandro et al., 2012; Johnson and Lewis,

2015), we assess empirically whether this potential incentive misalignment is an actual incentive misalignment. The actual presence of such a misalignment depends on the behavior of the ad platform as well as the relationship between consumers' inherent absolute purchase probabilities and the increase in purchase probabilities caused by ads. Our unique data set allows us to identify the causal impact of advertising on consumers' purchase probabilities and whether this causal impact is heterogeneous depending on the height of bids that a consumer receives from the ad platform. The setting of our experiment allows us to make use of the variation in bids without the introduction of selection bias into our sample as the ad platform is optimizing the campaigns for treatment and control group simultaneously.

We find that, in accordance with the incentives specified in CPA-based contracts, ad platforms target consumers with high absolute purchase probabilities. We do not find that these consumers react more positively to ad impressions. Using this empirical strategy, we are able to show that what we theoretically present as potential incentive misalignment turns out to be an actual misalignment. We do not find evidence of a positive correlation between consumers' purchase probabilities (as well as the bids they receive) and the increase in purchase probabilities caused by ads.

This work has significant practical implications for the digital advertising industry. While programmatic advertising and CPA-based contracts with ad platforms have become very popular among firms, we show that these contracts are not beneficial for firms. We show empirically that ad platforms follow the incentives specified in these contracts and target high purchase probability consumers to increase their profits. This behavior is harmful to firms as they pay more for ad impressions that are not more effective. Therefore, we present evidence that shows that firms need to change their ad allocation mechanisms with the aim to target more receptive consumers with ads.

The remainder of this paper is organized as follows. We first provide an overview of related literature in the area of behaviorally targeted advertising, ad effectiveness, and ad auctions. Next, we lay out our theoretical explanation for the potential incentive misalignment between firms and ad platforms. We then present the experimental design of our randomized field experiment followed by our econometric analysis. To conclude we summarize our theoretical contributions and practical implications of this study.

4.2 Related Literature

4.2.1 Behaviorally Targeted Advertising

Ad platforms' capability to use extensive amounts of individual-level data on consumers to target consumers has inspired research in the areas of both marketing and information systems. In behavioral targeting, also called retargeting, ad platforms utilize consumers' browsing behavior on firms' websites to readdress consumers with relevant products on external sites. Lambrecht and Tucker (2013) find that consumers are more responsive to generic retargeting ads instead of dynamic retargeting ads which display specific products that consumers have visited before. Only when consumers develop construed preferences, indicated by their visits to review websites, dynamic retargeting ads outperform less personalized generic ads. Other research finds that higher degrees of ad personalization with respect to product category and brand advertised in an ad, do outperform ads with lower degrees of ad personalization (Bleier and Eisenbeiss, 2015a). Nevertheless, ads with a high degree of ad personalization decrease in performance faster over time when losing their recency, the distance between a consumer's website visit and the display of an ad based on the behavior during that visit. Frick and Li (2016) find that retargeting ads that advertise a specific product do generally outperform ads that advertise less specific product categories. Furthermore, they suggest that socially targeted consumers might actually react negatively to personalized ads due to the conflict of personalizing ads for a single consumer with displaying names of friends in ads.

Commonly, studies in the area of retargeting compare the effectiveness of different types of ads. Effectiveness is in these cases measured as 'how many consumers addressed with a certain type of ad did conduct a certain action' (purchased, clicked on the ad, subscribed to a newsletter, etc.). Therefore, these studies do not necessarily measure the effect of ads on consumers' probability to conduct a certain action, the increase in the probability, but rather the absolute probability that a consumer will conduct a certain action given she was addressed with an ad. Such operationalization of effectiveness might be problematic as a higher absolute probability to conduct an action might not correlate with a higher impact of ads. In retargeting, consumers have an inherently high purchase probability, independent of the confrontation with ads (Dalessandro et al., 2012). This is the case as only consumers that have previously indicated their interest in a firm's product offerings, by visiting the firm's product pages, can be readdressed with this type of advertising. The generally high baseline

purchase probability of consumers, independent of being addressed with an ad, makes the context of retargeting especially interesting to investigate a potential incentive misalignment between firm and ad platform. Retargeting allows ad platforms to more easily identify high purchase probability consumers by making use of information on their browsing behavior. In this context, success measurement is an important topic as ad platforms can potentially more easily target high purchase probability consumers independent of the effect of ads and in consequence over-report success of retargeting campaigns.

4.2.2 Ad Effectiveness

While most of the research dealing with the design of targeted digital advertising focuses on the question of how certain elements in ads should be specified to lead to more positive consumer responses to these ads, the question of what the actual impact of these ads on a firm's performance is, is often neglected.

Dalessandro et al. (2012) discuss the issue of both ad effectiveness measurement and attribution and suggest that a good attribution system needs to align the incentives of both advertiser and advertising outlet which is contracted to serve ads. Last-click attribution, that is assigning credit for a conversion to the advertising channel that addressed a consumer last in the purchase process, is an example for these misaligned incentives. Here, advertising outlets have an incentive to confront consumers with ads as late as possible in their purchase process without taking an increase in consumers' purchase probabilities into account. Under a last-touch attribution system, success is attributed to the last ad channel that touches the consumer in the purchase process. Depending on the contract with the advertising firm, these channels can charge a fee from the firm for conversions of consumers addressed with ads.

The causal interpretation of an ad's effect on a consumer's purchase probability hinges on the possibility to identify this effect. The classic solution to identify the causal impact of ads on consumers' purchase probability are so-called hold-out experiments in which a certain percentage of consumers is randomly assigned to ads that are assumed to be orthogonal to the firm's ads, while the rest of the consumers are being exposed to the firm's regular ads. This approach is taken, since without serving ads to consumers in the control group they can technically not be identified as members of the control group. The difference between the consumers' purchase probabilities in treatment and control group can then be interpreted as the causal effect of ads on consumers' purchase probabilities.

This approach has been criticized for being expensive, making firms spend money on ads for consumers in the control group that are not contributing to the firms success, as well as for neglecting the defensive effect of ads (Johnson et al., 2017a). The defensive effect of ads describes that serving consumers unrelated ads does not represent the true counterfactual. In the real world, consumers might be exposed to ads of competitors that persuade them to purchase from the competitor instead of the focal firm. Therefore, ads can have a defensive mechanisms circumventing consumers to purchase from competitors that is not captured with these types of hold-out experiments. Johnson et al. (2017a) propose a methodology to overcome this limitation, called ghost ads, in which ad platforms are required to run hidden auctions for consumers in the control group to determine which consumer would have seen an ad if treated and should therefore be included into the control group. Although this methodology is clearly superior to a hold-out approach, it hinges on the fact that it needs to be implemented on the side of the ad platform. It may not be in the interest of ad platforms to implement such a system and carry their costs when they are the side currently benefiting from the incentive structure in programmatic advertising.

Related work has suggested to reward ad platforms based on the actually contributed uplift in purchase probabilities and therefore incentivize them to take this metric into account when bidding for ad spaces on behalf of a firm (Xu et al., 2016). Other authors showcase how to increase the precision in estimating the impact of online display advertising on offline sales of a retailer by including covariates, and even more drastically, by reducing noise in the control group by excluding observations for consumers that would have not been treated (Johnson et al., 2017b).

Other work focuses on the pricing models that are currently implemented in contracts between firms and ad platforms in which ad platforms are rewarded based on the absolute outcome, instead of the incremental increase in the outcome variable (Johnson and Lewis, 2015). We add to the discussion by empirically investigating the implications of contracts with incentives that are not in line with firms' true interest. We acknowledge the fact that the most trivial way to solve the potential incentive misalignment is to specify the contract between firm and ad platform differently. Nevertheless, we focus on the empirical assessment of whether the currently specified incentives do actually lead to an incentive misalignment between firm and ad platform.

4.2.3 Ad Auctions

Over the last years, the way display advertising slots are sold to advertisers has changed. Previously, advertisers were purchasing display ad slots in bulk, called 'direct' buying or guaranteed contracts³, and paid a fixed amount per ad shown to consumers. Nowadays, more and more ad slots are auctioned off in so-called real-time bidding auctions. Direct buying of ad spaces limits advertisers flexibility with respect to adjustments to traffic patterns or market conditions as agreements are made upfront (Balseiro et al., 2014). Using auctions to allocate ad spaces is considered more efficient, also as ad spaces are purchased by the advertiser that values the ad space the highest⁴ (Arnosti et al., 2016). Using auctions to sell digital ad space has been established in search engine advertising where advertisers bid for pre-defined keywords that users enter into the search field. Ad platforms use second-price auctions to sell off their ad space (Edelman et al., 2007). In the most simple form of these type of auctions, advertisers bid for an ad slot and are then ranked according to their willingness to pay. The advertiser with the highest bid then buys the ad slot paying the price matching the bid of the second highest bidder. It has been shown that the dominant strategy for such auctions, also called Vickrey-Clarke-Groves auctions, is to bid in accordance with one's true valuation (Clarke, 1971; Groves, 1973; Vickrey, 1961).

As both the volume of impressions has grown as well as the decision of whom to show an ad has increased in complexity, firms make use of bidding optimization tools provided by ad platforms to guide their ad allocation process. These optimization tools take firms' campaign objectives, overall willingness to pay for advertising, as well as detailed information of the respective consumer that can be addressed with an ad into account. With this information, ad platforms decide, on behalf of the firm, how much to bid for an ad impression.

4.3 Theory

Economic Value of an Ad Impression. When outsourcing the bidding for ad impressions to ad platforms, firms have the goal to pay no more for an ad impression than its actual value. The economic value of an ad impression i for a consumer j can

³An advertiser would individually negotiate a price for e.g. 100,000 impressions served over a pre-defined time frame.

⁴Hinging on advertisers knowing their true valuation of an ad impression.

be expressed by:

$$v_{i,j} = E(\pi_j|I = 1) - E(\pi_j|I = 0)$$

where $v_{i,j}$ represents the value of an ad impression i for consumer j , which we write as the difference between the expected value of profit when an ad is served to the expected value of profit when no ad is served. π_j is the outcome variable of interest and takes monetary values when a purchase occurs ($\pi_j > 0$) or equals zero in case no purchase occurs. I equals 1 if an impression is served to a consumer and zero otherwise. This equation can also be expressed as:

$$v_{i,j} = P(\pi_j > 0|I = 1) \times \pi_j - P(\pi_j > 0|I = 0) \times \pi_j$$

which equals:

$$v_{i,j} = \Delta P(\pi_j > 0)_i \times \pi_j$$

where $\Delta P(\pi_j > 0)_i$ is the increase in purchase probability caused by impression i for consumer j .

Nowadays, most online ad impressions are auctioned off in second-price sealed bid auctions (Arnosti et al., 2016). We know that for these types of auctions the optimal bid for an impression ($bid_{i,j,optimal}$) equals the firm's true valuation of the impression ($v_{i,j}$), given that one ad is auctioned off at a time (Edelman et al., 2007), as common in display advertising. This means that the firms optimal bidding strategy for an impression follows:

$$bid_{i,j,optimal} = v_{i,j} = \Delta P(\pi_j > 0)_i \times \pi_j$$

Incentive Misalignment for Ad Platforms. Although there seems to be a well established academic understanding of the economic value of advertising (Johnson et al., 2017a,b; Lewis and Rao, 2014; Li and Kannan, 2014), unfortunately, most ad platforms and advertisers cannot calculate or report the increase in a consumer's purchase probability caused by advertising, $\Delta P(\pi_j > 0)$. Thus the most obvious metrics that get reported by ad platforms are aggregate ad performance measures such as conversion rate - which is the percentage of consumers addressed with ads that conducted a predefined target action. For example, ad platforms report how many purchases were generated per consumer that was exposed to ads:

$$Performance = \frac{n_{purchases|I=1}}{n_{j|I=1}}$$

where $n_{purchases|I=1}$ is the number of purchases conducted by consumers that were addressed with ads and $n_{j|I=1}$ is the number of consumers that were addressed with ads. This measure is commonly referred to as conversion rate. The conversion rate measures the average purchase probability $P(\pi > 0)$ of consumers j that are addressed with ads. This operationalization of ad performance by ad platforms overstates the effect of ads, as some consumers would have conducted a purchase independent of being confronted with advertising. Both the actual increase in purchase probability $\Delta P(\pi > 0)$ as well as the profit contribution π are not typically incorporated into this measure.

Some ad platforms report profit contributed by advertising as:

$$Return_{reported} = \sum_{j=1}^j (P(\pi_j > 0) \times \pi_j)$$

assuming homogeneous profit contribution, where every purchase generates the same amount of profit, this equals:

$$Return_{reported} = P(\pi > 0) \times \sum_{j=1}^j \pi_j = P(\pi > 0) \times n_{j|I=1} \times \bar{\pi}$$

While the actual *Return* is:

$$Return_{actual} = \sum_{j=1}^j (\Delta P(\pi_j > 0) \times \pi_j) = \Delta P(\pi > 0) \times n_{j|I=1} \times \bar{\pi}$$

Firms aim to reward ad platforms for actual success contribution. This is also why firms have moved from paying for ad impressions, to paying for when consumers click on ads, to paying only for when consumers addressed with ads conduct a pre-defined target action, usually a purchase. Such an incentive structure is implemented in CPA-based contracts. Firms have now commonly implemented CPA-based contracts, in which they set this CPA value as their willingness to pay for a purchase reported by the ad platform. Nevertheless, firms start to understand that they should reward ad platforms at most for the increase in purchases caused by ads ($\Delta P(\pi > 0)$) and not for all purchases conducted by consumers that have been addressed with ads ($P(\pi > 0)$). Still, as the CPA-based contracts reward ad platforms based on $P(\pi > 0)$, firms face the need to apply a deflation factor to adjust for the platform's over-reporting of returns to not overpay the ad platform.

One common way to do so is to set the *CPA* value, representing a firm's willingness to pay for a purchase reported by the ad platform, below the absolute profit contribution of a purchase, as the ad platform triggered likely only part of consumers' purchase actions ($\Delta P(\pi_j > 0)$). Instead of paying the ad platform ($Return_{reported}$) firms set the *CPA* value to adjust for ad platforms over-reporting with the aim to set the maximum willingness to pay equal to $Return_{actual}$. The maximum *CPA* value can be defined as:

$$CPA_{max} = \frac{\Delta P(\pi > 0) \times \bar{\pi}}{P(\pi > 0)}$$

This restructuring of the formula for $Return_{reported}$ is possible under the assumption of homogeneous profit contributions ($\pi_j = \bar{\pi}$). Although this seems, at a first glance, like an adequate way for firms to adjust for the over-reporting of return on ad spend by the ad platform, this adjustment proves difficult in practice for several reasons:

1. Firms can usually not determine $\Delta P(\pi > 0)$ but in best cases have access to historical information from experiments conducted at some point in the past. Therefore, firms establish an approximate value for *CPA*.
2. The possibility to calculate a maximum *CPA* value hinges on the assumption that the above equation is static. It is more likely though, that the values for both $\Delta P(\pi > 0)$ and $P(\pi > 0)$ change dynamically. As firms cannot set their *CPA* value dynamically but need to pre-define a static *CPA* value this deflation approach can never offer an optimal solution.
3. Most notably, the deflation approach does not address the actual issue, the way that ad platforms are rewarded. Since the ad platform's revenue function remains:

$$Revenue_{ad\ platform} = P(\pi > 0) \times n_{j|I=1} \times CPA$$

the ad platform's interest to target consumers with high purchase probabilities independent of the effect of ads remains intact. As the revenue function of the ad platform does not contain $\Delta P(\pi > 0)$, there is no incentive for the ad platform to address consumers that are highly affected by ads. Increasing availability of individual-level consumer information and analytical capabilities allow ad platforms to identify consumers' inherent purchase probabilities and target them selectively with ads⁵.

⁵We show in Appendix A4.1 Table A4.1 that ad platforms seem to make use of consumer characteristics and behavior to define the bid that they place for an ad impression for a consumer.

However, whether such targeting by ad platforms is beneficial or harmful for firms is not straightforward. It depends on whether high purchase probability consumers are more or less receptive towards ads than low purchase probability consumers. In short, such targeting benefits firms only if $\Delta P(\pi_j > 0)$ is positively correlated with $P(\pi_j > 0)$. This becomes clear when looking at the formula that firms use to define their CPA_{max} . A sole increase in $P(\pi > 0)$ without a simultaneous increase in $\Delta P(\pi > 0)$ would mean that the identified CPA value is inflated if not readjusted. On the other hand, if $\Delta P(\pi > 0)$ increases (over-)proportionally with $P(\pi > 0)$ this targeting would actually be not harmful (beneficial). In short, if

$$corr(P(\pi_j > 0), \Delta P(\pi_j > 0)) \leq 0$$

the ad platform has an incentive to target consumers with ads, by bidding higher for their ad impressions, that are not more or even less receptive to these ads. The adequacy of the deflation approach, using adjusted CPA values, does highly depend on this correlation.

Incentive Misalignment with Heterogeneous Profit Contributions. For firms that have a heterogeneous portfolio of products with heterogeneous profit contributions π_j , the determination of a deflated CPA value becomes even more challenging. Commonly, a CPA value, the firm's willingness to pay an ad platform for a reported purchase, is defined across the product portfolio, not distinguishing between different products and their heterogeneous profit contributions. This essentially means that the firm is applying a binary reward system in which the question is simply whether a purchase occurred or not, instead of what the profit contribution of each purchased product is. To derive an estimate for a suitable CPA value the firm needs to make an assumption about the average profit contribution of products that will be purchased. This means, not only unknown variation in $\Delta P(\pi_j > 0)$ but also in profit contribution π_j makes the approximation more difficult. Firms cannot factorize out consumers' absolute purchase probability after having seen an ad in the formula for $Return_{reported}$ without making an assumption about the average profit contribution $\bar{\pi}$ to come up with an approximated CPA value. Such an approximation might become especially problematic when there is a negative correlation between π_j and $P(\pi_j > 0)$.

$$corr(P(\pi_j > 0), \pi_j) < 0$$

In this case, the ad platform has an incentive to target consumers that have lower profit contributions inflating the firm's *CPA* approximation. The assumption that there is a negative correlation between the purchase probability and the profit contribution seems plausible - consumers that are looking for high involvement (i.e. more expensive, complex, and durable) products are less likely to purchase compared to consumers that are in the market for low-involvement goods that usually require a less intensive information search (Gu et al., 2012). This means that for firms that offer heterogeneously priced products, the potential incentive misalignment might have an even more severe effect if ad platforms target consumers with lower profit contribution as they tend to be more likely to purchase. This might drive down the overall profit contribution generated by consumers being addressed with ads. Therefore, it is not only crucial for firms to think about the actual purchase probability uplift generated by an ad campaign but also whether the ad platform might have an incentive to target less profitable consumers as the ad platform is rewarded in a binary manner independent of the monetary value of a purchase.

4.4 Ad Allocation Process and Infrastructure

In this work, we describe the implications of incentive specifications in contracts between firms and, what we call, ad platforms that guide the allocation process of digital ads to consumers in programmatic advertising. These two parties represent the demand side of digital advertising, where firms purchase ad space and ad platforms act as the trading platform for this ad space. What we describe as ad platform consists of several distinct systems. While we are focusing on the case of an integrated ad platform that comprises of all these systems, there are several players in the industry that provide only part of the functionality of an integrated ad platform. As we are looking at an integrated ad platform that is encompassing several roles and functionalities in the ad allocation process, we refrain from dividing the ad platform into its distinct technical roles in the main part of the paper. Nevertheless, in this section we aim to give a better understanding of the ad allocation process from a more technical perspective.

Importantly, in this work we are not taking the supply side in programmatic digital advertising into account, consisting of publishers that sell ad space to firms via ad platforms. These publishers, e.g. news websites, that sell ad space, are also connected

to ad platforms, often through so called supply side platforms (SSP), to trade their ad space.

The delivery process of ads in programmatic digital advertising is technically complex involving several systems and providers that need to work together efficiently to allow the execution of serving ads in milliseconds from an ad slot becoming available to eventually serving an ad.

Commonly, the demand side of the ad allocation process is divided into the 4 main functional parts consisting of (1) *ad execution*, (2) *data enrichment*, (3) *analytics and measurement*, and (4) *content optimization* (Fisher et al., 2016). In the *ad execution* phase the so called *ad server* takes the responsibility to decide which ad will be served to a consumer on a publisher's website. Further, the *ad server* records consumers' reactions to ads. These consumer reactions are communicated to firms via reporting platforms that run under the *analytics and measurement* functionality. The *data enrichment* functionality allows firms to incorporate information on consumers that might be valuable in the process of deciding which consumer to address with ads. Such information can include internally recorded information, i.e. consumer behaviour on the firm's website, but also externally received information. This type of information is processed in the *data management platform (DMP)*. By using this data to categorize consumers, ad platforms can automatically define which consumers should be targeted and which ad creative should be shown to a respective consumer. The *content optimization* functionality encompasses a *content delivery network (CDN)*. This *CDN* is storing all ad creatives from which the *ad server* will eventually choose the creative matching the criteria defined.

So-called *demand side platforms (DSPs)* span over the functionalities (1) *ad execution*, (2) *data enrichment*, and (4) *content optimization* and are essentially the main platform that firms use to coordinate their programmatic advertising efforts. "A DSP is defined as an infrastructure that enables advertisers and agencies (buyers) to manage their media buying via a single platform" (Graham, 2017). Usually, a *DSP* is connected to several *ad exchanges / ad networks*, on which ad space is being traded, allowing firms to manage their ad distribution to different publishers centrally as ad inventory from several publishers can be accessed via a single platform. In our case, we focus on the ad allocation via auctions, where ads are being sold off using *real-time bidding*, compared to *programmatic direct* where ads are being traded directly in bulk for a fixed price without the use of auctions. In real-time bidding, the *DSP* submits bids for ad impressions on behalf of the firm where the *supply side platform (SSP)* determines the winner, based on the highest bid submitted. The *DSP* then

communicates the ad serving instructions with the *SSP* and makes sure that the right ad is delivered from the *CDN*. The *DSP* provides so-called bidding agent services (Johnson and Lewis, 2015) to the firm where the *DSP*, as part of what we call the *ad platform*, is optimizing the bidding for ad impressions on behalf of the firm.

When talking about an ad platform in this work, we refer to an integrated player that combines all of the above described functionalities on the demand side of programmatic advertising.

4.5 Payment and Pricing for Ad Impressions

In our data analysis section, we are referring to the cost per ad impression that the firm is paying instead of the *cost per acquisition (CPA)* that the firms specifies to incentivize the ad platform. In practice, ad platforms need to reconcile different types of contracts (e.g. CPM, CPC, CPA) that can simultaneously compete for the same ad impressions. To achieve this, the ad platform needs to rescale all different bidding systems to the cost per thousand impressions unit, allowing the ad platform to bid for impressions across different contracts by using the same unit. The adjustment to this common unit works as follows:

$$CPM = CPA \times CR \times 1,000$$

where *CR* stands for conversion rate, giving the average share of ad impressions that leads to a purchase. This allows the ad platform to rescale the from the firm defined willingness to pay for a reported purchase (*CPA*) to an average willingness to pay for 1,000 impressions (*CPM*). The ad platform then differentiates the bidding for individual ad impressions while keeping the average cost per 1,000 impressions (*CPM*) equal to the above described equation. This allows the ad platform to discriminate between the bids for ad impressions for individual consumers while staying within the boundaries of the firm-defined *CPA* value.

4.6 Experiment & Analysis

In collaboration with a major European e-retailer, we conducted a large-scale randomized field experiment to investigate the effectiveness of their retargeting advertising. Retargeting makes use of consumers' browsing behavior on the firm's website to show

products that a consumer has browsed on external websites (Bleier and Eisenbeiss, 2015a; Lambrecht and Tucker, 2013). This ad targeting strategy clearly introduces an upfront sample selection as consumers that are being targeted have shown interest in the firm’s products prior to being addressed with ads. This is also why this context is especially valuable to investigate a potential incentive misalignment between firm and ad platform. In this context, ad platforms have richer data on consumers that can be used to target consumers and optimize the bidding for ad impressions.

We randomly allocate consumers that have visited at least a product category page of our partner firm’s website to being treated with retargeting ads on external sites (80% probability of assignment) or with public service announcement (PSA) ads (20% probability of assignment), that advertise the donation to a charity⁶. We follow the commonly accepted assumption that PSA ads are orthogonal to our retargeting ads in their effect on consumers, allowing us to identify the consumers in our control group and measure the causal impact of retargeting ads (Johnson et al., 2017a). For the duration of the experiment consumers remain in the treatment or control group respectively.

Our partner firm selects a CPA optimization algorithm to guide the bidding for the ads on the ad platform. We are aware of a potential selection of consumers that might be introduced into our sample caused by the optimization algorithm. This can be the case if the optimization algorithm selects different types of consumers into our treatment and control group, biasing our ad effectiveness estimates. This happens if the bidding for the retargeting and PSA impressions is optimized separately. In this case, the ad platform has an incentive to select different types of consumers for the treatment and the control group, as different types of consumers are likely to respond positively to the retargeting ads than to the public service announcements (Johnson et al., 2017a). To circumvent this issue, we pay significant attention to setting up our campaigns in a way so that they are optimized jointly by the ad platform before we start the experiment. After the end of our experiment and our data collection, we run an extensive amount of randomization checks to make sure that consumers in the control and treatment groups are not systematically different and that the optimization was run jointly.

For our analysis we aggregate our data on consumer level ($n = 20,918$, $n_{treatment} = 16,734$, $n_{control} = 4,184$) (see Table 4.1 for descriptive statistics). In our randomization checks we find that there seems to be no significant differences between treatment

⁶We make sure that the share of consumers allocated to treatment (80%) and control group (20%) is consistent independent of the height of the bid. More details can be found in Appendix A4.2

and control group regarding both variables related to treatment and consumer characteristics (see Table 4.2). We find that the average bid (in Euro) ($mean_{bid} = .011$) ($\Delta M = .0003$, $t = 1.081$, $p = .280$), as well as the average cost for an ad impression (in Euro) ($mean_{cost} = .0036$) ($\Delta M = 4.29e - 6$, $t = .063$, $p = .950$) do not differ between control and treatment group. Next to that, the number of impressions served on average per consumer ($mean_{impressions} = 20.095$) does not differ between control and treatment group ($\Delta M = .213$, $t = 0.778$, $p = .437$).

We also investigate whether different types of consumers might have been allocated to treatment or control group. For this purpose, we assess consumer behavior before the first ad treatment, as this behavior is influenced by our experimental treatment and therefore not comparable after the first treatment. We find that the number of activities conducted on our retailer’s website before consumers’ first ad impression ($mean_{activities} = 5.826$) does not differ significantly between treatment and control group ($\Delta M = .045$, $t = .290$, $p = .772$). There is no difference in the number of visits between consumers in the treatment and control group ($mean_{visits} = 1.361$) ($\Delta M = .024$, $t = 1.268$, $p = .205$), a variable indicating how often consumers visit our partner firm’s website before their first experimental treatment. Also the time spent on our retailers site before the treatment (in minutes) ($mean_{activity\ duration} = 6.971$) does not differ significantly ($\Delta M = .247$, $t = .922$, $p = .357$). Further, variables such as the number of visited product categories, number of visited product pages, as well as the number of shopping cart visits, all characterizing consumers, do not differ significantly between control and treatment group. These findings make us confident that although the platform has been running an optimization algorithm for the bidding, no selection bias was introduced as treatment and control group were optimized simultaneously.

Table 4.1: Descriptive Statistics

Variable	Observations	Mean	Std. Dev.	Min	Max
impressions	20,918	20.0947	15.8761	1	111
purchase	20,918	0.0564	0.2307	0	1
ad treatment	20,918	0.8000	0.4000	0	1
bid	20,918	0.0109	0.0160	0.00054	0.09118
cost	20,918	0.0036	0.0039	0.00001	0.08946

We move on to investigate whether the ad platform’s targeting mechanisms which become explicit in the bids for impressions that we observe in our data follow an optimal strategy for the firm.

Table 4.2: Randomization Checks

Variable	Mean		t-Statistic	p-Value
	Control	Treatment		
bid	0.011	0.011	1.081	0.280
cost	0.004	0.004	0.063	0.950
impressions	19.924	20.137	-0.778	0.437
activities	5.790	5.835	-0.290	0.772
visits	1.341	1.366	-1.268	0.205
activity duration	6.773	7.020	-0.922	0.357
number of visited product categories	1.484	1.476	0.504	0.615
number of visited product pages	2.885	2.902	-0.191	0.849
number of shopping cart visits	0.100	0.101	-0.063	0.950

4.6.1 Bidding & Increase in Purchase Probability

We investigate whether the optimization algorithm implemented by the ad platform is bidding in the interest of the firm. First, we assess whether ad platforms target consumers with higher purchase probabilities in line with the incentive structure present in contracts between ad platforms and firms. Generally, our analysis requires variation in the endogenous variable for average bids for ad impressions placed by the ad platform on behalf of the firm. The histogram of the average bid for an ad impression for a consumer displays the variation in the bid variable (see Figure 4.1). A coefficient of variation (ratio of standard deviation to mean) larger than 1 ($CV = 1.470$) of the bid variable further points to variation in the height of bids placed by the ad platform.

Next, we investigate the average treatment effect of the ads in our experiment. Lastly, we test whether higher bids are placed for ads that have a higher impact on consumers' purchase probabilities.


To do that, we estimate a logit model that investigates whether consumers for whom the optimization algorithm does on average bid higher are more receptive to the ad treatment.

$$Purchase\ Probability_j^{Retargeting} = \frac{\exp(U_j^{Retargeting})}{\exp(U_j^{Retargeting}) + 1}$$

$$U_j^{Retargeting} = \beta_0 + \beta_1 bid_j + \beta_2 ad\ treatment_j + \beta_3 bid_j \times ad\ treatment_j + \epsilon_j$$

where bid_j gives the average bid for a consumer j over the duration of the experiment in Euro, $ad\ treatment_j$ represents a binary variable indicating whether a consumer was

Figure 4.1: Histogram Average Bid



addressed with retargeting ads ($ad\ treatment_j = 1$) or PSA ads ($ad\ treatment_j = 0$), and ϵ_j represents the idiosyncratic error term.

The coefficient of the interaction between bid_j and $ad\ treatment_j$ represents the focal aspect in this analysis. In case the ad platform does optimize the bidding in the interest of the firm, we would expect a positive and significant coefficient for the interaction term of bid_j and $ad\ treatment_j$.

Table 4.3 presents the results for our analysis. We find that, in line with the incentive structure in contracts between ad platforms and firms, ad platforms target consumers that are more likely to purchase by bidding higher for their ad impressions ($\beta_{bid} = 14.852$, $p < .001$). In Appendix A4.1 we present evidence that shows that the ad platform is using consumer characteristics to identify high purchase probability consumers. Overall, the ad treatment does significantly increase consumers' purchase probabilities ($\beta_{ad\ treatment} = .174$, $p = .032$) pointing towards the presence of a return on advertising spending. Nevertheless, we do not find evidence for a significant effect of the interaction between bids placed by the ad platform and the ad treatment ($\beta_{ad\ treatment \times bid} = .531$, $p = .880$). This means that while the optimization algorithm is bidding higher for consumers that are more likely to purchase, it fails to identify and bid higher for consumers that are more receptive towards ads. The non-significant interaction term of $ad\ treatment_j$ and bid_j also points towards no significant correlation between consumers' inherent purchase probabilities and their receptiveness towards ads rendering the bidding strategy adopted by the ad platform sub-optimal for the firm.

Table 4.3: Logit Regressions for Purchase Probability

VARIABLES	(1) purchase Logit	(2) purchase Logit	(3) purchase Logit
bid	14.852*** (1.377)	14.895*** (1.378)	14.466*** (3.158)
ad treatment		0.174** (0.081)	0.165* (0.098)
ad treatment \times bid			0.531 (3.510)
Constant	-3.071*** (0.038)	-3.213*** (0.077)	-3.206*** (0.089)
Observations	20,918	20,918	20,918

Standard errors in parentheses
 *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

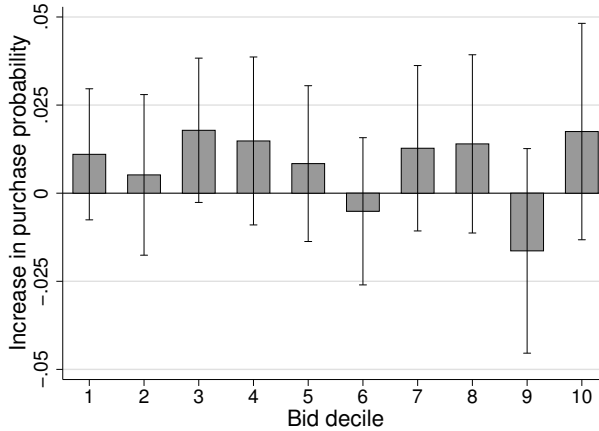
We run additional robustness checks in which we analyze the relationship between the maximum bid, the median bid, and the cumulative bid placed for an ad impression for consumer j and their impact on consumers' purchase probability (see Appendix A4.3). When operationalizing the ad platform's targeting behavior, visible in the ad platform's bidding for ad impressions, differently, we find consistent results.

To strengthen the argument that there is no significant correlation between consumers' inherent purchase probabilities and the increase in purchase probabilities caused by ads we build a predictive model estimating consumers' purchase probabilities prior to the ad treatment (see Appendix A4.4). The purpose of this analysis is to investigate the relationship between absolute purchase probabilities and the increase in purchase probabilities caused by ads more directly, instead of looking at the bid placed by ad platforms and the increase in purchase probabilities caused by ads. We use these predicted purchase probabilities to test whether consumers with a higher predicted purchase probabilities react more positively to our ad treatment. Consistent with our claim, we find that consumers with higher predicted purchase probabilities do not react more positively to the ad treatment.

To further assess the robustness of our findings, we investigate the possibility that we do not find a significant interaction effect in our analysis because our experiment has too low power. We assess the functional form of the relationship between bids placed by the ad platform and the increase in consumers' purchase probabilities. We plot the average increase in purchase probabilities caused by ads per consumer and

bid decile (see Figure 4.2). In case the ad platform is optimizing the bidding in the interest of the firm, we would expect a monotonically increasing trend in the average increase in purchase probability caused by ads with an increase in the bid decile. This graph points towards the absence of a relationship between the bidding conducted by the ad platform and the increase in consumers' purchase probabilities.

Figure 4.2: Average Purchase Probability Increase per Consumer and Bid Decile



Notably, firms do not pay their actual winning bid to serve an ad impression but are being charged the second highest bid in the ad auctions. Our results remain consistent when analyzing the impact of both the average cost per impression and the overall cost for impressions served to a consumer over the duration of the experiment (see Appendix A4.5). Firms are paying more for ad impressions that do not increase consumers' purchase probabilities more significantly. This means that firms pay more for ads that do not deliver significantly higher value to them.

One explanation that would justify the ad platform's bidding behavior is that the ad platform's optimization algorithm incorporates consumers' profit contributions (π_j) into its optimization. In case the profit contribution is negatively correlated with the increase in consumers' purchase probabilities:

$$\text{corr}(\pi_j, \Delta P(\pi_j > 0)) < 0$$

there might be a valid reason for the algorithm to not bid higher for more receptive consumers but instead target consumers with higher profit potential.

4.6.2 Bidding & Profit Contribution

To analyze whether the optimization algorithm does take profit contribution into account we investigate whether consumers that purchase more expensive products have received higher bids. The generated revenue seems to be an adequate measure for this analysis as our retailer's profit margin is a somewhat stable percentage of the generated revenue. The generated revenue per consumer can therefore be seen as a linear transformation of the profit contribution.

First, we analyze, for the consumers that conducted a purchase $n_{purchase} = 1,113$, whether the ad platform bid higher to serve ad impressions to consumers that generated higher revenue for the firm. We run the following OLS regression:

$$revenue_j = \beta_0 + \beta_1 bid_j + \epsilon_j$$

We do not find that more valuable consumers receive higher bids, pointing towards the optimization algorithm not taking the profit contribution into account (Table 4.4, column 1). Due to the positively skewed distribution of our revenue variable we re-run our OLS model using the log-transformed revenue variable as dependent variable (Table 4.4, column 2). Still, we do not find that the algorithm bids higher for more valuable consumers.

Table 4.4: OLS Regressions for Revenue

VARIABLES	(1) revenue OLS	(2) log(revenue) OLS
bid	135.092 (394.685)	2.175 (1.403)
Constant	255.859*** (12.494)	4.922*** (0.041)
Observations	1,113	1,113
R-squared	0.0001	0.0019

Robust standard errors in parentheses

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

Nevertheless, in this analysis we solely focus on the small share of consumers that conducted a purchase and the non-significant values might simply represent the challenge for the algorithm to predict not only the profit contribution but also whether

a consumer will purchase or not.⁷ Therefore, we move on to investigate whether the algorithm picks up earlier signs of profit contribution potential in a consumer's purchase process and incorporates them in its optimization.

We create a measure for the value of a product category within our sample. For each product page that the consumers in our sample browse, we collect the price of the featured product. We then calculate the average price of a product within a product category. In a next step, we use information on which categories a consumer has browsed to calculate a weighted average representing the potential for revenue contribution for an individual consumer.⁸ For example, if the average price for a product in the laptop category is 1,000 Euro and the average price for a product in the TV category is 500 Euro, and a consumer has browsed 5 pages in the laptop category and 5 pages in the TV category, we estimate the potential revenue contribution at 750 Euro. To check whether our measure for potential revenue contribution is a good predictor for revenue we run the following model:

$$revenue_j = \beta_0 + \beta_1 revenue\ potential_j + \epsilon_j$$

We find that our measure for revenue potential is a significant predictor of revenue contribution ($\beta_{revenue\ potential} = .018$, $p < .001$) (Table 4.5, column 1). In a next step, we analyze whether higher revenue potential does also predict the average height of a bid for a consumer that is placed by the optimization algorithm.

$$bid_j = \beta_0 + \beta_1 revenue\ potential_j + \epsilon_j$$

In column 2 of Table 4.5 we see that this is not the case. Our variable for revenue potential is not a significant positive predictor for a consumer's average bid for an ad impression. Therefore the ad platform's optimization algorithm does not abstain from targeting more receptive consumers for the sake of targeting consumers with higher profit contributions. Instead, our previously raised concern that the ad platform might target consumers not only with higher inherent purchase probability but also lower

⁷In Appendix A4.6 we show additional analysis confirming our results.

⁸The data used to derive the revenue potential variable was collected some time after the experiment had been executed. By this time, some products browsed by consumers were not available on the partner firm's website anymore. Therefore, we are missing values of the revenue potential for 137 consumers that browsed products in unpopular categories that were discontinued in between the experiment and collection of data to estimate the revenue potential of consumers. It seems reasonable to assume that there is no relationship between the likelihood that the sale of a product is discontinued and the targeting behavior of the ad platform that could influence our results. Therefore, we run the analysis with the slightly smaller sample.

revenue potential seems to be supported. We find a negative and significant coefficient for the effect of our revenue potential variable on the average bid placed by the ad platform on behalf of the firm ($\beta_{revenue\ potential} = -9.15e - 7, p = .049$). This finding indicates that, as consumers that are browsing for more expensive products are less likely to conduct a purchase, the ad platform targets consumers with lower revenue potential that have a higher purchase probability.

Table 4.5: OLS Regressions Revenue Potential

VARIABLES	(1) revenue OLS	(2) bid OLS
revenue potential	0.018*** (0.004)	-9.15e-7** (4.65e-7)
Constant	8.007*** (1.126)	1.11e-2*** (2.00e-4)
Observations	20,781	20,781
R-squared	0.0017	0.0002

Robust standard errors in parentheses

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

4.7 Theoretical Contributions

This study uniquely contributes to research in the area of behavioral targeting, ad effectiveness, and ad auctions. While previous studies pointed out the inadequacy of commonly used incentive schemes in programmatic advertising contracts as advertising platforms have an interest to serve ads to consumers that conduct a purchase independent of the effect of ads (Johnson and Lewis, 2015), we are able to investigate the implications of these incentives empirically. Our unique field experimental set up allows us to both identify the causal impact of digital advertising on consumers' purchase probabilities while simultaneously exploiting variation introduced in the bids for ad impressions by the ad platform that is targeting consumers.

We show that the ad platform is acting in accordance with incentives defined in the contract. Ad platforms target consumers with higher inherent purchase probabilities by placing higher bids on ad slots for these consumers. This finding confirms the presence of a potential incentive misalignment.

Up to now, it remained unclear how consumers with different inherent purchase probabilities differ with respect to their reaction to advertising. In line with this question, we show that firms efforts to mitigate this potential incentive misalignment by reducing rewards for purchases reported by the ad platform hinges on the correlation between the absolute inherent purchase probabilities of consumers and the increase in purchase probabilities caused by ads. In case this correlation is zero or non-positive, ad platforms have an interest to target consumers with ad impressions in a way that is sub-optimal for the focal firm.

To assess the correlation between absolute inherent purchase probabilities of consumers and their increase in purchase probabilities caused by ads we empirically assess whether the purchase probabilities of consumers that receive higher bids increase more significantly. This assessment is crucial when answering the question of whether an actual incentive misalignment between firm and ad platform is present in the context of programmatic advertising.

In our empirical analysis, we find evidence for the presence of the incentive misalignment that has negative implications for firms. While ads do generally increase consumers' purchase probabilities and ad platforms target consumers with ads that are more likely to purchase, independent of being served an ad impression, ad platforms do not target consumers that are more receptive to ads. This means, ad platforms do not target consumers whose purchase probabilities increase more significantly due to ad impressions. We find no evidence for a positive correlation between consumers' high inherent purchase probability and their increase in purchase probability caused by ads.

We further show, that due to the bid optimization conducted by the ad platform the firm is paying more for ad impressions that do not offer a higher return on ad investment. This finding renders the ad allocation process conducted by the ad platform on behalf of the firm sub-optimal for the firm. This work is the first to provide actual empirical evidence for the presence of an incentive misalignment between firms and ad platforms in programmatic advertising.

Our results remain consistent when considering the potential that ad platforms might take consumers' profit contribution potential in their ad allocation process into account. As ad platforms become better at using individual-level consumer data to predict consumers' purchase probability this problem might be enhanced.⁹

⁹We show in Appendix A4.1 Table A4.1 that the ad platform seems to make use of variables describing consumer characteristics and behavior to define their respective bids for an ad impression i for consumer j .

4.8 Practical Implications

The results of this research have major implications for firms that utilize automated ad allocation tools provided by ad platforms. While ad platforms claim to optimize the ad allocation process in the interest of firms, we find that this is the case only when there is a positive correlation between consumers' purchase probability and their increase in purchase probability caused by ad impressions. In our empirical context, we show that this is not the case, rendering the ad allocation process conducted by the ad platform on behalf of the firm sub-optimal.

Aggravatingly, firms usually do not have the experimental data available that is necessary to reveal this sub-optimal bidding behavior performed by an ad platform. This issue arises as the ad platform is controlling the whole ad allocation process including the selection of consumers, definition of bid, display of ads, as well as the report of success metrics. Firms usually have no access to both experimental data that allows for an investigation of the true impact of digital ads as well as the detailed data available on consumers to assess potential heterogeneity that could be exploited for targeting.

To resolve this incentive issue, ad platforms would need to disintegrate the different steps in the ad allocation process, giving the firms more control over and insights into whom to serve an ad at which price. Further, ad platforms would need to be rewarded based on incremental purchases instead of the absolute number of purchases conducted by consumers addressed with ads.

4.9 Limitations

Despite the fact that our study offers some interesting insights into the implications of incentives specified in contracts within programmatic advertising, some limitations need to be taken into consideration when interpreting our results.

When looking at the data that we analyze for our experiment, it is important to mention that we only observe bids for ad slot auctions won by the focal firm. This means that we only know how much the ad platform is bidding for impression i for consumer j if the firm wins the auction for the respective ad impression. As lower bids are less likely to win an auction for an ad impression, our data could be truncated at the lower end of bids. We might not observe lower bids placed by the ad platform as these auctions are not won by the firm. Therefore, it could be the case that consumers receiving higher bids for impressions are the ones more responsive to

ads. For consumers receiving low bids, that we are more likely to miss in our data as the firm is more likely to lose these ad auctions, the treatment could be less effective. Therefore, our data would simply not reveal a pattern where consumers that receive higher bids are more receptive to ads. Even though, theoretically, such a problem might exist, we think that this issue is of limited nature when looking at our data set.

First, we find that the coefficient of variation (ratio of standard deviation to mean) for the average bid for consumers is higher than 1 ($CV = 1.470$). A coefficient of variation higher than 1 indicates rather large variation in a variable. Therefore, even if we do not observe the lowest bids placed by the ad platform, there seems to be sufficient variation that can be exploited for our analysis. The question of whether ads with high bids have a higher impact on consumers' purchase probabilities – given that the ad auction is won – remains.

Further, when looking at the distribution of average bids per consumer placed by the ad platform (see Figure 4.1), we see that the bids we observe are clustered relatively close to zero, while we observe fewer high bids. The distribution of average bids that we observe does stand counter the argument that our data might be truncated at the lower end of bids as these is where our bids are concentrated.

Lastly, as we investigate the effectiveness of advertising for different bids in the context of retargeting, it seems questionable whether the ad platform would bid significantly lower for a certain type of consumer. The reason for this assumption is related to the fact that consumers that visit our partner firm's website – which is required to be eligible for treatment with retargeting ads which use consumers' browsing behavior for personalization – do have inherently higher purchase probabilities compared to consumers that do not visit the partner firm's website. This should mean that the ad platform does not have an incentive to bid very low for consumers that have indicated their interest in our partner firm's products, making it less likely to systematically lose auctions, truncating our data set.

Another issue we are facing is that our experimental design is not taking the so called 'defensive effect' of ads into account. As pointed out in former research (Johnson et al., 2017a), using PSAs as control group for the ad treatment does not represent the true counterfactual. The true counterfactual encompasses the possibility that consumers are confronted with ads of competitors. The defensive effect of ads describes that ads draw effectiveness not only from their direct impact on consumers' purchase probabilities but also from the fact that they prevent competitors from displaying their ads. Our experimental design does not take this defensive effect into account. There are two empirical challenges related to this limitation. First, as we do not observe the

defensive effect with our design, we might underestimate the overall impact of ads on consumers' purchase probabilities. This issue seems negligible in our context as we do find that ads generally increase consumers' purchase probabilities. The second issue is related to the identification of heterogeneous effects on consumers' purchase probabilities. In case the defensive effect is not symmetric for different bids, e.g. smaller for lower bids and bigger for larger bids, we might not be able to detect a significant increase in ad effectiveness with an increase in the average bid as we are not observing the defensive effect.

To handle this issue, we conduct some additional analyses in which we try to control for the competition for an ad impression (see Appendix A4.7). A measure for competition should allow us to assess to what extent a defensive effect might be asymmetric for different heights of bids. Both our measures for absolute and relative competition indicate that competition for ad impressions that receive higher bids is relatively lower. This indicates that if there is an asymmetric defensive effect, it is likely to decrease with an increase in bids. This could mean that we are not able to detect a decrease in ad effectiveness with higher bids, rendering the bidding optimization conducted by the ad platform even less favorable for the firm. When controlling for absolute and relative competition in our analysis, we find results consistent with our main analysis. Higher bids do not yield higher ad effectiveness. Therefore, we are confident that not being able to consider the defensive effect of ads in our analysis is a problem that is not significantly affecting the contributions of our work.

Appendices

A4.1 Bid Prediction with Consumer Characteristics & Behavioral Data

In Table A4.1 we use a set of consumer characteristics and behavioral variables to see how well we can predict the average *bid* for a consumer *j* placed by the ad platform. We find that these variables do significantly predict the height of the average bid and that we are able to explain 27.7% of the variation in the average *bid* and 33.3% of the variation in the log transformed version of the average *bid*. This points towards the ad platform making use of these type of variables, which characterize consumers, to decide on how much to bid for an ad impression.

Table A4.1: Bid Prediction with Behavioral Data

VARIABLES	(1)	(2)
	bid OLS	log(bid) OLS
category browsing	-0.001053*** (0.000040)	-0.046387*** (0.001512)
product browsing	-0.001338*** (0.000039)	-0.064929*** (0.001451)
focus country 1	0.001358*** (0.000039)	0.068786*** (0.001520)
focus country 2	0.001017*** (0.000038)	0.041928*** (0.001439)
number of activities	0.000049*** (0.000012)	0.000572 (0.000639)
number of visits	-0.000366*** (0.000057)	-0.011718*** (0.002907)
activity duration in minutes	0.000020*** (0.000006)	0.001932*** (0.000338)
number of visited categories	0.000212 * * (0.000083)	0.015485*** (0.004267)
Constant	0.012331*** (0.000199)	-4.901002*** (0.010082)
Observations	20,918	20,918
R-squared	0.276911	0.333471

Robust standard errors in parentheses

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

A4.2 Share of Consumers Allocated to Treatment and Control Group

To make sure that consumers in treatment and control group are consistently distributed per bid and in accordance with our randomization rule (80% of consumers allocated to treatment group, 20% of consumers allocated to control group) we analyze the allocation of consumers per bid decile. To do so, we run two types of analysis. First, we conduct a one-way analysis of variance (ANOVA) on our ad treatment variable investigating whether there are significant differences between the share of consumers that are allocated to the treatment for the different bid deciles. The analysis is not significant, $F(9, 20,908) = 1.37$ ($p = 0.197$), meaning that the share of consumers allocated to the treatment group per bid decile does not differ significantly.

Next, we run t-tests per bid-decile in which we test the null hypothesis whether the share of consumers allocated to the treatment group per respective bid decile differs significantly from 80%. Figure A4.1 plots our results. We find that there is no significant difference to the 80% allocation rule for any of or bid deciles.

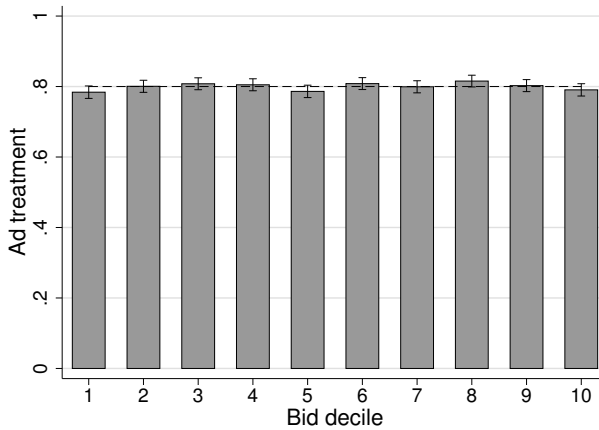


Figure A4.1: Share of Treated Consumers per Bid Decile

A4.3 Estimation of Purchase Probability Using Alternative Measures for Bid

We re-estimate our main model by replacing the average bid placed for ad impressions for a consumer by the maximum bid (see Table A4.2), the median bid (see Table A4.3), and the cumulative bid (see Table A4.4) placed for a consumer over the duration of the experiment. Our results are consistent with the main analysis.

Table A4.2: Logistic Regression Predicting Purchase Probability Using the Maximum Bid per Consumer

VARIABLES	(1) purchase Logit	(2) purchase Logit	(3) purchase Logit
max bid	11.713*** (1.107)	11.736*** (1.108)	13.240*** (2.491)
ad treatment		0.171** (0.081)	0.219** (0.109)
max bid \times ad treatment			-1.863 (2.781)
Constant	-3.131*** (0.041)	-3.270*** (0.079)	-3.310*** (0.099)
Observations	20,918	20,918	20,918

Standard errors in parentheses
 *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table A4.3: Logistic Regression Predicting Purchase Probability Using the Median Bid per Consumer

VARIABLES	(1) purchase Logit	(2) purchase Logit	(3) purchase Logit
median bid	13.896*** (1.338)	13.926*** (1.339)	12.352*** (3.188)
ad treatment		0.171** (0.081)	0.143 (0.096)
median bid × ad treatment			1.926 (3.514)
Constant	-3.049*** (0.037)	-3.188*** (0.076)	-3.165*** (0.087)
Observations	20,918	20,918	20,918

Standard errors in parentheses
 *** p<0.01, ** p<0.05, * p<0.1

Table A4.4: Logistic Regression Predicting Purchase Probability Using the Cumulative Bid per Consumer

VARIABLES	(1) purchase Logit	(2) purchase Logit	(3) purchase Logit
cumulative bid	0.202*** (0.060)	0.203*** (0.060)	0.096 (0.153)
ad treatment		0.166** (0.081)	0.136 (0.089)
cumulative bid × ad treatment			0.129 (0.166)
Constant	-2.924*** (0.034)	-3.059*** (0.075)	-3.034*** (0.081)
Observations	20,918	20,918	20,918

Standard errors in parentheses
 *** p<0.01, ** p<0.05, * p<0.1

A4.4 Prediction of Purchase Probabilities

We decide to predict the purchase probabilities of consumers in our sample to both investigate whether the bid placed by the ad platform is correlated with this prediction as well as whether we are able to find heterogeneous treatment effects of the ad treatment when comparing consumers with differently predicted purchase probabilities. To do so, we randomly assign 80% of our data to train our predictive model, keeping the remaining 20% as validation set. In a first step, we select a set of variables that offers the highest area under the curve (*AUC*) when applied to our validation data (see Table A4.5). We use logit model in the variable selection stage to select the variables that maximize the *AUC* in our validation set.

Table A4.5: Descriptive Statistics for Variables in Predictive Model

VARIABLES	(1) N	(2) mean	(3) sd	(4) min	(5) max
impressions after shopping basket visit	20,918	0.908	5.193	0	86
impressions after category page visit	20,918	4.034	10.09	0	81
impressions after product page visit	20,918	15.15	15.71	0	104
activity duration	20,918	6.971	15.50	0	656.9
number of visited categories	20,918	1.478	0.927	1	20
check-out page visits	20,918	0.0105	0.156	0	7
product page visits	20,918	2.899	5.092	0	168
homepage visits	20,918	1.027	2.536	0	47
shopping cart visits	20,918	0.101	0.838	0	32
unweighted revenue potential	20,918	1,620	3,359	0	165,489
country 1	20,918	0.694	0.461	0	1
country 2	20,918	0.295	0.456	0	1

In a next step, we compare the logit model's performance with the performance of a probit model (see Figure A4.2). We find that, when looking at the *AUC* values, the probit model is outperforming the logit model. The performance difference between the logit and probit model is significant ($Prob > chi2 = .0186$).

We predict the purchase probabilities for all consumers in the sample using the probit model. We use these predicted purchase probabilities to first see whether they can predict the average bid per consumer that is placed by the ad platform (see Table A4.6, column 1). We find that our purchase probability prediction variable is significant when regressing it upon the average bid placed by the ad platform ($\beta_{\text{predicted purchase probability}} = .129, p < .001$). This once more points towards the




Figure A4.2: Receiver Operating Characteristic (ROC) Curve

ad platform making use of consumer characteristics to determine the height of the respective bid for an ad impression.

Then we investigate whether we are able to find a heterogeneous treatment effect of our ads when comparing consumers with different predicted purchase probabilities (see Table A4.6, column 2). Consistent with our main analysis (see Table 4.3), we find that our ad treatment does significantly increase consumers' likelihood to purchase ($\beta_{ad\ treatment} = .287$, $p = 0.03$). At the same time, our predicted purchase probabilities (by construction) significantly predict whether a consumer will conduct a purchase ($\beta_{predicted\ purchase\ probability} = 9.570$, $p < 0.001$). Most importantly, we cannot find evidence for the presence of a heterogeneous treatment effect when looking at the interaction of the ad treatment with our variable for predicted purchase probabilities instead of the average bid per consumer that is generated by the ad platform.

Table A4.6: Impact of Predicted Purchase Probability

VARIABLES	(1) bid OLS	(2) purchase Logit
ad treatment		0.287** (0.132)
predicted purchase probability	0.129*** (0.003)	9.570*** (1.584)
ad treatment × predicted purchase probability		-2.214 (1.710)
Constant	0.004*** (0.000)	-3.557*** (0.122)
Observations	20,918	20,918
R-squared	0.071	

Standard errors in parentheses

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

A4.5 Alternative Cost Measures

In our main analysis (see Table 4.3), we investigate whether consumers that receive an on average higher bid from the ad platform are more receptive towards ads. Nevertheless, as the ad platform is operating a second-price auction, the firm is not actually paying the bid placed by the ad platform but the second highest bid that is placed on behalf of another firm that is competing for a respective impression. Therefore, we investigate the robustness of our model by focusing on the actual average cost incurred per impression and consumer.

We run our logit model for consumers' purchase probability using the actual average cost per impression and consumer that the firm is paying instead of the bid that is placed by the ad platform on behalf of the firm.

$$Purchase\ Probability_j^{Retargeting} = \frac{\exp(U_j^{Retargeting})}{\exp(U_j^{Retargeting}) + 1}$$

$$U_j^{Retargeting} = \beta_0 + \beta_1 cost_j + \beta_2 ad\ treatment_j + \beta_3 ad\ treatment_j \times cost_j + \epsilon_j$$

where, compared to our main analysis, we replace bid_j with $cost_j$ which gives the average cost of an impression for a consumer in Euro over the duration of the experiment. Again, in case the ad platform would optimize the bidding in the interest of the firm, we would expect a positive and significant coefficient for the interaction term of $ad\ treatment_j$ and $cost_j$. More expensive impressions for a consumer should have a higher impact on the purchase probability. Nevertheless, our results remain consistent (see Table A4.7): Generally, the firm is on average paying more to serve impressions to consumers that are more likely to convert ($\beta_{cost} = 44.645$, $p < .001$). Ads have a significant positive impact on consumers' purchase probability ($\beta_{ad\ treatment} = .166$, $p = .040$). Nevertheless, the on average more expensive impressions for consumers are not justified by a higher impact of ads on consumers' purchase probability ($\beta_{ad\ treatment \times bid} = -4.210$, $p = .738$).

Table A4.7: Logit Regressions for Impact of Actual Incurred Average Cost on Purchase Probability

VARIABLES	(1) purchase Logit	(2) purchase Logit	(3) purchase Logit
cost	44.645*** (4.990)	44.639*** (4.983)	48.035*** (11.310)
ad treatment		0.166** (0.081)	0.185* (0.099)
ad treatment \times cost			-4.210 (12.608)
Constant	-3.056*** (0.038)	-3.191*** (0.077)	-3.206*** (0.089)
Observations	20,918	20,918	20,918

Standard errors in parentheses
 *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Additionally, we run our logit model estimating the impact of the overall cost for all impressions served to a consumer over the duration of the experiment instead of the average cost (see Table A4.8). This way we take the variation in the number of impressions per consumer into account that is driving the overall costs that the firm

needs to pay per consumer. We find that ad impressions are not more effective for consumers for which the firm is investing more in advertising.

These analyses show that the firm is paying more for ad impressions that do not yield a higher return, which would be represented by a more significant increase in purchase probabilities of consumers for whose ad impressions the firm is paying more.

Table A4.8: Logit Regressions for Impact of Actual Incurred Overall Cost on Purchase Probability

VARIABLES	(1) purchase Logit	(2) purchase Logit	(3) purchase Logit
overall cost	0.647** (0.268)	0.644** (0.268)	1.194* (0.638)
ad treatment		0.164** (0.081)	0.214** (0.098)
treatment \times overall cost			-0.657 (0.704)
Constant	-2.925*** (0.037)	-3.058*** (0.076)	-3.100*** (0.089)
Observations	20,918	20,918	20,918

Standard errors in parentheses
 *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

A4.6 Revenue Regression

In the main part of this study, we investigate whether the ad platform bids higher for consumers that generate higher revenue for the ad platform. We do this using only the observations of consumers that actually conduct a purchase instead of the whole sample of consumers. In this section, we explain why we decided to conduct this analysis this way.

Most of the consumers in our sample do not conduct a purchase and therefore contribute no revenue ($n_{purchase} = 1,113$; $n_{no\ purchase} = 19,805$). We need to be cautious about how to investigate the moderating effect of bid height on ad effectiveness (revenue generation) for two reasons.

First, such an analysis mixes up two predictions the ad platform needs to make when deciding whom to target: (1) Who will purchase? (2) How much will they spend? These two questions are heavily intertwined as predicting who will purchase correctly will lead to the ad platform automatically targeting consumers that generate higher revenue compared to consumers that do not conduct a purchase.

Second, in a case with heavily skewed revenue generated by the consumers this is especially problematic. The heavily skewed distribution of revenue and impact of outliers might cause the results to rather represent a case of chance than actual performance of the ad platform. The distribution of the revenue per consumer in our sample is heavily positively skewed (see Figure A4.3). We have some very extreme outliers ($revenue_{max} = 3,789.8$). With a dependent variable that skewed these outliers strongly affect the estimation results in our regression while they might occur by chance. According to Tukey (1977)'s definition of outliers, all our observations with $revenue > 0$ are considered outliers (see Figure A4.4).




Figure A4.3: Histogram Revenue





Figure A4.4: Box Plot Revenue

We therefore investigate the results from a more descriptive perspective. We plot the average uplift in revenue per consumer for every bid decile in our sample (see Figure A4.5). This graph depicts the difference in the average revenue generated between treated and non-treated consumers for each bid decile, which is basically representing the effect of the interaction between $ad\ treatment_j$ and bid_j on $revenue_j$. We see here that there seems to be no clear increasing trend with an increase in the bid placed by the ad platform. In contrast, this graph seems to rather suggest a decreasing trend in advertising's contribution to revenue with an increase in the average bid (especially when not considering bid decile 8 and 10). Further, we see

that the impact of bid decile 10 seems to be particularly strong and positive. Such an extreme value is likely to drive the results in the regression.

Figure A4.5: Average Revenue Uplift per Consumer and Bid Decile



To remedy the effect of the heavily skewed distribution of revenue, we run our analysis using the log transformed value of revenue¹⁰. Figure A4.6 displays the histogram of the log-transformed revenue variable, in which the distribution of revenue values moves closer together. This transformation decreases the impact of extreme revenue values in the regression.

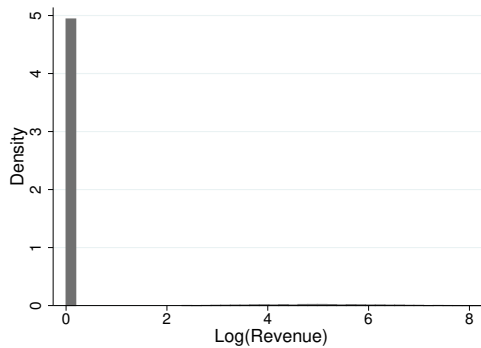


Figure A4.6: Histogram Log(Revenue)

¹⁰We actually make use of the $\log(\text{revenue}_j + 1)$ as we would otherwise lose the observations with zero revenue.

Table A4.9 gives the results with the log-transformed revenue variable as dependent variable. Consistent with our main analysis, we find that there is no significant interaction effect between treatment and bid. Once more, this suggests that the ad platform does not target consumers that generate a higher revenue caused by ads.

Table A4.9: Log(Revenue) OLS Regression

VARIABLES	(1)	(2)	(3)
	log(revenue) OLS	log(revenue) OLS	log(revenue) OLS
ad treatment	0.039** (0.019)	0.041** (0.019)	0.024 (0.022)
bid		5.592*** (0.689)	4.431*** (1.269)
ad treatment \times bid			1.479 (1.503)
Constant	0.233*** (0.017)	0.170*** (0.018)	0.183*** (0.019)
Observations	20,918	20,918	20,918
R-squared	0.0002	0.0063	0.0064

Robust standard errors in parentheses

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

A4.7 Controlling for the Effect of Competition for Ad Impressions

A limitation of our work is that we compare our ad treatment to the orthogonal charity ad treatment. We use this method to identify the causal impact of ads as we would otherwise not be able to recognize the consumers in our control group. Previously, researchers have argued that this method to identify the causal impact of ads suffers from the limitation that the true counterfactual of an ad treatment would potentially include ads from competitors (Johnson et al., 2017a). This means that ads might have a defensive effect originating in the preemption of the display of competitors' ads. In our experimental setup we are not able to capture this defensive effect of ads as our

counterfactual consists of charity ads that are orthogonal to our treatment excluding competitors' ads¹¹.

To address this issue we generate a measure for how high the competition for an ad impression is. We do this by subtracting the actual average cost for impressions for a consumer j from the average bid placed for the consumer. The underlying idea is that if the difference between the bid and cost is smaller, there is higher competition for ad impressions for the respective consumer. Figure A4.7 shows the absolute difference between bid and cost per bid decile. A higher difference, indicated by a higher bar, would represent lower competition for a consumer. This means that when looking at the absolute difference between cost and bid we find that for consumers that receive on average higher bids there is less competition as other firms that target these consumers bid lower for these ad impressions.




Figure A4.7: Difference Between Average Bid and Cost per Bid Decile

In a next step, we include this difference between the average bid and cost into our main analysis. We call this variable *absolute competition_j* and rerun our logit model estimating the purchase probability of consumer j . Table A4.10 give the results of this analysis. We find that a higher value for *absolute competition_j*, which indicates less competition, is correlated with lower purchase probabilities of consumers. This makes sense as these consumers receive lower bids from firms competing for the purchase of a

¹¹Consumers might still see ads of competitors but we are not able to directly control for this. Furthermore, this effect is in our case symmetric between control and treatment group while with a counterfactual including competitors' ads it would be stronger for the control group.

respective ad impression as they likely have not indicated their purchase interest with these firms. Such behavior correlates with a lower overall purchase probability. Most importantly, when looking at the focus coefficient of this study, $bid_j \times ad\ treatment_j$, that indicates whether consumers that receive on average higher bids for their ad impressions are more receptive to ads, we still find no significant coefficient. This means that these consumers are not more receptive even when controlling for competition for ad impressions.

Table A4.10: Logistic Regression Predicting Purchase Probability Controlling for Competition for Bids

VARIABLES	(1) purchase Logit	(2) purchase Logit	(3) purchase Logit	(4) purchase Logit
absolute competition	14.756*** (1.549)	-21.720*** (6.093)	-21.611*** (6.096)	-21.606*** (6.101)
bid		33.648*** (5.358)	33.588*** (5.359)	33.524*** (6.178)
ad treatment			0.171** (0.081)	0.170* (0.099)
bid \times ad treatment				0.074 (3.560)
Constant	-3.011*** (0.035)	-3.120*** (0.040)	-3.259*** (0.078)	-3.258*** (0.091)
Observations	20,918	20,918	20,918	20,918

Standard errors in parentheses
 *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Nevertheless, it is not directly clear whether the operationalization of competition for ad impressions in an absolute manner is sensible. We decide to rerun the analysis using a relative value for competition taking the height of the average bid into consideration. In detail, we operationalize the variable measuring competition for ad impressions as $relative\ competition_j = \frac{bid - cost}{bid}$. This way we correct the difference between bid and cost for the height of the average bid placed for a consumer. Figure A4.8 shows the relative difference between cost and bid per bid decile. Again, a higher value for $relative\ competition_j$ indicates less competition. Once more, we find that competition for impressions seems to decrease for consumers that receive on average higher bids.

We move on to re-estimate the logistic regression estimating purchase probabilities by including our variable $relative\ competition_j$. Consistent with our main results, we find that there is no significant coefficient for the interaction between $bid_j \times ad\ treatment_j$,

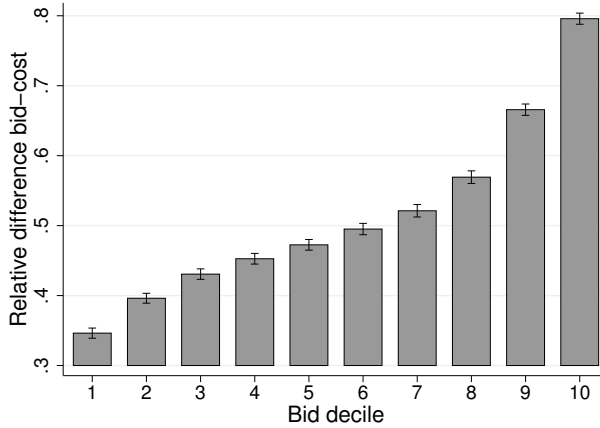


Figure A4.8: Relative Difference Between Average Bid and Cost per Bid Decile

rendering our results robust to the concern of the defensive effect of ads that we cannot directly consider in our experimental design.

Table A4.11: Logistic Regression Predicting Purchase Probability Controlling for Relative Competition for Bids

VARIABLES	(1) purchase Logit	(2) purchase Logit	(3) purchase Logit	(4) purchase Logit
relative_competition	0.570*** (0.138)	-0.130 (0.156)	-0.132 (0.156)	-0.132 (0.156)
bid		15.605*** (1.645)	15.660*** (1.646)	15.248*** (3.283)
ad treatment			0.174** (0.081)	0.166* (0.098)
bid × ad treatment				0.509 (3.501)
Constant	-3.180*** (0.080)	-3.013*** (0.079)	-3.153*** (0.104)	-3.147*** (0.113)
Observations	20,918	20,918	20,918	20,918

Standard errors in parentheses

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

Conclusions

In this dissertation we investigate the implications of advertising personalization for firms, consumers, and ad platforms. We focus on the implications for, what we consider, the three main stakeholders of personalized advertising on the demand side of advertising.

Our research is guided by the research question:

How do firms' advertising personalization strategies affect consumer responses and how can these consumer responses be assessed using ad platforms' reporting systems?

We conduct three studies to answer this research question. When assessing the implications of firms' advertising personalization strategies, we find that more specific ad personalization leads to more positive consumer responses. Nevertheless, such a positive effect of more specific ad personalization can be harmed when combining ad personalization with other advertising targeting techniques, e.g. social targeting. In social targeting firms make use of consumers' social connections to identify them as members of a relevant target group and make these social connections explicit in the ad text. We find that consumers that are socially targeted with personalized ads respond less positive. This boundary condition of advertising personalization should be incorporated in firms' personalization strategies.

Despite the notion that higher levels of personalization come at the cost of increased consumer privacy concerns (Awad and Krishnan, 2006; Sutanto et al., 2013), we find that more specific ad personalization leads to more positive consumer responses to

personalized ads. We further investigate the role of consumer privacy concerns and the personalization privacy paradox for advertising by looking into their effect on consumers' visual attention towards ads. Previous research suggests that firms might use consumers' inclination to focus on negative information stimuli, i.e. the negativity bias (Kanouse, 1984; Kätsyri et al., 2016), to attract consumers' attention towards ads that cause privacy concerns. Our research shows that consumers tend to decrease their overall attention towards ads if they experience higher privacy concerns, leading to less positive consumer responses to personalized ads. Firms' ad personalization strategy should take into account that consumers' attention cannot be attracted with intrusive ads that trigger consumer privacy concerns.

To enable the personalization of advertising, firms need to have access to information on individual consumers allowing for the inference of consumer preferences. The requirement to have access to this data creates a selection in terms of which consumers can be targeted with personalized ads. In case no information on an individual consumer is available, this consumer cannot be targeted with a personalized ad. At the same time, this selection carries the risk that firms target consumers with ads that have already indicated an interest in the firm's products, i.e. have a higher propensity to purchase, leading to only a marginal effect of ads on consumers' likelihood to purchase from the firm. We find that ad platforms are incentivized to address consumers that have an inherently high probability to purchase from the firm, independent of the effect of ads. Our empirical analysis then shows that consumers that are being targeted by ad platforms, i.e. high purchase probability consumers, are not more receptive to ads. This means that the purchase probability of these consumers is not increasing significantly stronger than for consumers that are targeted to a lower degree by the ad platform. Economically speaking, the way the ad allocation is governed currently leads to firms paying more for ads that do not deliver a higher value. We find empirical evidence for an incentive misalignment between firm and ad platform that is harmful for the firm. Firms need to take these underlying economic relationships with ad platforms into account when assessing consumer responses to personalized ads to correctly assess ad performance.

5.1 Synthesis of Findings

This dissertation was conducted with the aim to provide insights on the implications of advertising personalization for firms, consumers, and ad platforms. We conducted this

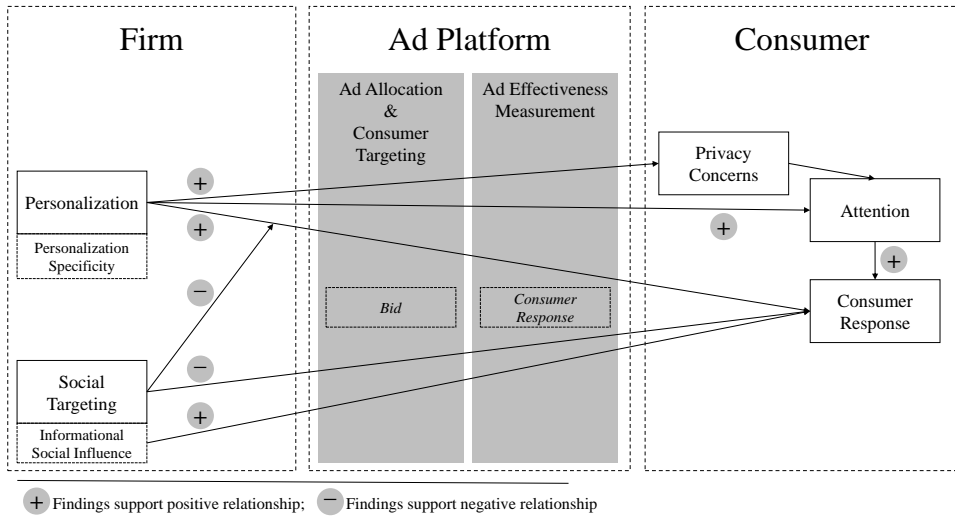


Figure 5.1: Synthesis of Findings

research with the intention to provide a holistic view on the interdependence of these stakeholders within the context of advertising personalization. Additionally, we think that this research can inform policy makers that need to make regulatory decisions regarding the use of consumer data for ad personalization purposes. The synthesis of findings from our three studies investigating different aspect of the relationships of firms, consumers, and ad platforms, allows us to arrive at more general insights. We depict the integration of the research conducted for this dissertation in Figure 5.1.

Our research makes clear that firms need to make strategic choices when defining their personalization strategies with the aim to positively influence consumers' responses to their personalized ads. Firms can benefit from increasing personalization capabilities, provided by ad platforms, to reach consumers. Furthermore, firms can benefit by making use of opportunities to measure the effectiveness of their personalization strategies and inform their strategic choices.

Our research provides support for personalization leading to more positive consumer responses to ads. We find that next to personalized ads compared to non-personalized ads, higher levels of personalization specificity lead to more positive consumer responses. Higher levels of ad personalization allow firms to address consumers with ads that are closer to consumers' preferences leading to more positive consumer responses. These findings resolve contradictory findings with respect to the adequate level of

ad personalization (Bleier and Eisenbeiss, 2015a; Lambrecht and Tucker, 2013) by providing evidence from carefully designed experiments.

Nevertheless, not all personalization features lead to more positive consumer responses. We find that socially targeted personalized ads lead to less positive consumer responses. This effect is enhanced with higher levels of personalization specificity. Social targeting negatively moderates the relationship between personalization specificity and consumer response. We challenge previous research that pointed towards positive implications of social targeting on consumers' responses to ads (Bakshy et al., 2012; Tucker, 2016).

Interestingly, we find in our lab experiment that informational social influence does positively affect consumer response. Although these findings from the field and the lab seem to contradict each other on a first glance there several reasons for the differences in how consumers respond to social targeting in the field experiment (chapter 2) and information social influence in the lab experiment (chapter 3).

First, social targeting (chapter 2) and informational social influence (chapter 3) are operationalized differently. Social targeting does not represent an experimental treatment and instead can be seen as a consumer characteristic that is made explicit in the ad text. Social targeting (chapter 2) is made explicit to consumers by displaying their friends' names in the ad (in case these friends are connected to the advertiser). Informational social influence (chapter 3) is an experimental treatment and not represented by friends' names but by their profile pictures that are displayed to experiment participants on top of the ad. One reason why informational social influence, operationalized as social endorsements, leads to more positive consumer responses while social targeting decreases consumer responses could be that consumers are more intrigued by the profile pictures of friends. This form of social endorsement is rather new, compared to including friends' names, potentially triggering a more positive perception.

Next to that, it remains challenging to create a realistic environment in a lab setting (chapter 3). A lack of external validity might potentially influence how consumers perceive the social endorsements. As the lab environment seems more artificial, safe, and the presented scenario not fully realistic, consumers might respond differently to the inclusion of social elements in ads. We tried our best to create a realistic lab environment but simply face the common limitation of external validity of lab experiments.

Lastly, there are vast differences between the type of products advertised in the field experiment (chapter 2) and the lab experiment (chapter 3). In the field experiment, our

partner firm advertises consumer electronics to consumers. We argue that the inclusion of friends names in the ad text conflicts with consumers' perception that an ad is uniquely personalized for them. In contrast, in the lab, we advertise privacy sensitive products, a gambling application as well as a STD treatment. The difference in product categories might influence how consumers perceive the social component in the ad. Especially, as consumers are unlikely to signal their identity with privacy sensitive products, offsetting the conflict between personalization and social endorsements.

Firms compete online for consumers' attention with other advertisers as well as organic content such as news or social media posts. Personalization offers the means to increase ad relevance and attract consumers' attention. We find evidence for the positive effect of personalization on consumers' attention towards an ad. Furthermore, we find that consumers' attention is partially mediating the relationship between ad personalization and consumer response. The more attention consumers dedicate towards an ad the more likely they are to respond positively to the ad. Therefore, our research supports the positive role of consumer attention for ad effectiveness (Wedel and Pieters, 2007).

Nevertheless, consumers remain concerned about the use of their information for personalization purposes. While ad platforms track consumer responses, such as clicking on an ad or purchasing after having seen an ad, ad platforms do not provide a direct measure that allows firms to assess consumer privacy concerns. Although firms can infer which type of ad performs better or worse, attributing the superior performance of an ad to a mechanism, for example that ads cause lower privacy concerns, remains very challenging. Insights from lab experiments can help firms to overcome such limitations. We find that personalization does lead to an increase in consumer privacy concerns. Nevertheless, in our setting, we do not find evidence for a direct effect of consumer privacy concerns on consumer response to personalized ads. Therefore, we do not draw a connection between privacy concerns and consumer response in our model (see Figure 5.1). We assume that this might be the case as participants in our lab experiment did not perceive privacy to be a big concern in our controlled lab setting.

Further, we find no evidence for a direct effect of privacy concerns on attention. Interestingly, we find evidence for a negative serial mediation effect from personalization through privacy concerns and attention to consumer response. Therefore we draw an arrow from privacy concerns to attention to depict this path. As we do not find a direct effect from privacy concerns to attention, we draw the arrow without mathematical sign (see Figure 5.1). This serial mediation effect points towards privacy concerns

playing a role as mediator in our context. Next to that, this finding provides evidence against the potential role of a negativity bias that attracts consumers' attention towards information stimuli that trigger negative emotions (Fiske, 1980; Kanouse, 1984; Kätsyri et al., 2016; Smith et al., 2003). Firms cannot make use of intrusive ads, that attract consumers' attention, to improve the performance of their ads.

Figure 5.1 also displays the powerful role of ad platforms that mediate the relationship between consumers and firms and act as gatekeeper regarding information on ad allocation and ad effectiveness. Ad platforms that automate the ad allocation for firms and allow firms to adjust ad content for the individual consumer provide the necessary technical infrastructure that builds the foundation for the personalization process. To serve digital advertising to consumers, firms typically outsource the ad allocation process to ad platforms. While ad platforms can, due to their scale, provide an efficient technical solution for the ad allocation process, firms need to rely on the ad platform both conducting the ad allocation in their interest and reporting ad effectiveness accurately. We are able to retrieve insights on the ad platform's targeting by observing the bids the ad platform places for ad impressions for different consumers. Next to that, we obtain our consumer response variable for chapter 2 and 4 from the reporting systems of ad platforms. The ad platforms allow us to operationalize the variable consumer response as clicks and purchases, but also as consumer responses such as comments, likes and shares of an ad.

Our research shows that, under the currently common contract structures, ad platforms conduct the ad allocation in a way that is not aligned with the firm's interest. Ad platforms are incentivized to target consumers with high purchase probabilities independent of the effect of ads on their purchase probabilities. We advance related research (Johnson and Lewis, 2015), by showing empirically that ad platforms adhere to these incentives and that consumers that are targeted by ad platforms are not more receptive to ads. This means that the ad allocation process is conducted sub-optimally for the firm that would like to target consumers that are more receptive to the firm's ads.

Personalization of advertising coincides with the selection of consumers that can be addressed with personalized ads. This targeting of consumers is conducted by the ad platform by making use of information on consumers to infer consumers' preferences. If there is no information available on certain consumers, firms cannot serve personalized ads to these consumers. Firms need to take this selection into account when assessing the performance of their personalization strategy and how they reward ad platforms.

When focusing solely on the ad personalization strategy without properly evaluating its performance, firms run the risk to use a sub-optimal strategy.

Ad platforms should consider which advertising features are beneficial for firms. We find that socially targeted consumers respond less positive to personalized advertising. In light of increasing consumer concerns regarding their privacy, ad platforms might benefit both consumers and firms by abandoning advertising features that make use of consumers' information and do not improve ad performance. Such a step would likely be in the interest of regulators, as becomes apparent for example in the European General Data Protection Regulation (GDPR). The GDPR aims to establish standards that define how firms have to manage and can process consumers' data (EUR-Lex, 2016). What became apparent during our research process is that a major share of consumers' personal information is being processed by ad platforms and not firms. Often times, firms do not have access to this personal information (e.g. Facebook Ad Platform) but make use of services of ad platforms that process consumer information in the background.

Generally, our research implies that a firm's personalization strategy needs to incorporate not only how consumers respond to different types of personalized ads but also the economic relationships underlying the personalization process. Only by taking a more holistic perspective on advertising personalization, taking relationships with both consumers and ad platforms into account, firms become able to both define their personalization strategy as well as assess its performance adequately.

5.2 Academic Contributions

With the present research we contribute to the literature on advertising personalization and economics of advertising. We are confident that our results promote both academic research and the practical understanding. Such understanding is required to define and implement personalization strategies as well as assess the performance of personalized ads.

In Chapter 2, we offer firms insights on how to define their personalization strategy to increase consumers' likelihood to click ads and purchase after being confronted with ads. More precisely, with the help of a large scale field experiment, we find that more specific ad personalization, recommending specific products instead of less specific product categories, leads to more positive consumer responses to ads. Further, we present evidence that when personalized ads are socially targeted, their performance

decreases. With these findings, we add to the theoretical discussion on adequate levels of ad personalization. Results from previous research offered inconsistent insights in that some work suggested lower levels of advertising personalization to outperform more personalized ads (Lambrecht and Tucker, 2013), whereas other research finds that higher levels of ad personalization increase ad performance (Bleier and Eisenbeiss, 2015a). Next to that, we are able to provide empirical evidence for the conceptual conflict of socially targeting personalized ads. As consumers expect a personalized ad message is uniquely personalized for them, including peers' names in the ad text that endorse an ad conflicts with the idea of unique personalization. Therefore, we challenge the previously established understanding of generally positive implications of social targeting for advertising (Bakshy et al., 2012; Tucker, 2016). To the best of our knowledge, we are the first to empirically investigate personalized ads in combination with social targeting.

In Chapter 3, we show that firms have an interest to increase consumers' attention towards their ads. Attention acts, as previously pointed out in marketing research (Wedel and Pieters, 2007), as positive mediator between ad personalization and ad performance. Theoretically, we find arguments for both an increase and a decrease in consumers' attention towards ads with an increase in consumer privacy concerns triggered by firms' use of consumers' information to personalize ads. The negativity bias in consumers' attention, describes that consumers tend to direct their attention towards negatively perceived stimuli (Kanouse, 1984; Kätsyri et al., 2016). This suggests that consumers dedicate more attention to ads when experiencing higher privacy concerns, as privacy concerns trigger negative feelings of vulnerability within consumers (Aguirre et al., 2015). Contradicting this notion, previous research in marketing found that attention acts as enabler for advertising performance as consumers need to dedicate attention to ads to process them (Ho and Bodoff, 2014). Our results confirm the previously established positive role of attention for advertising performance and show that the negativity bias does not play a significant role when consumers assess personalized ads. We show that a strategy in which firms try to catch consumers' attention with intrusive ads that trigger privacy concerns, does not pay off as these ads decrease consumers' overall attention towards ads. With this work we contribute to research investigating the personalization privacy paradox in an advertising context.

In Chapter 4, we find that the contracts that are currently implemented in programmatic advertising to govern the ad allocation process between firms and ad platforms are not favorable for firms. We show empirically that due to the currently specified incentives in these contracts, firms end up paying more for ad impressions that do

not deliver a higher return on ad investment. Ad platforms target consumers with ads, by bidding higher for their ad impressions, that are not more receptive to these ads. With this work we contribute to the literature on economics of advertising. We present theoretical arguments for the potential presence of an incentive misalignment in programmatic advertising. This misalignment is contingent on both the ad platform's actual behavior as well as the correlation between absolute and incremental purchase probabilities of consumers. Next, with our unique empirical setting, we are able to present evidence for the actual presence of this incentive misalignment. While previous research has suggested that current contracts implemented in programmatic advertising might not be beneficial for firms (Johnson and Lewis, 2015), we are able to present actual empirical evidence for the presence of this incentive misalignment in programmatic advertising.

5.3 Practical Relevance

In this dissertation, we focus on the implications of personalized digital advertising for firms, consumers, and ad platforms. We consider these to be the three main stakeholders on the demand side of personalized advertising (where the supply side encompasses publishers that display ads to consumers on their websites and sell this ad space to firms via ad platforms). We provide several insights of high practical relevance that can help firms to improve their ad personalization strategies and assess ad performance more adequately. We help consumers to understand the benefits and costs of personalized advertising. Furthermore, we allow ad platforms to assess limitations of targeting mechanisms and incentive structures in contracts governing their relationships with firms.

Firms have identified the potential of personalized digital advertising but simultaneously struggle with how to define and implement their personalization strategies (Adobe Systems Inc., 2014). In chapter 2, we present insights for firms on how to define their personalization strategies to increase ad effectiveness. Chapter 3, suggests to firms that more intrusive ads are not an adequate mean to attract consumers' attention to ads. Lastly, chapter 4 cautions firms about the currently implemented contracts between firms and ad platforms in programmatic advertising. Under the currently specified incentives in these contracts ad platforms optimize the ad allocation process in a way that is not in the interest of firms. Firms need to reevaluate how they

structure the economic relationship with ad platforms to allow for a more favorable ad allocation process.

Consumers can benefit from being addressed with more relevant personalized ads instead of generic ads. In chapter 2, we confirm this notion and show that consumers respond more positively to more personalized ads. At the same time, we show in chapter 3, that consumers are concerned about their privacy when addressed with personalized ads. These privacy concerns represent an economic cost that decreases the benefits of ad personalization for consumers and limits the amount of attention consumers are willing to dedicate to ads.

Ad platforms act as mediator between firms, that want to display ads, and consumers, that are the recipients of these ads. Chapter 2 offers insights into the commonly implemented advertising strategy of social networking sites. These sites operate using an advertising revenue model (Schumann, 2014) and can therefore be considered as ad platforms in the context of social media. We find that socially targeting consumers with personalized ads backfires. Therefore, we are pointing ad platforms to reconsider their default strategy to socially target consumers when ads are personalized. In chapter 4, we show that ad platforms are optimizing the ad allocation process in programmatic advertising rather in their own than the firm's interest. Ad platforms need to be cautious about the rising awareness of this behavior on the side of firms and the implications on their business model.

While we focus on the implications of personalized advertising for firms, consumers, and ad platforms in this dissertation, we also offer insights that are valuable for policy makers. Chapter 3 emphasizes the importance for policy makers to stay aware of consumers' privacy concerns originating in firms' use of consumers' personal information for personalization purposes. Where firms go too far in the use of this information, policy makers might need to intervene to protect consumers. Consumers have generally been shown to prefer environments with less advertising and are less willing to give access to their personal information for personalization of ads compared to the personalization of other services (Awad and Krishnan, 2006; Sutanto et al., 2013).

5.4 Limitations

The aim of this dissertation is to provide an integrated and more holistic view on the implications of advertising personalization for firms, consumers, and ad platforms.

Although we are able to shed light on the different relationships between these stakeholders we want to acknowledge that this dissertation does not contain a study that focuses on all three stakeholders equally at the same time. Instead, we focus on specific relationships to be able to dive deeper into the implications of personalization strategies and the assessment of consumer responses. We are confident that insights from the different studies can still be combined to derive broader insights.

For both chapter 2 and 4 we face the difficulty that we need to rely on data provided by ad platforms for our empirical analysis. It remains challenging for firms to assess to what extent ad platforms report accurate data. Economically, ad platforms have an interest to over-report the success of marketing campaigns to incentivize firms to keep investing in their marketing budgets. We are aware of the limitation regarding the data we are using in this dissertation. Nevertheless, the focus of this dissertation is less on the quality and accuracy of data from ad platforms but focuses more on firms' ad personalization strategies and the economic relationship between firms and ad platforms. We are confident that our analyses can provide valuable insights regarding these topics.

A major difficulty in field research is to not only identify the causal effect of an experimental treatment but to disentangle the underlying causal mechanisms. While field experiments offer the opportunity to run an experiment in a realistic environment and use much larger sample sizes it is difficult to gain insights on these causal mechanisms as consumers can seldomly be questioned about their underlying motivations and perceptions. The difficulty to identify causal mechanisms most severely affects chapter 2 of this dissertation. We conducted an extensive amount of additional analysis with the aim to overcome this limitation and provide better insights on underlying mechanisms.

5.5 Future Research

This dissertation lies in the intersection of marketing and information systems research. We see this intersection as very fruitful and would like to encourage researchers to continue to combine theories and insights from both fields to generate and disseminate new knowledge. Nowadays, information and communication technologies enable complex marketing operations. Therefore, the combination of information systems with marketing methodologies deserves further attention.

There is a plethora of research opportunities within the area of advertising personalization. Below we want to present some of the directions that we find promising based on the past five years of experience in this research area.

The ability to personalize ads often times heavily depends on the availability of information on consumers underlying the applied personalization mechanisms. The necessity of having access to information on consumers, leads to a 'natural' sample selection. Only consumers for whom the advertiser has information available can be addressed with personalized ads. Firstly, we call for more research that more clearly distinguishes the concept of ad personalization from ad targeting. It remains unclear where the boundary between these two terms lies and when the terms actually explain the same marketing technique. Next, further research should address the question to what extent the targeting of consumers does actually lead to higher returns on advertising investments for firms. In chapter 4, we have presented evidence that ad platforms tend to address consumers with ads that are more likely to purchase independent of the effect that ads have on these consumers' purchase probabilities. When targeting consumers for whom an ad is highly relevant, firms run the risk to address consumers that would have purchased from them independent of seeing an ad. A firm's decision which consumers to target with advertising also includes the decision which consumers to exclude from the target population. Commonly, consumers are excluded because they are not interested in the advertised product. But at the same time, firms should exclude consumers that purchase anyways, independent of a treatment with ads. We call for more research investigating an adequate level of advertising targeting balancing both the exclusion of non-relevant consumers as well as overly-relevant consumers.

In chapter 3 of this dissertation, we combine several research methodologies to investigate how consumers respond to ads. More precisely, we assess the role of consumers' attention within the personalization privacy paradox. While for this study, we were still required to conduct this research within a lab environment, recent technological developments allow researchers to potentially assess consumers' attention in the field. Latest developments in both software and hardware for smart phones enable consumers to interface with their devices via eye movements. As the accuracy of this data collection method increases, it allows researchers to collect eye movement data in the field. This renders the necessity to put experiment participants in a lab setting, which is often expensive and appears unnatural to participants, obsolete. We suggest researchers to monitor opportunities that become available with the development of this technology.

In chapter 4, we present empirical evidence for the fact that currently implemented contracts between firms and ad platforms, governing the ad allocation process, are not favorable for firms. While the identification of such an economic incentive misalignment is critical to trigger change, further research focusing on the economic implications for adjusted contract structures can help the ad industry to switch to a setting with aligned incentives. While researchers have suggested systems that allow ad platforms to accurately track return on ad investment (Johnson et al., 2017a), we have seen limited interest in the ad industry to implement such systems. Economic research investigating the welfare implications for consumers, firms, and ad platforms can help to identify a market structure in which a switch to a system with aligned incentives is feasible.

While we focus on the application of personalization methodologies for marketing purposes in this dissertation, personalization can be used to adjust different types of services to consumers' preferences. This ranges from personalization of music play lists in music streaming services to personalized information matching consumers' preferences on search engines. Often times, managers within large internet corporations decide how this personalization is deployed with the aim to maximize the value of their business. At the same time, such personalization algorithms, that determine which content users of these services see at what time, have significant social implications. When confronted with only information confirming one's opinion, consumers' chance to challenge their preconceived opinions decreases. The social welfare implications of such personalization algorithms requires further attention. We hope more research will be conducted investigating the potential negative social welfare implications of personalized services.

References

- Adobe Systems Inc. Real-Time Marketing Insights Study. 2014. URL https://offers.adobe.com/en/na/marketing/landings/_46316_real_time_marketing_insights_study.html.
- AdRoll. State of Performance Marketing - North America & Europe. Technical report, AdRoll, 2017. URL <https://www.adroll.com/en-AU/resources/guides-and-reports/state-of-performance-marketing-us>.
- E. Aguirre, D. Mahr, D. Grewal, K. de Ruyter, and M. Wetzels. Unraveling the Personalization Paradox: The Effect of Information Collection and Trust-Building Strategies on Online Advertisement Effectiveness. *Journal of Retailing*, 91(1):34–49, 2015.
- E. Aguirre, A. Roggeveen, D. Grewal, and M. Wetzels. The Personalization-Privacy Paradox: Implications for New Media. *Journal of Consumer Marketing*, 33(2):98–110, 2016.
- A. Ansari and C. Mela. E-Customization. *Journal of Marketing Research*, 40(2):131–145, 2003.
- S. Aral, L. Muchnika, and A. Sundararajan. Distinguishing Influence-Based Contagion from Homophily-Driven Diffusion in Dynamic Networks. *Proceedings of the National Academy of Sciences*, 106(51):21544–21549, 2009.
- S. Aral and D. Walker. Creating Social Contagion Through Viral Product Design: A Randomized Trial of Peer Influence in Networks. *Management Science*, 57(9):1623–1639, 2011.

- N. Arnosti, M. Beck, and P. Milgrom. Adverse Selection and Auction Design for Internet Display Advertising. *The American Economic Review*, 106(10):2852–2866, 2016.
- N. Arora, X. Dreze, A. Ghose, J. D. Hess, R. Iyengar, B. Jing, Y. Joshi, V. Kumar, N. Lurie, S. Neslin, S. Sajeesh, M. Su, N. Syam, J. Thomas, and Z. J. Zhang. Putting One-to-One Marketing to Work: Personalization, Customization, and Choice. *Marketing Letters*, 19(3-4):305–321, 2008.
- N. F. Awad and M. S. Krishnan. The Personalization Privacy Paradox: An Empirical Evaluation of Information Transparency and the Willingness to Be Profiled Online for Personalization. *MIS Quarterly*, 30(1):13–28, 2006.
- T. H. Baek and M. Morimoto. Stay Away From Me. *Journal of Advertising*, 41(1):59–76, 2012.
- E. Bakshy, D. Eckles, R. Yan, and I. Rosenm. Social Influence in Social Advertising: Evidence from Field Experiments. *Proceedings of the 13th ACM Conference on Electronic Commerce*, 1(212):146–161, 2012.
- S. R. Balseiro, J. Feldman, V. Mirrokni, and S. Muthukrishnan. Yield Optimization of Display Advertising with Ad Exchange. *Management Science*, 60(12):2886–2907, 2014.
- H. Bang and B. W. Wojdyski. Tracking Users’ Visual Attention and Responses to Personalized Advertising Based on Task Cognitive Demand. *Computers in Human Behavior*, 55:867–876, 2016.
- G. Bansal, F. M. Zahedi, and D. Gefen. The Impact of Personal Dispositions on Information Sensitivity, Privacy Concern and Trust in Disclosing Health Information Online. *Decision Support Systems*, 49(2):138–150, 2010.
- F. Bélanger and R. E. Crossler. Privacy in the Digital Age: A Review of Information Privacy Research in Information Systems. *MIS Quarterly*, 35(4):1017–1041, 2011.
- J. Berger and C. Heath. Who Drives Divergence? Identity Signaling, Outgroup Dissimilarity, and the Abandonment of Cultural Tastes. *Journal of Personality and Social Psychology*, 95(3):593–607, 2008.
- A. Bleier and M. Eisenbeiss. Personalized Online Advertising Effectiveness: The Interplay of What, When, and Where. *Marketing Science*, 34(5):669–688, 2015a.
- A. Bleier and M. Eisenbeiss. The Importance of Trust for Personalized Online Advertising. *Journal of Retailing*, 91(3):390–409, 2015b.
- E. Bonabeau. The Perils of the Imitation Age. *Harvard Business Review*, 82(6):45–54, 2004.

- M. B. Brewer. The Social Self: On Being the Same and Different at the Same Time. *Personality and Social Psychology Bulletin*, 17(5):475–482, 1991.
- R. Burnkrant and A. Cousineau. Informational and Normative Social Influence in Buying Behavior. *Journal of Consumer Research*, 2(3):206–215, 1975.
- D. E. Campbell and R. T. Wright. Shut-Up I Don't Care: Understanding the Role of Relevance and Interactivity on Customer Attitudes Toward Repetitive Online Advertising. *Journal of Electronic Commerce Research*, 9(1):62–76, 2008.
- C. Chan, J. Berger, and L. Van Boven. Identifiable but Not Identical: Combining Social Identity and Uniqueness Motives in Choice. *Journal of Consumer Research*, 39(3):561–573, 2012.
- R. Chellappa and R. Sin. Personalization Versus Privacy: An Empirical Examination of the Online Consumer's Dilemma. *Information Technology and Management*, 6(2):181–202, 2005.
- E. H. Clarke. Multipart Pricing of Public Goods. *Public Choice*, 11(1):17–33, 1971.
- M. T. Clements. Direct and Indirect Network Effects: Are They Equivalent? *International Journal of Industrial Organization*, 22(5):633–645, 2004.
- B. Dalessandro, C. Perlich, O. Stitelman, and F. Provost. Causally Motivated Attribution for Online Advertising. *Proceedings of the Sixth International Workshop on Data Mining for Online Advertising and Internet Economy*. ACM., 2012.
- R. Desimone and J. Duncan. Neural Mechanisms of Selective Visual Attention. *Annual Review of Neuroscience*, 18(1):193–222, 1995.
- M. Deutsch and H. B. Gerard. A Study of Normative and Informational Social Influences Upon Individual Judgement. *Journal of Abnormal Psychology*, 51(3):629–636, 1955.
- T. Dinev and P. Hart. Internet Privacy Concerns and Their Antecedents - Measurement Validity and a Regression Model. *Behaviour & Information Technology*, 23(6):413–422, 2004.
- X. Drèze and F.-X. Husherr. Internet Advertising: Is Anybody Watching? *Journal of Interactive Marketing*, 17(4):8–23, 2003.
- B. B. Edelman, M. Ostrovsky, and M. Schwarz. Internet Advertising and the Generalized Second-Price Auction: Selling Billions of Dollars Worth of Keywords. *The American Economic Review*, 97(1):242–259, 2007.
- EMarketer. Personalization Drives Engagement, Conversions - eMarketer. 2015. URL <http://www.emarketer.com/Article/Personalization-Drives-Engagement-Conversions/1012914>.

- eMarketer. eMarketer Updates Worldwide Internet and Mobile User Figures - eMarketer. 2017a. URL <https://www.emarketer.com/content/emarketer-updates-worldwide-internet-and-mobile-user-figures>.
- eMarketer. US Time Spent with Media: eMarketer's Updated Estimates for 2017 - eMarketer. 2017b. URL <https://www.emarketer.com/Report/US-Time-Spent-with-Media-eMarketers-Updated-Estimates-2017/2002142>.
- EUR-Lex. EUR-Lex - 32016R0679 - EN - EUR-Lex. 2016. URL https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=uriserv:OJ.L_.2016.119.01.0001.01.ENG&toc=OJ:L:2016:119:TOC.
- L. T. Fisher, C. Boyle, A. Rotondo, and B. Yeager. The New Display Ad Tech Stack. Technical Report May, eMarketer, 2016.
- S. T. Fiske. Attention and Weight in Person Perception: The Impact of Negative and Extreme Behavior. *Journal of Personality and Social Psychology*, 38(6):889–906, 1980.
- C. Forman. The Corporate Digital Divide: Determinants of Internet Adoption. *Management Science*, 51(4):641–654, 2005.
- T. W. Frick and T. Li. Personalization in Social Retargeting - A Field Experiment. In *Proceedings International Conference on Information Systems*, pages 1–20. 2016.
- E. Gal-Or, R. Gal-Or, and N. Penmetsa. The Role of User Privacy Concerns in Shaping Competition Among Platforms. *Information Systems Research*, In Advance:1–25, 2018.
- M. E. Goldberg and Jon Hartwick. The Effects of Advertiser Reputation and Extremity of Advertising Claim on Advertising Effectiveness. *The Journal of Consumer Research*, 17(2):172–179, 1990.
- B. Gordon, F. Zettelmeyer, N. Bhargava, and D. Chapsky. A Comparison of Approaches to Advertising Measurement: Evidence from Big Field Experiments at Facebook. *Working Paper*, pages 1–47, 2018.
- K. Graham. Programmatic Advertising: Key Terms, Examples and Trends Explained. 2017. URL <https://www.monetizemore.com/blog/programmatic-advertising-explained/>.
- A. S. Greenberg. The Role of Visual Attention in Internet Advertising: Eleven Questions and a Score of Answers. *Journal of Advertising Research*, 52(4):400–404, 2012.
- T. Groves. Incentives in Teams. *Econometrica*, 41(4):617–631, 1973.

- B. Gu, J. Park, and P. Konana. The Impact of External Word-of-Mouth Sources on Retailer Sales of High-Involvement Products. *Information Systems Research*, 23(1):182–196, 2012.
- A. F. Hayes. *Introduction to Mediation, Moderation, and Conditional Process Analysis: A Regression-Based Approach*. Guilford Press, 2013.
- G. Hervet and K. Gue. Is Banner Blindness Genuine? Eye Tracking Internet Text Advertising. *Applied Cognitive Psychology*, 25(5):708–716, 2011.
- S. Y. Ho and D. Bodoff. The Effects of Web Personalization on User Attitude and Behavior: An Integration of the Elaboration Likelihood Model and Consumer Search Theory. *MIS Quarterly*, 38(2):497–520, 2014.
- S. Hoeffler and D. Ariely. Constructing Stable Preferences: A Look Into Dimensions of Experience and Their Impact on Preference Stability. *Journal of Consumer Psychology*, 8(2):113–139, 1999.
- D. L. Hoffman and T. P. Novak. Marketing in Hypermedia Computer-Mediated Environments: Conceptual Foundations. *Journal of Marketing*, 60(3):50–68, 1996.
- M. Irvine. Facebook Ad Benchmarks for YOUR Industry [New Data] — WordStream. 2018. URL <https://www.wordstream.com/blog/ws/2017/02/28/facebook-advertising-benchmarks>.
- G. A. Johnson and R. A. Lewis. Cost Per Incremental Action: Efficient Pricing of Advertising. *Working Paper*, pages 1–36, 2015.
- G. A. Johnson, R. A. Lewis, and E. I. Nubbemeyer. Ghost Ads: Improving the Economics of Measuring Ad Effectiveness. *Journal of Marketing Research*, 54(6):867–884, 2017a.
- G. A. Johnson, R. A. Lewis, and D. H. Reiley. When Less Is More: Data and Power in Advertising Experiments. *Marketing Science*, 36(1):43–53, 2017b.
- S. Kalyanaraman and S. Sundar. The Psychological Appeal of Personalized Content in Web Portals: Does Customization Affect Attitudes and Behavior? *Journal of Communication*, 56(1):110–132, 2006.
- D. E. Kanouse. Explaining Negativity Biases in Evaluation and Choice Behavior: Theory and Research. *Advances in Consumer Research*, 11:703–708, 1984.
- J. Kätsyri, T. Kinnunen, K. Kusumoto, P. Oittinen, and N. Ravaja. Negativity Bias in Media Multitasking: The Effects of Negative Social Media Messages on Attention to Television News Broadcasts. *Plos One*, 11(5):1–21, 2016.
- M. L. Katz and C. Shapiro. Network Externalities, Competition, and Compatibility. *The American Economic Review*, 75(3):424–440, 1985.

- M. L. Katz and C. Shapiro. Systems Competition and Network Effects. *Journal of Economic Perspectives*, 8(2):93–115, 1994.
- F. D. Keyzer, N. Dens, and P. D. Pelsmacker. Is This for Me? How Consumers Respond to Personalized Advertising on Social Networking Sites. *Journal of Interactive Advertising*, 15(2):124–134, 2015.
- Y. J. Kim and J. Han. Why Smartphone Advertising Attracts Customers: A Model of Web Advertising, Flow, and Personalization. *Computers in Human Behavior*, 33:256–269, 2014.
- S. Komiak and I. Benbasat. The Effects of Personalization and Familiarity on Trust and Adoption of Recommendation Agents. *MIS Quarterly*, 30(4):941–960, 2006.
- M. Köster, M. Rüth, K.-C. Hamborg, and K. Kaspar. Effects of Personalized Banner Ads on Visual Attention and Recognition Memory. *Applied Cognitive Psychology*, 29(2):181–192, 2015.
- K. Kwahk and X. Ge. The Effects of Social Media on E-Commerce: A Perspective of Social Impact Theory. *45th Hawaii International Conference on System Sciences*, pages 1814–1823, 2012.
- D. LaBerge. *Attentional Processing: The Brain's Art of Mindfulness*. Harvard University Press, 1995.
- A. Lambrecht and C. Tucker. When Does Retargeting Work? Information Specificity in Online Advertising. *Journal of Marketing Research*, 50(5):561–576, 2013.
- D.-J. Lee, J.-H. Ahn, and Y. Bang. Managing Consumer Privacy Concerns in Personalization: A Strategic Analysis of Privacy Protection. *MIS Quarterly*, 35(2):423–444, 2011.
- J. Lee and J.-H. Ahn. Attention to Banner Ads and Their Effectiveness: An Eye-Tracking Approach. *International Journal of Electronic Commerce*, 17(1):119–137, 2012.
- J.-M. Lee and J.-Y. Rha. Personalization-Privacy Paradox and Consumer Conflict with the Use of Location-Based Mobile Commerce. *Computers in Human Behavior*, 63:453–462, 2016.
- W. Leven. *Blickverhalten von Konsumenten: Grundlagen, Messung und Anwendung in der Werbeforschung*. Physica-Verlag, Heidelberg, 1991.
- R. Lewis, J. Rao, and D. Reiley. Here, There, and Everywhere: Correlated Online Behaviors Can Lead to Overestimates of the Effects of Advertising. *Proceedings of the 20th International Conference on World Wide Web. ACM*, pages 157–166, 2011.

- R. A. Lewis and J. M. Rao. The Unfavorable Economics of Measuring the Returns to Advertising. *The Quarterly Journal of Economics*, 130(4):1941–1973, 2014.
- H. Li and P. K. Kannan. Attributing Conversions in a Multichannel Online Marketing Environment: An Empirical Model and a Field Experiment. *Journal of Marketing Research*, 51(1):40–56, 2014.
- S. J. Liebowitz and S. E. Margolis. Network Externality: An Uncommon Tragedy. *Journal of Economic Perspectives*, 8(2):133–150, 1994.
- C. Liu. Worldwide Ad Spending. Technical Report October, eMarketer, 2016. URL http://www.strathcom.com/wp-content/uploads/2016/11/eMarketer_Worldwide_Ad_Spending-eMarketers_Updated_Estimates_and_Forecast_for_20152020.pdf.
- Y. Liu and J. Sutanto. Buyers’ Purchasing Time and Herd Behavior on Deal-of-the-Day Group-Buying Websites. *Electronic Markets*, 22(2):83–93, 2012.
- X. Luo, M. Andrews, Z. Fang, and C. W. Phang. Mobile Targeting. *Management Science*, 60(7):1738–1756, 2013.
- M. Malheiros, C. Jennett, S. Patel, S. Brostoff, and M. Sasse. Too Close for Comfort: A Study of the Effectiveness and Acceptability of Rich-Media Personalized Advertising. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems. ACM, 2012*, 2012.
- A. McCarthy. Most Marketers Worldwide Have Personalization Strategies in Place - eMarketer. 2017. URL <https://www.emarketer.com/Article/Most-Marketers-Worldwide-Have-Personalization-Strategies-Place/1015299>.
- C. McNair. Worldwide Ad Spending: eMarketer’s Updated Estimates and Forecast for 2016-2021 - eMarketer. 2017. URL <https://www.emarketer.com/Report/Worldwide-Ad-Spending-eMarketers-Updated-Estimates-Forecast-20162021/2002145>.
- A. McStay. *Digital Advertising*. Palgrave Macmillan, 2010.
- S. Meredith. Facebook-Cambridge Analytica: A Timeline of the Data Hijacking Scandal. 2018. URL <https://www.cnn.com/2018/04/10/facebook-cambridge-analytica-a-timeline-of-the-data-hijacking-scandal.html>.
- W. Moe. Buying, Searching, or Browsing: Differentiating Between Online Shoppers Using In-Store Navigational Clickstream. *Journal of Consumer Psychology*, 13(1):29–39, 2003.

- L. Muchnik, S. Aral, and S. J. Taylor. Social Influence Bias: A Randomized Experiment. *Science*, 647(2013):647–651, 2014.
- M. Palmeira and G. Spassova. The Specificity Heuristic: Consumer Evaluation of Expert Recommendation. *Advances in Consumer Research*, 40:653–654, 2012.
- C. W. Park and V. P. Lessig. Students and Housewives: Differences in Susceptibility to Reference Group Influence. *Journal of Consumer Research*, 4(2):102–110, 1977.
- R. Pieters, E. Rosbergen, and M. Wedel. Visual Attention to Repeated Print Advertising: A Test of Scanpath Theory. *Journal of Marketing Research*, 36(4):424–438, 1999.
- R. Pieters and M. Wedel. Attention Capture and Transfer in Advertising: Brand, Pictorial, and Text-Size Effects. *Journal of Marketing*, 68(2):36–50, 2004.
- R. Pieters and M. Wedel. Goal Control of Attention to Advertising: The Yarnbus Implication. *Journal of Consumer Research*, 34(2):224–233, 2007.
- R. Pieters and M. Wedel. Ad Gist: Ad Communication in a Single Eye Fixation. *Marketing Science*, 31(1):59–73, 2012.
- R. Pieters, M. Wedel, and R. Batra. The Stopping Power of Advertising: Measures and Effects of Visual Complexity. *Journal of Marketing*, 74(5):48–60, 2010.
- F. Provost, B. Dalessandro, R. Hook, X. Zhang, and A. Murray. Audience Selection for On-line Brand Advertising: Privacy-friendly Social Network Targeting. *Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, pages 707–716, 2009.
- E. Rosbergen, R. Pieters, and M. Wedel. Visual Attention to Advertising: A Segment-Level Analysis. *Journal of Consumer Research*, 24(3):305–314, 1997.
- N. S. Sahni, S. C. Wheeler, and P. K. Chintanguta. Personalization in Email Marketing: The Role of Noninformative Advertising Content. *Marketing Science*, 37(2):236–258, 2018.
- J. Schumann. Targeted Online Advertising: Using Reciprocity Appeals to Increase Acceptance Among Users of Free Web Services. *Journal of Marketing*, 78(1):59–75, 2014.
- H. Sheng, F. F.-H. Nah, and K. Siau. An Experimental Study on Ubiquitous Commerce Adoption: Impact of Personalization and Privacy Concerns. *Journal of the Association for Information Systems*, 9(6):344–376, 2008.
- I. Simonson. Determinants of Customers’ Responses to Customized Offers: Conceptual Framework and Research Propositions. *Journal of Marketing*, 69(1):32–45, 2005.

- H. J. Smith, S. J. Milberg, S. J. Burke, and O. N. Hall. Information Privacy: Measuring Individuals' Concerns about Organizational Practices. *MIS Quarterly*, 20(2):167–196, 1996.
- N. K. Smith, J. T. Cacioppo, J. T. Larsen, and T. L. Chartrand. May I Have Your Attention, Please: Electrocortical Responses to Positive and Negative Stimuli. *Neuropsychologia*, 41(2):171–183, 2003.
- C. Snyder and H. Fromkin. *Uniqueness: The Pursuit of Difference*. Plenum, New York, 1980.
- J. H. Song, H. Y. Kim, S. Kim, S. W. Lee, and J.-H. Lee. Effects of Personalized E-Mail Messages on Privacy Risk: Moderating Roles of Control and Intimacy. *Marketing Letters*, 27(1):89–101, 2016.
- M. Sun, X. M. Zhang, and F. Zhu. Nonconformity in Online Social Networks. *Working Paper*, pages 1–31, 2017.
- J. Sutanto, E. Palme, C.-H. Tan, and C. W. Phang. Addressing the Personalization-Privacy Paradox: An Empirical Assessment from a Field Experiment on Smartphone Users. *MIS Quarterly*, 37(4):1141–1164, 2013.
- K. Tam and S. Ho. Understanding the Impact of Web Personalization on User Information Processing and Decision Outcomes. *MIS Quarterly*, 30(4):865–890, 2006.
- K. Y. Tam and S. Y. Ho. Web Personalization as a Persuasion Strategy: An Elaboration Likelihood Model Perspective. *Information Systems Research*, 16(3):271–291, 2005.
- S. Thirumalai and K. K. Sinha. To Personalize or Not to Personalize Online Purchase Interactions : Implications of Self-Selection by Retailers. *Information Systems Research*, 24(3):683–708, 2013.
- J. Treistman and J. P. Gregg. Visual, Verbal, and Sales Responses to Print Ads. *Journal of Advertising Research*, 19(4):41–47, 1979.
- C. Tucker. Social Networks, Personalized Advertising and Privacy Controls. *Journal of Marketing Research*, 51(5):546–562, 2014.
- C. Tucker. Social Advertising : How Advertising that Explicitly Promotes Social Influence Can Backfire. *Working Paper*, pages 1–28, 2016.
- J. W. Tukey. *Exploratory Data Analysis*. Addison-Wesley, 1977.
- S. Utz and N. C. Kramer. The Privacy Paradox on Social Network Sites Revisited: The Role of Individual Characteristics and Group Norms. *Cyberpsychology: Journal of Psychosocial Research on Cyberspace*, 3(2):1–10, 2009.

- W. Vickrey. Counterspeculation, Auctions, and Competitive Sealed Tenders. *The Journal of Finance*, 16(1):8–37, 1961.
- H. A. Wadlinger and D. M. Isaacowitz. Positive Mood Broadens Visual Attention to Positive Stimuli. *Motivation and Emotion*, 30(1):87–99, 2006.
- S. Wattal, R. Telang, T. Mukhopadhyay, and P. Boatwright. What’s in a Name? Impact of Use of Customer Information in E-Mail Advertisements. *Information Systems Research*, 23(3):679–697, 2012.
- M. Wedel and R. Pieters. Eye Fixations on Advertisements and Memory for Brands: A Model and Findings. *Marketing Science*, 19(4):297–312, 2000.
- M. Wedel and R. Pieters. A Review of Eye-Tracking Research in Marketing. In *Review of Marketing Research*, pages 123–147. Emerald Group Publishing Limited, 2007.
- T. B. White, D. L. Zahay, H. Thorbjørnsen, and S. Shavitt. Getting Too Personal: Reactance to Highly Personalized Email Solicitations. *Marketing Letters*, 19(1):39–50, 2008.
- Wishpond. How to Track Referrals Using Google Analytics. 2014. URL <http://corp.wishpond.com/facebook-marketing-resources-for-business/best-practices-to-performance-marketing-on-facebook/>.
- B. Xiao and I. Benbasat. E-Commerce Product Recommendation Agents: Use, Characteristics, and Impact. *MIS Quarterly*, 31(1):137–209, 2007.
- H. Xu, X. Luo, J. Carroll, and M. Rosson. The Personalization Privacy Paradox: An Exploratory Study of Decision Making Process for Location-Aware Marketing. *Decision Support Systems*, 51(1):42–52, 2011.
- J. Xu, X. Shao, J. Ma, K.-C. Lee, H. Qi, and Q. Lu. Lift-Based Bidding in Ad Selection. In *Proceedings of the 30th Conference on Artificial Intelligence (AAAI 2016)*, pages 651–657. 2016.
- M. S. Yadav, K. de Valck, T. Hennig-Thurau, D. L. Hoffman, and M. Spann. Social Commerce: A Contingency Framework for Assessing Marketing Potential. *Journal of Interactive Marketing*, 27(4):311–323, 2013.
- C. Y. Yoo. Implicit Memory Measures for Web Advertising Effectiveness. *Journalism and Mass Communication Quarterly*, 84(1):7–23, 2007.
- B. Zhang, N. Wang, and H. Jin. Privacy Concerns in Online Recommender Systems: Influences of Control and User Data Input. In *SOUPS ’14: Proceedings of the Tenth Symposium On Usable Privacy and Security*. 2014.
- J. Zhang and M. Wedel. The Effectiveness of Customized Promotions in Online and Offline Stores. *Journal of Marketing Research*, 46(4):190–206, 2009.

Back Matter

English Summary

The increasing availability of detailed data on consumers' characteristics online allows firms to personalize advertising to the preferences of these consumers. The personalization of ad messages offers firms tremendous potential. If done right, firms can address consumers with ad messages that are considered more relevant leading to more positive consumer responses to ads. Firms understand the potential of personalized advertising and aim to positively affect their bottom line with the personalization of ads. Simultaneously, firms struggle with how to design and implement personalization strategies. Supposedly, personalized advertising leads to an increase in firms' return on advertising investment. Nevertheless, firms face the challenge to correctly measure and assess advertising effectiveness to inform their marketing decisions. With this research, we advance the understanding of ad personalization and its implications for firms, consumers, and ad platforms.

With the help of a large-scale field experiment, addressing 198,234 individual consumers with personalized advertising, we present evidence for how firms should design their personalization strategies. We find that high levels of personalization specificity pay off for firms. At the same time, firms need to take the relationship of ad personalization with other advertising features into account when personalizing ads. We show that socially targeting personalized ads, where names of consumers' friends are included in the ad text, leads to less positive consumer responses to personalized ads.

Firms need to be aware that the use of consumers' information to personalize ads can trigger consumer privacy concerns. These privacy concerns negatively influence consumers' responses to personalized ads. To advance the understanding of privacy concerns in ad personalization, we conduct a lab experiment using eye tracking technology to assess the role of consumers' attention when confronted with personalized ads. Our findings show that firms cannot use intrusive ads, which cause consumer privacy concerns, to attract consumers' attention. Such a strategy is harmful as it decreases consumers' overall attention towards ads, eventually leading to less positive consumer responses.

In programmatic advertising firms outsource the ad allocation process, the decision to which consumer to serve an ad impression and how much to pay for this impression, to ad platforms. An examination of how ad platforms handle the ad allocation process reveals that contracts between firms and ad platforms might not be in the economic interest of firms. We show theoretically that ad platforms have an incentive to target consumers that are more likely to purchase independent of the effect of ads on their purchase probabilities. We conduct a large field experiment in which we analyze an ad platform's bidding behavior for 20,918 individual consumers over the duration of seven weeks. Our empirical analysis shows that the incentives specified in contracts between ad platforms and firms lead to an incentive misalignment that is harmful for firms. While ads generally increase consumers' likelihood to purchase, firms pay more for ads that are not providing higher value to them.

Nederlandse Samenvatting

De toenemende beschikbaarheid van gedetailleerde gegevens over de kenmerken van online consumenten stelt bedrijven in staat om advertenties te personaliseren. Het personaliseren van reclameboodschappen biedt bedrijven enorm veel potentieel, omdat het bedrijven in staat stelt advertenties te tonen die door consumenten als relevant worden beschouwd, wat leidt tot een positievere reactie op de advertentie. Hoewel bedrijven het potentieel van het personaliseren van advertenties begrijpen en ernaar streven om hun prestaties hiermee zoveel mogelijk te verbeteren, hebben ze tegelijkertijd moeite om hiervoor strategieën te ontwerpen en te implementeren. Bedrijven moeten de effectiviteit van reclame adequaat meten en beoordelen om marketingbeslissingen goed te onderbouwen. Met dit onderzoek bevorderen we het begrip van het personaliseren van advertenties en de bijbehorende implicaties voor bedrijven, consumenten en advertentieplatformen.

Op basis van een grootschalig veldexperiment met 198.234 individuele consumenten presenteren we bewijs voor de manier waarop bedrijven hun personaliseringsstrategieën moeten ontwerpen. Vergaande personalisatie van advertenties loont voor bedrijven. Tegelijkertijd dienen bedrijven de relatie tussen het personaliseren van advertenties en andere advertentiefuncties bij het personaliseren van advertenties in beschouwing te nemen. We laten zien dat gepersonaliseerde advertenties die gebruik maken van social targeting, waarbij namen van vrienden van de consumenten worden opgenomen in de advertenties, leiden tot minder positieve reacties van consumenten op advertenties.

Bedrijven moeten zich ervan bewust zijn dat het gebruik van consumenteninformatie om advertenties te personaliseren privacy-zorgen kan oproepen bij consumenten. Deze zorgen hebben een negatieve invloed op de reacties van consumenten op gepersonaliseerde advertenties. Om het begrip van privacy-zorgen in het personaliseren van advertenties te vergroten, voerden we een laboratorium experiment uit met eye-tracking om de rol van de aandacht van consumenten te beoordelen wanneer ze geconfronteerd worden met gepersonaliseerde advertenties. Onze bevindingen laten zien dat bedrijven geen opdringerige advertentie zouden moeten gebruiken, daar dit privacy-zorgen juist aanwakkert. Dergelijke advertenties zijn schadelijk omdat de algemene aandacht van consumenten voor een advertentie afneemt en uiteindelijk verkleint dit de kans dat consumenten op advertenties klikken.

Vaak besteden bedrijven de toewijzing van advertenties (de beslissing aan welke consument een bepaalde advertentie wordt getoond en voor welke prijs) uit aan advertentieplatformen. Onderzoek naar de wijze waarop advertentieplatformen dit aanpakken laat zien dat de huidige contracten tussen bedrijven en advertentieplatformen mogelijk niet in het economisch belang van bedrijven zijn. We tonen theoretisch aan dat advertentieplatformen een belang hebben zich te richten op consumenten die waarschijnlijk al een hogere kans tot aanschaf hebben, ongeacht het effect van de advertenties op deze kans. We hebben een groot veldexperiment uitgevoerd waarin we het gedrag van een advertentieplatform gericht op 20.918 individuele consumenten gedurende een periode van zeven weken analyseren. Onze empirische analyse toont

aan dat de gespecificeerde beloningen in de contracten conflicteren met de belangen van bedrijven en schadelijk voor hen zijn. Terwijl advertenties over het algemeen bij de consument de kans tot aanschaf vergroten, betalen bedrijven meer voor advertenties die hen geen grotere waarde bieden.

About the Author



Thomas Walter Frick was born in the year 1987 in Hechingen, Germany. He received a Bachelor of Science in Business Administration from Ludwig-Maximilians-Universität Munich and a Master of Science in Business Information Management from Rotterdam School of Management, Erasmus University. In his Master's thesis, Thomas investigated the implications of social media for physical and digital music consumption.

In 2013, Thomas started his PhD research at Rotterdam School of Management, Erasmus University within the Business Information Management section of the department of Technology and Operations Management.

In his research Thomas focuses on the implications of personalized digital advertising for firms, consumers, and ad platforms. In most of his research Thomas uses experiment methodology, with a focus on field experiments, to uncover causal relationships within the context of personalized digital advertising. Thomas is passionate about collaborating with firms and translating academic findings to business insights.

Thomas attended multiple international conferences to present his research, including ICIS, WISE, INFORMS, ISMS MARKETING SCIENCE CONFERENCE and others.

Since 1st of August 2018, Thomas is working as Assistant Professor in the Department of Digitalization at Copenhagen Business School.

Author Portfolio

Peer Reviewed Conference Proceedings

Frick, T.W. & Li, T. (2016). Personalization in Social Retargeting: A Field Experiment. In *International Conference on Information Systems (ICIS)*, Dublin, Ireland.

Tsekouras, D., **Frick, T.W.** & Li, T. (2016). Don't Take It Personally: The Effect of Explicit Targeting in Advertising Personalization. In *International Conference on Information Systems (ICIS)*, Dublin, Ireland.

Submitted Papers

“Don't Take it Personally: Investigating the Effect of Explicit Behavioral Targeting and Ad Message Framing” with Dimitrios Tsekouras and Ting Li
under journal review.

“Personalization Specificity in Social Retargeting - A Field Experiment” with Ting Li
under journal review.

Working Papers

“Pay For What You Get - Incentive Misalignments in Programmatic Advertising: Evidence from a Randomized Field Experiment” with Rahul Telang and Rodrigo Belo

“Social Influence and Visual Attention in the Personalization Privacy Paradox: An Eye Tracking Study” with Ting Li and Paul Pavlou

“Challenges in Performance Measurement for Digital Advertising” with Rodrigo Belo

PhD Courses

Advanced Statistical Methods

Applied Econometrics

Experimental Methods in Business Research

Mediation, Moderation, and Conditional Process Modeling

Programming

Panel Data Econometrics: Theory and Practice

Publishing Strategy

Statistical Methods

Statistical Modeling of Emerging Data Sets

Teaching

As Main Lecturer and Coordinator

Business Architecture & Consultancy (MSc 2015, 2016)
Master's Theses Supervision, MSc Business Information Management

As Teaching Assistant

Business Architecture & Consultancy (MSc 2013, 2014)
Information Strategy (MSc 2013, 2014)
Big Data & Analytics (MSc 2015)

Invited Guest Lectures

Driving Digital and Social Strategy (BSc 2014, 2015, 2016)
Customer Centric Digital Commerce (MSc 2014, 2015)

Conferences

Frick, T.W., Telang, R. & Belo R. (2018). Pay For What You Get - Incentive Misalignments in Programmatic Advertising: Evidence from a Randomized Field Experiment. In *The Economics of Information and Communication Technologies, ZEW Conference, Centre for European Economics Research*, Mannheim, Germany.

Frick, T.W., Telang, R. & Belo R. (2018). Pay For What You Get - Incentive Misalignments in Programmatic Advertising: Evidence from a Randomized Field Experiment. In *Statistical Challenges in eCommerce Research Symposium 2018 (SCECR)*, Rotterdam, The Netherlands.

Frick, T.W., Telang, R. & Belo R. (2017). Pay For What You Get - Incentive Misalignments in Programmatic Advertising: Evidence from a Randomized Field Experiment. In *Workshop on Information Systems and Economics 2017 (WISE)*, Seoul, South Korea.

Frick, T.W. & Telang, R. (2017). Pay For What You Get - Incentive Misalignments in Programmatic Advertising: Evidence from a Randomized Field Experiment. In *Conference on Digital Experimentation 2017 (CODE)*, Boston, US.

Frick, T.W. & Telang, R. (2017). Pay For What You Get - Incentive Misalignments in Programmatic Advertising: Evidence from a Randomized Field Experiment. In *NBER Summer Institute 2017, Economics of Information Technology and Digitization Workshop*, Boston, US.

Frick, T.W., Li, T. & Pavlou, P. (2017). The Role of Social Influence and Attention in the Personalization Privacy Paradox: An Eye Tracking Study. In *Statistical Challenges in eCommerce Research Symposium (SCECR)*, Hanoi, Vietnam.

Frick, T.W. & Li, T. (2016). Personalization in Social Retargeting: A Field Experiment. In *International Conference on Information Systems (ICIS)*, Dublin, Ireland.

Tsekouras, D., Frick, T.W. & Li, T. (2016). Don't Take It Personally: The Effect of Explicit Targeting in Advertising Personalization. In *International Conference on Information Systems (ICIS)*, Dublin, Ireland.

Frick, T.W., Li, T. & Pavlou, P. (2016). Investigating the Impact of Social Influence on the Personalization-Privacy Paradox: An Eye Tracking Study. In *INFORMS Annual Meeting*, Nashville, US.

Frick, T.W. & Li, T. (2016). Social Retargeting: A Field Experiment. In *Statistical Challenges in eCommerce Research Symposium (SCECR)*, Naxos, Greece.

Frick, T.W. & Li, T. (2016). Social Retargeting: A Field Experiment. In *The Economics of Information and Communication Technologies, ZEW Conference, Centre for European Economics Research*, Mannheim, Germany.

Frick, T.W. & Li, T. (2015). Social Retargeting A Randomized Field Experiment. In *37th ISMS Marketing Science Conference*, Baltimore, US.

Frick, T.W. & Li, T. (2015). Understanding Information Privacy Concerns in Social Advertising: An Eye Tracking Study. In *37th ISMS Marketing Science Conference*, Baltimore, US.

Frick, T.W., Tsekouras, D. & Li, T. (2014). For Whom, How, and Where? Examining the Relationship between Social Media and Music Album Performance. In *Statistical Challenges in eCommerce Research Symposium (SCECR)*, Tel Aviv, Israel.

Frick, T.W., Tsekouras, D. & Li, T. (2014). The Times They Are A-Changin: Examining the Impact of Social Media on Music Album Performance. In *Annual Meeting of the Academy of Management*, Philadelphia, US.

Frick, T.W. & Tsekouras, D. (2013). The Times They Are A-Changin: Examining the Impact of Social Media on Music Album Sales and Piracy. In *7th ICT Conference Paris on The Economics of Information and Communications Technologies*, Paris, France.

The ERIM PhD Series

The ERIM PhD Series contains PhD dissertations in the field of Research in Management defended at Erasmus University Rotterdam and supervised by senior researchers affiliated to the Erasmus Research Institute of Management (ERIM). All dissertations in the ERIM PhD Series are available in full text through the ERIM Electronic Series Portal: <http://repub.eur.nl/pub>. ERIM is the joint research institute of the Rotterdam School of Management (RSM) and the Erasmus School of Economics at the Erasmus University Rotterdam (EUR).

Dissertations in the last five years

Abbink, E.J., *Crew Management in Passenger Rail Transport*, Promotors: Prof. L.G. Kroon & Prof. A.P.M. Wagelmans, EPS-2014-325-LIS, <http://repub.eur.nl/pub/76927>

Acar, O.A., *Crowdsourcing for Innovation: Unpacking Motivational, Knowledge and Relational Mechanisms of Innovative Behavior in Crowdsourcing Platforms*, Promotor: Prof. J.C.M. van den Ende, EPS-2014-321-LIS, <http://repub.eur.nl/pub/76076>

Akin Ates, M., *Purchasing and Supply Management at the Purchase Category Level: Strategy, structure and performance*, Promotors: Prof. J.Y.F. Wynstra & Dr E.M. van Raaij, EPS-2014-300-LIS, <http://repub.eur.nl/pub/50283>

Akpinar, E., *Consumer Information Sharing*, Promotor: Prof. A. Smidts, EPS-2013-297-MKT, <http://repub.eur.nl/pub/50140>

Alexander, L., *People, Politics, and Innovation: A Process Perspective*, Promotors: Prof. H.G. Barkema & Prof. D.L. van Knippenberg, EPS-2014-331-S&E, <http://repub.eur.nl/pub/77209>

Alexiou, A., *Management of Emerging Technologies and the Learning Organization: Lessons from the Cloud and Serious Games Technology*, Promotor: Prof. S.J. Magala, Prof. M.C. Schippers and Dr I. Oshri, EPS-2016-404-ORG, <http://repub.eur.nl/pub/93818>

Almeida e Santos Nogueira, R.J. de, *Conditional Density Models Integrating Fuzzy and Probabilistic Representations of Uncertainty*, Promotors: Prof. U. Kaymak & Prof. J.M.C. Sousa, EPS-2014-310-LIS, <http://repub.eur.nl/pub/51560>

Bannouh, K., *Measuring and Forecasting Financial Market Volatility using High-frequency Data*, Promotor: Prof. D.J.C. van Dijk, EPS-2013-273-F&A, <http://repub.eur.nl/pub/38240>

Ben-Menahem, S.M., *Strategic Timing and Proactiveness of Organizations*, Promotors: Prof. H.W. Volberda & Prof. F.A.J. van den Bosch, EPS-2013-278-S&E, <http://repub.eur.nl/pub/39128>

Benschop, N., *Biases in Project Escalation: Names, frames & construal levels*, Promotors: Prof. K.I.M. Rhode, Prof. H.R. Commandeur, Prof. M. Keil & Dr A.L.P. Nuijten, EPS-2015-375-S&E, <http://repub.eur.nl/pub/79408>

Berg, W.E. van den, *Understanding Salesforce Behavior using Genetic Association Studies*, Promotor: Prof. W.J.M.I. Verbeke, EPS-2014-311-MKT, <http://repub.eur.nl/pub/51440>

Beusichem, H.C. van, *Firms and Financial Markets: Empirical Studies on the Informational Value of Dividends, Governance and Financial Reporting*, Promotors: Prof. A. de Jong & Dr G. Westerhuis, EPS-2016-378-F&A, <http://repub.eur.nl/pub/93079>

Blik, R. de, *Empirical Studies on the Economic Impact of Trust*, Promotors: Prof. J. Veenman & Prof. Ph.H.B.F. Franses, EPS-2015-324-ORG, <http://repub.eur.nl/pub/78159>

Boons, M., *Working Together Alone in the Online Crowd: The Effects of Social Motivations and Individual Knowledge Backgrounds on the Participation and Performance of Members of Online Crowdsourcing Platforms*, Promotors: Prof. H.G. Barkema & Dr D.A. Stam, EPS-2014-306-S&E, <http://repub.eur.nl/pub/50711>

Brazys, J., *Aggregated Marcoeconomic News and Price Discovery*, Promotor: Prof. W.F.C. Verschoor, EPS-2015-351-F&A, <http://repub.eur.nl/pub/78243>

Byington, E., *Exploring Coworker Relationships: Antecedents and Dimensions of Interpersonal Fit, Coworker Satisfaction, and Relational Models*, Promotor: Prof. D.L. van Knippenberg, EPS-2013-292-ORG, <http://repub.eur.nl/pub/41508>

Cancurtaran, P., *Essays on Accelerated Product Development*, Promotors: Prof. F. Langerak & Prof. G.H. van Bruggen, EPS-2014-317-MKT, <http://repub.eur.nl/pub/76074>

Caron, E.A.M., *Explanation of Exceptional Values in Multi-dimensional Business Databases*, Promotors: Prof. H.A.M. Daniels & Prof. G.W.J. Hendrikse, EPS-2013-296-LIS, <http://repub.eur.nl/pub/50005>

Carvalho, L. de, *Knowledge Locations in Cities: Emergence and Development Dynamics*, Promotor: Prof. L. Berg, EPS-2013-274-S&E, <http://repub.eur.nl/pub/38449>

Cranenburgh, K.C. van, *Money or Ethics: Multinational corporations and religious organisations operating in an era of corporate responsibility*, Promotors: Prof. L.C.P.M. Meijs, Prof. R.J.M. van Tulder & Dr D. Arenas, EPS-2016-385-ORG, <http://repub.eur.nl/pub/93104>

Consiglio, I., *Others: Essays on Interpersonal and Consumer Behavior*, Promotor: Prof. S.M.J. van Osselaer, EPS-2016-366-MKT, <http://repub.eur.nl/pub/79820>

Cox, R.H.G.M., *To Own, To Finance, and To Insure - Residential Real Estate Revealed*, Promotor: Prof. D. Brounen, EPS-2013-290-F&A, <http://repub.eur.nl/pub/40964>

Darnihamedani, P., *Individual Characteristics, Contextual Factors and Entrepreneurial Behavior*, Promotors: Prof. A.R. Thurik & S.J.A. Hessels, EPS-2016-360-S&E, <http://repub.eur.nl/pub/93280>

Deng, W., *Social Capital and Diversification of Cooperatives*, Promotor: Prof. G.W.J. Hendrikse, EPS-2015-341-ORG, <http://repub.eur.nl/pub/77449>

Depecik, B.E., *Revitalizing brands and brand: Essays on Brand and Brand Portfolio Management Strategies*, Promotor: Prof. G.H. van Bruggen, Dr Y.M. van Everdingen and Dr M.B. Ataman, EPS-2016-406-MKT, <http://repub.eur.nl/pub/93507>

Dollevoet, T.A.B., *Delay Management and Dispatching in Railways*, Promotor: Prof. A.P.M. Wagelmans, EPS-2013-272-LIS, <http://repub.eur.nl/pub/38241>

Duyvesteyn, J.G., *Empirical Studies on Sovereign Fixed Income Markets*, Promotors: Prof. P. Verwijmeren & Prof. M.P.E. Martens, EPS-2015-361-F&A, <http://hdl.handle.net/1765/79033>

Duursema, H., *Strategic Leadership: Moving Beyond the Leader-Follower Dyad*, Promotor: Prof. R.J.M. van Tulder, EPS-2013-279-ORG, <http://repub.eur.nl/pub/39129>

Elemes, A., *Studies on Determinants and Consequences of Financial Reporting Quality*, Promotor: Prof. E. Peek, EPS-2015-354-F&A, <http://hdl.handle.net/1765/79037>

Ellen, S. ter, *Measurement, Dynamics, and Implications of Heterogeneous Beliefs in Financial Markets*, Promotor: Prof. W.F.C. Verschoor, EPS-2015-343-F&A, <http://repub.eur.nl/pub/78191>

Erlemann, C., *Gender and Leadership Aspiration: The Impact of the Organizational Environment*, Promotor: Prof. D.L. van Knippenberg, EPS-2016-376-ORG, <http://repub.eur.nl/pub/79409>

Eskenazi, P.I., *The Accountable Animal*, Promotor: Prof. F.G.H. Hartmann, EPS-2015-355-F&A, <http://repub.eur.nl/pub/78300>

Evangelidis, I., *Preference Construction under Prominence*, Promotor: Prof. S.M.J. van Osselaer, EPS-2015-340-MKT, <http://repub.eur.nl/pub/78202>

Faber, N., *Structuring Warehouse Management*, Promotors: Prof. M.B.M. de Koster & Prof. A. Smidts, EPS-2015-336-LIS, <http://repub.eur.nl/pub/78603>

Fernald, K., *The Waves of Biotechnological Innovation in Medicine: Interfirm Cooperation Effects and a Venture Capital Perspective*, Promotors: Prof. E. Claassen, Prof. H.P.G. Pennings & Prof. H.R. Commandeur, EPS-2015-371-S&E, <http://hdl.handle.net/1765/79120>

Fisch, C.O., *Patents and trademarks: Motivations, antecedents, and value in industrialized and emerging markets*, Promotors: Prof. J.H. Block, Prof. H.P.G. Pennings & Prof. A.R. Thurik, EPS-2016-397-S&E, <http://repub.eur.nl/pub/94036>

Fliers, P.T., *Essays on Financing and Performance: The role of firms, banks and board*, Promotors: Prof. A. de Jong & Prof. P.G.J. Roosenboom, EPS-2016-388-F&A, <http://repub.eur.nl/pub/93019>

Fourne, S.P., *Managing Organizational Tensions: A Multi-Level Perspective on Exploration, Exploitation and Ambidexterity*, Promotors: Prof. J.J.P. Jansen & Prof. S.J. Magala, EPS-2014-318-S&E, <http://repub.eur.nl/pub/76075>

Gaast, J.P. van der, *Stochastic Models for Order Picking Systems*, Promotors: Prof. M.B.M. de Koster & Prof. I.J.B.F. Adan, EPS-2016-398-LIS, <http://repub.eur.nl/pub/93222>

Glorie, K.M., *Clearing Barter Exchange Markets: Kidney Exchange and Beyond*, Promotors: Prof. A.P.M. Wagelmans & Prof. J.J. van de Klundert, EPS-2014-329-LIS, <http://repub.eur.nl/pub/77183>

Hekimoglu, M., *Spare Parts Management of Aging Capital Products*, Promotor: Prof. R. Dekker, EPS-2015-368-LIS, <http://repub.eur.nl/pub/79092>

Heyde Fernandes, D. von der, *The Functions and Dysfunctions of Reminders*, Promotor: Prof. S.M.J. van Osselaer, EPS-2013-295-MKT, <http://repub.eur.nl/pub/41514>

Hogenboom, A.C., *Sentiment Analysis of Text Guided by Semantics and Structure*, Promotors: Prof. U. Kaymak & Prof. F.M.G. de Jong, EPS-2015-369-LIS, <http://repub.eur.nl/pub/79034>

Hogenboom, F.P., *Automated Detection of Financial Events in News Text*, Promotors: Prof. U. Kaymak & Prof. F.M.G. de Jong, EPS-2014-326-LIS, <http://repub.eur.nl/pub/77237>

Hollen, R.M.A., *Exploratory Studies into Strategies to Enhance Innovation-Driven International Competitiveness in a Port Context: Toward Ambidextrous Ports*, Promotors: Prof. F.A.J. Van Den Bosch & Prof. H.W. Volberda, EPS-2015-372-S&E, <http://repub.eur.nl/pub/78881>

Hout, D.H. van, *Measuring Meaningful Differences: Sensory Testing Based Decision Making in an Industrial Context; Applications of Signal Detection Theory and Thurstonian Modelling*, Promotors: Prof. P.J.F. Groenen & Prof. G.B. Dijksterhuis, EPS- 2014-304-MKT, <http://repub.eur.nl/pub/50387>

Houwelingen, G.G. van, *Something To Rely On*, Promotors: Prof. D. de Cremer & Prof. M.H. van Dijke, EPS-2014-335-ORG, <http://repub.eur.nl/pub/77320>

Hurk, E. van der, *Passengers, Information, and Disruptions*, Promotors: Prof. L.G. Kroon & Prof. P.H.M. Vervest, EPS-2015-345-LIS, <http://repub.eur.nl/pub/78275>

Iseger, P. den, *Fourier and Laplace Transform Inversion with Applications in Finance*, Promotor: Prof. R. Dekker, EPS-2014-322-LIS, <http://repub.eur.nl/pub/76954>

Jaarsveld, W.L. van, *Maintenance Centered Service Parts Inventory Control*, Promotor: Prof. R. Dekker, EPS-2013-288-LIS, <http://repub.eur.nl/pub/39933>

Khanagha, S., *Dynamic Capabilities for Managing Emerging Technologies*, Promotor: Prof. H.W. Volberda, EPS-2014-339-S&E, <http://repub.eur.nl/pub/77319>

Kil, J., *Acquisitions Through a Behavioral and Real Options Lens*, Promotor: Prof. H.T.J. Smit, EPS-2013-298-F&A, <http://repub.eur.nl/pub/50142>

Klooster, E. van't, *Travel to Learn: the Influence of Cultural Distance on Competence Development in Educational Travel*, Promotors: Prof. F.M. Go & Prof. P.J. van Baalen, EPS-2014-312-MKT, <http://repub.eur.nl/pub/51462>

Koendjibiarie, S.R., *The Information-Based View on Business Network Performance: Revealing the Performance of Interorganizational Networks*, Promotors: Prof. H.W.G.M. van Heck & Prof. P.H.M. Vervest, EPS-2014-315-LIS, <http://repub.eur.nl/pub/51751>

Koning, M., *The Financial Reporting Environment: The Role of the Media, Regulators and Auditors*, Promotors: Prof. G.M.H. Mertens & Prof. P.G.J. Roosenboom, EPS-2014-330-F&A, <http://repub.eur.nl/pub/77154>

Konter, D.J., *Crossing Borders with HRM: An Inquiry of the Influence of Contextual Differences in the Adoption and Effectiveness of HRM*, Promotors: Prof. J. Paauwe & Dr L.H. Hoeksema, EPS-2014-305-ORG, <http://repub.eur.nl/pub/50388>

Korkmaz, E., *Bridging Models and Business: Understanding Heterogeneity in Hidden Drivers of Customer Purchase Behavior*, Promotors: Prof. S.L. van de Velde & Prof. D. Fok, EPS-2014-316-LIS, <http://repub.eur.nl/pub/76008>

Krämer, R., *A license to mine? Community organizing against multinational corporations*, Promotors: Prof. R.J.M. van Tulder & Prof. G.M. Whiteman, EPS-2016-383-ORG, <http://repub.eur.nl/pub/94072>

Kroezen, J.J., *The Renewal of Mature Industries: An Examination of the Revival of the Dutch Beer Brewing Industry*, Promotor: Prof. P.P.M.A.R. Heugens, EPS-2014- 333-S&E, <http://repub.eur.nl/pub/77042>

Kysucky, V., *Access to Finance in a Cross-Country Context*, Promotor: Prof. L. Norden, EPS-2015-350-F&A, <http://repub.eur.nl/pub/78225>

Lee, C.I.S.G., *Big Data in Management Research: Exploring New Avenues*, Promotors: Prof. S.J. Magala & Dr W.A. Felts, EPS-2016-365-ORG, <http://repub.eur.nl/pub/79818>

Legault-Tremblay, P.O., *Corporate Governance During Market Transition: Heterogeneous responses to Institution Tensions in China*, Promotor: Prof. B. Krug, EPS-2015-362-ORG, <http://repub.eur.nl/pub/78649>

Lenoir, A.S., *Are You Talking to Me? Addressing Consumers in a Globalised World*, Promotors: Prof. S. Puntoni & Prof. S.M.J. van Osselaer, EPS-2015-363-MKT, <http://repub.eur.nl/pub/79036>

Leunissen, J.M., *All Apologies: On the Willingness of Perpetrators to Apologize*, Promotors: Prof. D. de Cremer & Dr M. van Dijke, EPS-2014-301-ORG, <http://repub.eur.nl/pub/50318>

Li, D., *Supply Chain Contracting for After-sales Service and Product Support*, Promotor: Prof. M.B.M. de Koster, EPS-2015-347-LIS, <http://repub.eur.nl/pub/78526>

Li, Z., *Irrationality: What, Why and How*, Promotors: Prof. H. Bleichrodt, Prof. P.P. Wakker, & Prof. K.I.M. Rohde, EPS-2014-338-MKT, <http://repub.eur.nl/pub/77205>

Liang, Q.X., *Governance, CEO Identity, and Quality Provision of Farmer Cooperatives*, Promotor: Prof. G.W.J. Hendrikse, EPS-2013-281-ORG, <http://repub.eur.nl/pub/39253>

Liket, K., *Why 'Doing Good' is not Good Enough: Essays on Social Impact Measurement*, Promotors: Prof. H.R. Commandeur & Dr K.E.H. Maas, EPS-2014-307-STR, <http://repub.eur.nl/pub/51130>

Loos, M.J.H.M. van der, *Molecular Genetics and Hormones: New Frontiers in Entrepreneurship Research*, Promotors: Prof. A.R. Thurik, Prof. P.J.F. Groenen, & Prof. A. Hofman, EPS-2013-287-S&E, <http://repub.eur.nl/pub/40081>

Lu, Y., *Data-Driven Decision Making in Auction Markets*, Promotors: Prof. H.W.G.M. van Heck & Prof. W. Ketter, EPS-2014-314-LIS, <http://repub.eur.nl/pub/51543>

Ma, Y., *The Use of Advanced Transportation Monitoring Data for Official Statistics*, Promotor: Prof. L.G. Kroon and Dr J. van Dalen, EPS-2016-391-LIS, <http://repub.eur.nl/pub/80174>

Manders, B., *Implementation and Impact of ISO 9001*, Promotor: Prof. K. Blind, EPS-2014-337-LIS, <http://repub.eur.nl/pub/77412>

Mell, J.N., *Connecting Minds: On The Role of Metaknowledge in Knowledge Coordination*, Promotor: Prof. D.L. van Knippenberg, EPS-2015-359-ORG, <http://hdl.handle.net/1765/78951>

Meulen, D. van der, *The Distance Dilemma: the effect of flexible working practices on performance in the digital workplace*, Promotors: Prof. H.W.G.M. van Heck & Prof. P.J. van Baalen, EPS-2016-403-LIS, <http://repub.eur.nl/pub/94033>

Micheli, M.R., *Business Model Innovation: A Journey across Managers' Attention and Inter-Organizational Networks*, Promotor: Prof. J.J.P. Jansen, EPS-2015-344-S&E, <http://repub.eur.nl/pub/78241>

Milea, V., *News Analytics for Financial Decision Support*, Promotor: Prof. U. Kaymak, EPS-2013-275-LIS, <http://repub.eur.nl/pub/38673>

Moniz, A., *Textual Analysis of Intangible Information*, Promotors: Prof. C.B.M. van Riel, Prof. F.M.G de Jong & Dr G.A.J.M. Berens, EPS-2016-393-ORG, <http://repub.eur.nl/pub/93001>

Mulder, J., *Network design and robust scheduling in liner shipping*, Promotors: Prof. R. Dekker & Dr W.L. van Jaarsveld, EPS-2016-384-LIS, <http://repub.eur.nl/pub/80258>

Naumovska, I., *Socially Situated Financial Markets: A Neo-Behavioral Perspective on Firms, Investors and Practices*, Promotors: Prof. P.P.M.A.R. Heugens & Prof. A. de Jong, EPS-2014-319-S&E, <http://repub.eur.nl/pub/76084>

Neerijnen, P., *The Adaptive Organization: the socio-cognitive antecedents of ambidexterity and individual exploration*, Promotors: Prof. J.J.P. Jansen, P.P.M.A.R. Heugens & Dr T.J.M. Mom, EPS-2016-358-S&E, <http://repub.eur.nl/pub/93274>

Oord, J.A. van, *Essays on Momentum Strategies in Finance*, Promotor: Prof. H.K. van Dijk, EPS-2016-380-F&A, <http://repub.eur.nl/pub/80036>

Pennings, C.L.P., *Advancements in Demand Forecasting: Methods and Behavior*, Promotors: Prof. L.G. Kroon, Prof. H.W.G.M. van Heck & Dr J. van Dalen, EPS-2016-400-LIS, <http://repub.eur.nl/pub/94039>

Peters, M., *Machine Learning Algorithms for Smart Electricity Markets*, Promotor: Prof. W. Ketter, EPS-2014-332-LIS, <http://repub.eur.nl/pub/77413>

Porck, J., *No Team is an Island: An Integrative View of Strategic Consensus between Groups*, Promotors: Prof. P.J.F. Groenen & Prof. D.L. van Knippenberg, EPS- 2013-299-ORG, <http://repub.eur.nl/pub/50141>

Pronker, E.S., *Innovation Paradox in Vaccine Target Selection*, Promotors: Prof. H.J.H.M. Claassen & Prof. H.R. Commandeur, EPS-2013-282-S&E, <http://repub.eur.nl/pub/39654>

Protzner, S., *Mind the gap between demand and supply: A behavioral perspective on demand forecasting*, Promotors: Prof. S.L. van de Velde & Dr L. Rook, EPS-2015-364-LIS, <http://repub.eur.nl/pub/79355>

Pruijssers, J.K., *An Organizational Perspective on Auditor Conduct*, Promotors: Prof. J. van Oosterhout & Prof. P.P.M.A.R. Heugens, EPS-2015-342-S&E, <http://repub.eur.nl/pub/78192>

Retel Helmrich, M.J., *Green Lot-Sizing*, Promotor: Prof. A.P.M. Wagelmans, EPS- 2013-291-LIS, <http://repub.eur.nl/pub/41330>

Rietdijk, W.J.R., *The Use of Cognitive Factors for Explaining Entrepreneurship*, Promotors: Prof. A.R. Thurik & Prof. I.H.A. Franken, EPS-2015-356-S&E, <http://repub.eur.nl/pub/79817>

Rietveld, N., *Essays on the Intersection of Economics and Biology*, Promotors: Prof. A.R. Thurik, Prof. Ph.D. Koellinger, Prof. P.J.F. Groenen, & Prof. A. Hofman, EPS-2014-320-S&E, <http://repub.eur.nl/pub/76907>

Rösch, D., *Market Efficiency and Liquidity*, Promotor: Prof. M.A. van Dijk, EPS-2015-353-F&A, <http://repub.eur.nl/pub/79121>

Roza, L., *Employee Engagement in Corporate Social Responsibility: A collection of essays*, Promotor: L.C.P.M. Meijs, EPS-2016-396-ORG, <http://repub.eur.nl/pub/93254>

Rubbaniy, G., *Investment Behaviour of Institutional Investors*, Promotor: Prof. W.F.C. Verschoor, EPS-2013-284-F&A, <http://repub.eur.nl/pub/40068>

Schoonees, P., *Methods for Modelling Response Styles*, Promotor: Prof.dr P.J.F. Groenen, EPS-2015-348-MKT, <http://repub.eur.nl/pub/79327>

Schouten, M.E., *The Ups and Downs of Hierarchy: the causes and consequences of hierarchy struggles and positional loss*, Promotors: Prof. D.L. van Knippenberg & Dr L.L. Greer, EPS-2016-386-ORG, <http://repub.eur.nl/pub/80059>

Shahzad, K., *Credit Rating Agencies, Financial Regulations and the Capital Markets*, Promotor: Prof. G.M.H. Mertens, EPS-2013-283-F&A, <http://repub.eur.nl/pub/39655>

Smit, J., *Unlocking Business Model Innovation: A look through the keyhole at the inner workings of Business Model Innovation*, Promotor: H.G. Barkema, EPS-2016-399-S&E, <http://repub.eur.nl/pub/93211>

Sousa, M.J.C. de, *Servant Leadership to the Test: New Perspectives and Insights*, Promoters: Prof. D.L. van Knippenberg & Dr D. van Dierendonck, EPS-2014-313-ORG, <http://repub.eur.nl/pub/51537>

Spliet, R., *Vehicle Routing with Uncertain Demand*, Promotor: Prof. R. Dekker, EPS-2013-293-LIS, <http://repub.eur.nl/pub/41513>

Staatd, J.L., *Leading Public Housing Organisation in a Problematic Situation: A Critical Soft Systems Methodology Approach*, Promotor: Prof. S.J. Magala, EPS-2014-308-ORG, <http://repub.eur.nl/pub/50712>

Stallen, M., *Social Context Effects on Decision-Making: A Neurobiological Approach*, Promotor: Prof. A. Smidts, EPS-2013-285-MKT, <http://repub.eur.nl/pub/39931>

Szatmari, B., *We are (all) the champions: The effect of status in the implementation of innovations*, Promoters: Prof J.C.M & Dr D. Deichmann, EPS-2016-401-LIS, <http://repub.eur.nl/pub/94633>

Tarakci, M., *Behavioral Strategy: Strategic Consensus, Power and Networks*, Promoters: Prof. D.L. van Knippenberg & Prof. P.J.F. Groenen, EPS-2013-280-ORG, <http://repub.eur.nl/pub/39130>

Tuijl, E. van, *Upgrading across Organisational and Geographical Configurations*, Promotor: Prof. L. van den Berg, EPS-2015-349-S&E, <http://repub.eur.nl/pub/78224>

Tuncdogan, A., *Decision Making and Behavioral Strategy: The Role of Regulatory Focus in Corporate Innovation Processes*, Promoters: Prof. F.A.J. van den Bosch, Prof. H.W. Volberda, & Prof. T.J.M. Mom, EPS-2014-334-S&E, <http://repub.eur.nl/pub/76978>

Uijl, S. den, *The Emergence of De-facto Standards*, Promotor: Prof. K. Blind, EPS-2014-328-LIS, <http://repub.eur.nl/pub/77382>

Vagias, D., *Liquidity, Investors and International Capital Markets*, Promotor: Prof. M.A. van Dijk, EPS-2013-294-F&A, <http://repub.eur.nl/pub/41511>

Valogianni, K., *Sustainable Electric Vehicle Management using Coordinated Machine Learning*, Promoters: Prof. H.W.G.M. van Heck & Prof. W. Ketter, EPS-2016-387-LIS, <http://repub.eur.nl/pub/93018>

Veelenturf, L.P., *Disruption Management in Passenger Railways: Models for Timetable, Rolling Stock and Crew Rescheduling*, Promotor: Prof. L.G. Kroon, EPS-2014-327-LIS, <http://repub.eur.nl/pub/77155>

Venus, M., *Demystifying Visionary Leadership: In search of the essence of effective vision communication*, Promotor: Prof. D.L. van Knippenberg, EPS-2013-289-ORG, <http://repub.eur.nl/pub/40079>

Vermeer, W., *Propagation in Networks: The impact of information processing at the actor level on system-wide propagation dynamics*, Promotor: Prof. P.H.M. Vervest, EPS-2015-373-LIS, <http://repub.eur.nl/pub/79325>

Versluis, I., *Prevention of the Portion Size Effect*, Promotors: Prof. Ph.H.B.F. Franses & Dr E.K. Papias, EPS-2016-382-MKT, <http://repub.eur.nl/pub/79880>

Vishwanathan, P., *Governing for Stakeholders: How Organizations May Create or Destroy Value for their Stakeholders*, Promotors: Prof. J. van Oosterhout & Prof. L.C.P.M. Meijs, EPS-2016-377-ORG, <http://repub.eur.nl/pub/93016>

Visser, V.A., *Leader Affect and Leadership Effectiveness: How leader affective displays influence follower outcomes*, Promotor: Prof. D.L. van Knippenberg, EPS-2013-286-ORG, <http://repub.eur.nl/pub/40076>

Vries, J. de, *Behavioral Operations in Logistics*, Promotors: Prof. M.B.M de Koster & Prof. D.A. Stam, EPS-2015-374-LIS, <http://repub.eur.nl/pub/79705>

Wagenaar, J.C., *Practice Oriented Algorithmic Disruption Management in Passenger Railways*, Promotors: Prof. L.G. Kroon & Prof. A.P.M. Wagelmans, EPS-2016-390-LIS, <http://repub.eur.nl/pub/93177>

Wang, P., *Innovations, status, and networks*, Promotors: Prof. J.J.P. Jansen & Dr V.J.A. van de Vrande, EPS-2016-381-S&E, <http://repub.eur.nl/pub/93176>

Wang, T., *Essays in Banking and Corporate Finance*, Promotors: Prof. L. Norden & Prof. P.G.J. Roosenboom, EPS-2015-352-F&A, <http://repub.eur.nl/pub/78301>

Wang, Y., *Corporate Reputation Management: Reaching Out to Financial Stakeholders*, Promotor: Prof. C.B.M. van Riel, EPS-2013-271-ORG, <http://repub.eur.nl/pub/38675>

Weenen, T.C., *On the Origin and Development of the Medical Nutrition Industry*, Promotors: Prof. H.R. Commandeur & Prof. H.J.H.M. Claassen, EPS-2014-309-S&E, <http://repub.eur.nl/pub/51134>

Wolfswinkel, M., *Corporate Governance, Firm Risk and Shareholder Value*, Promotor: Prof. A. de Jong, EPS-2013-277-F&A, <http://repub.eur.nl/pub/39127>

Yang, S., *Information Aggregation Efficiency of Prediction Markets*, Promotor: Prof. H.W.G.M. van Heck, EPS-2014-323-LIS, <http://repub.eur.nl/pub/77184>

Ypsilantis, P., *The Design, Planning and Execution of Sustainable Intermodal Port-hinterland Transport Networks*, Promotors: Prof. R.A. Zuidwijk & Prof. L.G. Kroon, EPS-2016-395-LIS, <http://repub.eur.nl/pub/94375>

Yuferova, D., *Price Discovery, Liquidity Provision, and Low-Latency Trading*, Promotors: Prof. M.A. van Dijk & Dr D.G.J. Bongaerts, EPS-2016-379-F&A, <http://repub.eur.nl/pub/93017>

Zaerpour, N., *Efficient Management of Compact Storage Systems*, Promotor: Prof. M.B.M. de Koster, EPS-2013-276-LIS, <http://repub.eur.nl/pub/38766>

Zuber, F.B., *Looking at the Others: Studies on (un)ethical behavior and social relationships in organizations*, Promotor: Prof. S.P. Kaptein, EPS-2016-394-ORG, <http://repub.eur.nl/pub/94388>

The personalization of advertising offers firms tremendous potential. If done right, firms can address consumers with more relevant ads, leading to more positive consumer responses. Nevertheless, firms are struggling with how to design personalization strategies and face the challenge to correctly assess advertising effectiveness. With this research, we advance the understanding of advertising personalization and its implications for firms, consumers, and ad platforms.

With the help of a large-scale field experiment, we present evidence for how firms should design their personalization strategies. We find that high levels of personalization specificity pay off for firms. At the same time, socially targeting personalized ads, where names of consumers' friends are included in the ad text, leads to less positive consumer responses.

To advance the understanding of privacy concerns in advertising personalization, we conduct a lab experiment using eye tracking technology. Our findings reveal that firms cannot use intrusive ads, that cause privacy concerns, to attract consumers' attention. Such a strategy is harmful as it decreases consumers' overall attention towards ads, eventually leading to less positive consumer responses.

An examination of contracts between firms and ad platforms exposes that these contracts might not be in the economic interest of firms. We conduct a large field experiment and our analysis reveals that currently implemented contracts between ad platforms and firms lead to an incentive misalignment that is harmful for firms. While ads generally increase consumers' likelihood to purchase, firms pay more for ads that are not providing higher value to them.

ERiM

The Erasmus Research Institute of Management (ERIM) is the Research School (Onderzoekschool) in the field of management of the Erasmus University Rotterdam. The founding participants of ERIM are the Rotterdam School of Management (RSM), and the Erasmus School of Economics (ESE). ERIM was founded in 1999 and is officially accredited by the Royal Netherlands Academy of Arts and Sciences (KNAW). The research undertaken by ERIM is focused on the management of the firm in its environment, its intra- and interfirm relations, and its business processes in their interdependent connections.

The objective of ERIM is to carry out first rate research in management, and to offer an advanced doctoral programme in Research in Management. Within ERIM, over three hundred senior researchers and PhD candidates are active in the different research programmes. From a variety of academic backgrounds and expertises, the ERIM community is united in striving for excellence and working at the forefront of creating new business knowledge.

ERiM

ERIM PhD Series **Research in Management**

Erasmus University Rotterdam (EUR)
Erasmus Research Institute of Management
Mandeville (T) Building
Burgemeester Oudlaan 50
3062 PA Rotterdam, The Netherlands

P.O. Box 1738
3000 DR Rotterdam, The Netherlands
T +31 10 408 1182
E info@erim.eur.nl
W www.erim.eur.nl