Aalto University
School of Business

UNDERSTANDING THE VALUE OF ONLINE AD TARGETING – ADVERTISER'S PERSPECTIVE

Master's Thesis
Kalle Tupamäki
Aalto University School of Business
Information and Service Management
Spring 2019

Author	Kalle Tupamäki
Title of thesis	Understanding the Value of Online Ad Targeting – Advertiser's Perspective
Degree	Master of Science in Economics and Business Administration
Degree programme	Information and Service Management
Thesis advisor(s)	Virpi Tuunainen
Year of approval 2019	**Number of pages** 82 — **Language** English

Abstract

Online advertising is a huge, growing industry and a key part of the growth is the improving efficiency of online advertising due to ad targeting. Targeting is a process that includes collecting data from various sources to help match advertisements with consumers who are most likely to be interested and to purchase the advertised products. While the consensus is that targeting can be used to provide value for advertisers, and for consumers, the incremental value is difficult to quantify.

In this paper, we take a look at some of the problems advertisers face when using targeting and present some existing solutions to mitigate these challenges. Earlier literature is used to exemplify how advertisers can effectively identify desired target audiences, choose the optimal methods for targeting to reach the selected audience and measure the success of their campaigns. One of the key challenges for advertisers however, appears to be difficulties understanding the platforms they use for online advertising.

Participant observation, a method often used for ethnographic research is presented here as a way for advertisers to learn about the ad platforms and the success of their campaigns. In business, participant observation has been used to analyze the quality of in-store services, but here its use is developed and demonstrated in an online environment. To exemplify how the method can be used in an online advertising context, this paper presents the findings of a study, where ads were observed on YouTube. The study is a combination of market research and marketing research. First, ads from four different markets are observed to find out whether market size and the level of data privacy regulations affect how advertisements appear to consumers. Then, additional observations are made in two of the markets, to see how disabling personalized advertising affects how adverts are distributed to consumers.

The study presented here, is intended as a proof-of-concept and advertisers can use the developed method in a multitude of ways to find out valuable insight about online advertising platforms. Participant observation can be used for example to better understand the setup process of targeted advertising campaigns, to optimize one's campaigns and to examine how one's competitors appear to be running their campaigns. In addition to numerous business use cases like these, the method can be used for academic purposes, as presented in this paper.

Keywords Online advertising, targeting, participant observation

Aalto-yliopisto, P.O. BOX 11000, 00076 AALTO
www.aalto.fi
Maisterintutkinnon tutkielman tiivistelmä

Aalto-yliopisto
Kauppakorkeakoulu

Tekijä Kalle Tupamäki	
Työn nimi Understanding the Value of Online Ad Targeting – Advertiser's Perspective	
Tutkinto Kauppatieteiden maisteri	
Koulutusohjelma Information and Service Management	
Työn ohjaaja(t) Virpi Tuunainen	
Hyväksymisvuosi 2019	**Sivumäärä** 82 · **Kieli** Englanti

Tiivistelmä

Internet-mainonta on suuri, kasvava ala ja markkinoinnin tehokkuuden parantaminen mainonnan kohdistamisen avulla on merkittävässä roolissa alan kasvun kannalta. Kohdistaminen on prosessi, joka sisältää datan keräämisen lukuisista lähteistä, mitä käytetään hyödyksi pyrkimyksessä näyttää mainokset sellaisille kuluttajille, jotka ovat kiinnostuneita mainostetuista tuotteista. Vaikka alalla ollaan melko yksimielisiä siitä, että kohdistaminen voi luoda arvoa sekä mainostajille että kuluttajille, sen tuoman lisäarvon määrää on vaikea mitata.

Tässä tutkimuksessa tuomme esille muutamia tyypillisiä ongelmia, joita mainostajat kohtaavat käyttäessään kohdistamista sekä esittelemme ratkaisuja näihin ongelmiin. Aiemmassa kirjallisuudessa on käsitelty kattavasti, kuinka mainostajat voivat tunnistaa halutut kohdeyleisöt, valita optimaaliset kohdistamismetodit sekä mitata kohdistettujen kampanjoidensa menestymistä. Aiemmassa kirjallisuudessa ei ole kuitenkaan käsitelty sitä, miten mainostajat voivat oppia ymmärtämään käyttämiään mainosalustoja, mikä vaikuttaa olevan iso ongelma alalla.

Osallistuva havainnointi, jota käytetään usein etnografisissa tutkimuksissa, esitetään tässä tutkimuksessa keinona, jolla mainostajat voivat tutkia mainosalustojen toimintaa ja kampanjoidensa menestymistä. Liiketoiminnassa osallistuvaa havainnointia on käytetty muun muassa palvelujen laadun tutkimisessa kaupoissa, mutta tässä tutkimuksessa sitä sovelletaan online-ympäristössä. Tässä paperissa esitetyt tulokset YouTubessa havaituista mainoksista kerätystä datasta ilmentävät, kuinka osallistuvaa havainnointia voidaan käyttää internet-mainonnan kontekstissa. Tässä työssä esitetty esimerkkitutkimus sisältää piirteitä sekä markkina- että markkinointitutkimuksesta. Ensin, mainoksia havainnoidaan neljältä eri markkina-alueelta, tarkoituksena selvittää, miten markkinan koko sekä yksityisyydensuojaan liittyvän lainsäädännön taso vaikuttavat siihen, minkälaisia mainoksia kuluttaja näkee. Tämän jälkeen, lisähavaintoja tehdään kahdelta markkina-alueelta, tarkoituksena tutkia, miten personoidun mainonnan kieltäminen vaikuttaa mainoksien distribuutioon.

Tässä työssä esitetty osallistuvan havainnoinnin tutkimus on tarkoitettu todisteeksi metodin toimivuudesta internet-markkinoinnin tutkimisessa. Mainostajat voivat käyttää metodia lukuisiin eri tarkoituksiin ymmärtääkseen paremmin käyttämiään mainosalustoja. Osallistuvaa havainnointia voidaan käyttää esimerkiksi ymmärtämään kampanjoiden valmisteluissa tehtäviä päätöksiä ja niiden vaikutuksia kampanjaan, kampanjoiden optimointiin sekä selvittämään, millaisia kampanjoita kilpailijat käyttävät. Lukuisten liiketoiminnan käyttötapausten lisäksi metodia voidaan käyttää myös akateemiseen tutkimukseen, kuten tässä työssä tehty tutkimus osoittaa.

Avainsanat Internet-mainonta, kohdistaminen, osallistuva havainnointi

Table of Contents

List of Tables

List of Figures

1 Introduction

Online advertising is a massive global industry that has been increasing rapidly over the last few decades. According to Statista (2018), yearly global digital ad spending amounted to an estimated $229B in 2017 and will increase to $335B by 2020. The growth is projected to continue in the foreseeable future due to three main reasons. Firstly, advertisers are understanding the value of online presence better, which results in more competition for valuable ad spots *(more demand)*, increasing their prices. Secondly, the size of the market keeps growing *(more supply)*, as the global internet usage keeps increasing. Finally, improving ad targeting methods are increasing the value of online advertising *(improved efficiency)*. As the growth in both supply and demand is likely to plateau sooner or later, objectively the third reason is the most intriguing to study, as it is more likely to provide sustainable growth in the industry.

My bachelor's thesis (Tupamäki, 2017) focused on uncovering the different methods of online ad targeting and discussed some of the current trends in the targeting industry. The results of an included empirical study suggested that ad targeting still has a lot of room for improvement. This is despite targeting services often costing much more than plain ad spots. This raises the question, how can companies estimate the value provided by targeting, to choose whether to buy it as a service. Thus, this paper looks at online ad targeting from the perspective of advertisers.

For advertisers, making the decision to pay for targeting services ultimately boils down to measuring or estimating the value of targeting. There are two sides to this. One is whether targeting efforts actually result in successful targeting (i.e. ads are shown to people who have an interest in the advertised product). The other relates to measuring the value of targeting to the company that is advertising (i.e. companies should not pay more than $1 for targeting an advertisement, if it doesn't increase profits by at least $1).

Both issues are important and interesting to study. The companies that are advertising usually can't affect the first issue, as targeting services are typically provided by another company, for example Google (Gill et al., 2013). This makes the second question even more important for most companies. How do companies estimate the effects of their ad targeting campaigns? Because they are done digitally, online ad campaigns could theoretically be tracked very accurately, for instance by using cookies. In practice, regulations, including the EU's recent General Data Protection Regulation (GDPR), limit this significantly however. Targeting service providers like Google offer advertisers some

key statistics (e.g. click-through rate) about their campaigns (Bien, n.d.). The literature review part of this study intends to find out how companies can validate these statistics independently, as some earlier literature exists on measuring the performance of targeted ad campaigns and even more research is focused on the accountability of online advertising in general. After this, a research method, which could help companies regarding the first issue, is developed and demonstrated. While it does not allow them to change the targeting services provided by others, it can potentially be used to understand the services better, which in turn will help them make more knowledgeable decisions regarding the use of targeting services.

Because online data privacy regulations differ from country to country, companies that operate online and have customers from many different countries need to comply with various laws. As data collection is becoming increasingly more regulated, especially recently in the EU due to the GDPR, the online advertising industry is changing. Some of these changes are discussed in the literature review and building on that, the empirical research studies how, if at all, the use of personalized targeting is changing due to new regulations. This is done by comparing advertising campaigns from various markets, some of which have stricter data privacy laws than others.

Online marketing campaign performance data is typically provided by the same party that provides the advertisers with advertising services (e.g. targeting), resulting in a conflict of interest, which undermines the trustworthiness of the data for the advertisers. As the service providers stand to benefit from offering a good service, it is in their best interest to make the performance data seem as good as possible. A framework is needed to help advertisers estimate the effects of online ad targeting independently, so that they can assess if they are getting what they are paying for.

A good example of how companies are not always made aware of what they are paying for is the "YouTube Adpocalypse". In 2017, Advertisers buying pre-roll video ad spots on YouTube realized that their advertisements were shown on videos that included unethical material, such as violence (Tupamäki, 2017). While Google's targeting algorithm likely worked as it was intended to, showing ads to those consumers who are most likely to be interested in the advertised products, regardless of the context in which the ad was shown, it is easy to understand why companies would not like their advertisements and products linked to such content. This scandal caused numerous companies to stop advertising on the site and Google was quick to implement a plan to fix things by better recognizing non-advertiser-friendly content. Though the commotion around the topic

calmed down for a while, new issues, which will be looked into further later in this paper, surfaced in early 2019.

Over the last few years, online privacy has become an increasingly discussed topic. As online ad targeting is largely dependent on user data, the public's growing concern about their online privacy is changing the way some companies approach online advertising. Companies that are accused of mishandling consumer data can quickly get a bad reputation leading to long-lasting detrimental effects, such as decreased sales. For example, some of Facebook's recent privacy scandals and other accusations have caused many people to stop using the social media platform. Lost trust can be extremely difficult to regain. For example, Facebook's data collection and data usage methods have been updated vigorously and in a relatively transparent way, but the public and the media still seem to have mixed opinions regarding the situation (Jenblat, 2018).

1.1 Research Problem

Earlier research suggests that a small number of data aggregators (i.e. companies that collect user data and use it to sell targeting services) are responsible for an increasing share of the ad targeting industry (Krishnamurthy & Wills, 2009). These aggregators include search engines such as Google and social media sites like Facebook (Gill et al., 2013). Google provides advertisers performance data about their campaigns, but this is somewhat limited by privacy regulations (Google, n.d.). For digital marketing to be financially justifiable, its effects must be measurable. This raises the question of the accuracy and credibility of the data, resulting in my research problem:

Online marketing campaign performance data is typically provided by the same party that provides the advertisers with advertising services (e.g. targeting), resulting in a conflict of interest, which undermines the trustworthiness of the data for the advertisers.

1.2 Research Questions

To tackle to research problem, the key question that needs to be answered is:

- How can companies accurately, and preferably independently, measure and/or estimate the value created by targeting in their online ad campaigns?

In addition to this, there are topical privacy-related issues currently at play in the online advertising industry. Furthermore, some advertisers have expressed their frustration at the lack of transparency in online ad targeting services. Therefore, this study aims to also answer the following questions about these issues:

- Have tighter online privacy regulations affected how companies approach and use ad targeting?
- How can advertisers evaluate the services provided by ad networks, with the intention of making better marketing decisions?

1.3 Motivation

In addition to my prior research, which suggested that there is still a lot of room to improve in online ad targeting (Tupamäki, 2017), there are multiple reasons why the topic is of interest. While the online advertising industry is interesting to study just based on its size and growth alone, it is even more so at this time, because regulations are starting to catch up with the collection and usage of online data. As ad targeting is typically done based on consumer data, it is undeniably affected by new regulations like the GDPR. As the EU's GDPR was implemented in 2018 and many other countries have since scrambled to update their privacy regulations as well, not much research has yet been done on how this affects the online advertising industry. Data scandals like Cambridge Analytica's data harvesting on Facebook with the purpose of political influence have also made consumers more aware of the scope and power of online data collection. This, among other reasons discussed later, is leading to increased use of privacy-protecting software, which can potentially have a massive impact on how online advertising is done in the future.

Another motivator for this study is personal experience as a consumer. While ad targeting works well in many cases, there are still plenty of situations that leaves one wondering why they are repeatedly advertised products and services that they have no interest in buying or using. This study aims to shed some light on some of the curious advertising decisions made by companies and possibly offer some advice for companies that are looking to improve their knowledge on online ad targeting.

1.4 Background

The purpose of this section is to briefly explain some of the key concepts of the online advertising industry. Firstly, it is important to understand who is involved in the advertising process. Earlier research on online advertising typically lists advertisers, ad networks and ad publishers as the main entities (Evans, 2009; Toubiana et al., 2010). Advertisers are trying to get their ads shown in relevant contexts to the right consumers (e.g. target audience). The ad networks are usually the most powerful players, who connect advertisers and publishers and offer them services, such as targeting, and tools like campaign management software. Publishers are the websites that are trying to monetize their content by having advertisements shown alongside it in varying formats. It is important to note that companies can take on multiple roles at once. For example, large search engine companies like Google and Baidu act in many situations as both ad networks and publishers. Some studies, particularly focusing on privacy-related issues, also emphasize the role of the web user, or consumer, in the online advertising process (Papaodyssefs et al. 2015). As this paper discusses targeting, which is in many cases heavily dependent on consumer data, the consumer's role in the advertising process is considered significant in this study.

At this point, it is also important to define what is meant by ad targeting in this paper. Instead of focusing on a single method of targeting, for example behavioral targeting, which is based on tracking consumers' browsing habits to understand their interests (Beales, 2010), this research applies a broader definition encompassing various different targeting methods. In order to create a better understanding of the value of online ad targeting, it must be recognized that different targeting methods serve different purposes. The value provided by a targeting method depends on the situation it is used in. For example, if a consumer is searching for plumbing services on a search engine, it is likely not useful to serve them advertisements for sporting goods in that situation, even if it can be determined that the consumer is in general interested in sports by using behavioral targeting. Instead location-based targeting and contextual advertising are more appropriate in this case, meaning that the consumer is shown advertisements for geographically-relevant plumbing service providers. For this reason, for the purposes of this research paper, online ad targeting is broadly defined as using any available data with the intention of better matching consumers and advertisements, to create value for advertisers, consumers and to those who provide the targeting services.

1.5 Structure of the Thesis

The literature review in Section 2 begins by showcasing some of the challenges companies face when setting up targeted ad campaigns and explains how companies can overcome these challenges. It then discusses the building blocks of online ad performance measurement, focusing on how the value of targeting can be singled out from these measurements. It also considers whether advertisers can use certain methods independently, or if they need the help of the advertising network or a third-party. In doing this, Section 2 aims to provide answers for the main research question. The literature review also touches on topical issues like tightening privacy regulations and some software-based problems that the online advertising industry is facing.

In Section 3, a research model is developed, which helps us examine how market differences affect how online advertising and targeting are done in different locations. The use of this model is demonstrated in Section 4. It is also proposed that this kind of methodology can be used by companies to evaluate the services provided by some advertising networks. Section 5 presents the results of the empirical research, which are discussed further in Section 6. Section 6 also includes a small case study about the YouTube Adpocalypse, which aims to emphasize the need for advertisers to better understand the online platforms they use for advertising. The limitations and suggestions for further research are also presented in this section. Section 7 presents the research summary and highlights the implications and contributions of the research.

2 Literature Review

In this Section, we will first discuss some of the difficulties companies face in online ad targeting. Then we use existing literature to find and propose solutions to these issues. In Sections 2.4 and 2.5 some topical issues for online advertising and targeting are discussed.

2.1 Challenges Faced in Online Ad Targeting

Defining successful targeting can be a tricky matter, as the definition of success varies in each situation. An easier way to tackle the issue at hand, is to consider the opposite – situations that are obvious failures. Although it is impossible to create a comprehensive list of everything that can go wrong with targeting, highlighting typical issues, their reasons and ways to avoid them is a good starting point. As targeting is a tool used for improving marketing campaigns, examining mistakes made in marketing in general is a good place to start, as they have been studied more over the years.

Tedesco (2006) and Koeppel (2005) mention several similar issues companies face regarding their marketing efforts. The most notable ones are not having a diverse enough marketing mix, not targeting to a specific audience and not understanding the lifetime value of customers. Raynor's (1992) research on niche marketing proposes that companies have to be careful when determining their target audience, to avoid becoming too dependent on one niche target market and to ensure that the niche in question is viable from a business point of view. All of these issues are also present in online ad targeting. Diversifying the marketing mix in this context means not leaning too heavily on one targeting method. Understanding the lifetime value of customers is also dependent on knowing how and when to use different targeting methods. Problems defining the target audience, whether it be too broad or too narrow, are a key issue in the targeting process. Procter and Gamble (P&G), one of the worlds top advertising and marketing spenders, provides a good example of the difficulty of determining a suitable target audience. The following case also brings forth other difficulties companies face when dealing with online ad targeting.

In 2016, P&G decided to reduce the amount of targeted advertising done on Facebook (Terlep & Seetharaman, 2016). The decision was made based on internal marketing performance measurements, which deemed that the company was hurting its sales by targeting too vigorously, limiting the overall reach of their campaigns. At this point, P&G's plan was to maintain the same level of spending on Facebook marketing, but

instead of targeting, the company intended to use cheaper non-targeted advertising to increase reach (ibid). However, less than one year later, P&G decided to significantly reduce its digital ad spending, when internal testing suggested that digital marketing was not as effective as expected (Bruell & Terlep, 2017). One of their key competitors, Unilever, also made similar decision around the same time (Dua, 2017). Both companies expressed concerns about the transparency of online advertising and the performance measurements provided to them by ad agencies and ad networks (ibid). P&G also emphasized the importance of understanding the platforms and websites that the company intends to advertise on (Neff, 2017). After further investigation, the company found that very short average ad viewing times and excessive repetition were the grassroot level issues in their targeted campaigns, which led to them not being as effective as intended (Vranica, 2018).

As the largest ad spender in the world, it is understandable that P&G would have difficulties keeping up with the performance of all of its campaigns and marketing efforts. However, it also means that the company employs numerous marketing professionals and even with their expertise as an asset, the company has difficulties understanding the diverse domain of digital advertising. Smaller companies might be able to focus more on fewer campaigns, but they lack the expertise a big player like P&G has, making it even more difficult for them to understand different online advertising platforms for example. It is important to note that P&G is definitely not the only company that is having problems with understanding the value of online advertising, but its case is emphasized here to show that even the large players that might have some influence on and power against the massive ad networks face problems. The following list compiles the typical challenges companies face regarding online ad targeting, which have been identified based on prior literature and P&G's well-documented troubles.

1. Identifying the target audience
2. Choosing how to use targeting
3. Measuring the performance of targeting
4. Understanding the online advertising platforms

P&G is not the only company that struggles with understanding how online advertising platforms work. One notable example of a situation where advertisers did not fully understand the platform they were advertising on is the so-called YouTube Adpocalypse.

Case – YouTube Adpocalypse

In 2017, numerous large companies stopped advertising on YouTube, when it was brought into their attention that their advertisements appeared on videos that contained violent or otherwise inappropriate content. Advertisers did not realize that by using personalized targeting in their YouTube ad campaigns, their advertisements would be shown to consumers that fit their target audience, regardless of the content they were currently watching (Cunningham, 2018). Understandably, many advertisers felt it was Google's responsibility to ensure that advertisements would not be shown next to questionable content, which could result in a bad brand image for the advertising company. As advertising generates a large part of the platform's and its content creators' revenues, a large-scale ad boycott was a devastating blow to the platform, hence resulting in the situation being referred to as the Adpocalypse.

The situation was seemingly resolved by the end of 2017 with better content evaluation processes, new guidelines for content creators and better campaign setup controls for advertisers (Stanford, 2018). However, a similar situation arose in early 2019. This time, a video posted on YouTube explained how the platforms algorithms were possibly being used by sexual predators to find sexually suggestive content of minors (Bergen et al., 2019). While the content of the videos was meant to be innocent, the comment sections of these videos contained predatory comments. Once again, many large advertisers, notably Disney and Nestle (Bergen et al., 2019), pulled their advertisements from the platform. Again, Google responded firmly and swiftly, initially by disabling commenting on numerous videos that feature children. Content creators were disappointed by this approach, as many of them are dependent on strong communities around their content, part of which is commenting on their videos. Additionally, the platform is implementing a new moderation algorithm that scans through comments and is better at recognizing and removing predatory comments than prior solutions (Hale, 2019).

Although the platform owner has to take a bulk of the responsibility when it comes to monitoring what happens on the platform, advertisers should not turn a blind eye and then act offended when unfortunate situations appear. On a platform as massive as YouTube, real-time content surveillance is incredibly difficult, if not impossible.

Understanding the platform where you advertise is key for avoiding unwanted situations like the adpocalypses on YouTube. Companies have to understand that there are many factors outside of their control in these campaigns, some of which can have harmful effects on their brands. While these adpocalypses received a lot of media attention,

YouTube is not the only platform that is susceptible to this kind of incidents. Essentially all platforms that allow the general public to post content, run the risk of ending up in a situation where ads are shown in unfavorable contexts.

2.2 Overcoming the Challenges

Especially when it comes to the first two challenges listed in Section 2.1, it is of utmost importance to understand what one wants to accomplish by using targeting. Whether the purpose is to increase sales, to increase brand awareness, to educate (potential) customers or anything else a company might aim to accomplish with a marketing campaign, it affects the optimal targeting strategy. Section 2.2.1 discusses different types of goals marketing campaigns can have. Marketers must then create strategies that enable them to reach these goals. These strategies include for example selecting who to target with the campaigns.

Once marketers have an idea of what kind of consumers they want to reach, they need to find those consumers by using data. Section 2.2.2 goes through a selection of data collection methods typically used for targeting, listing suggested use cases for the methods and discussing their pros and cons. It is important to note, that multiple methods are often used simultaneously to improve the effectiveness of targeting.

2.2.1 Well-Defined Campaign Strategy

As discussed, the first task for advertisers is to understand what they want to accomplish with targeted advertising. Again, a good place to begin, is to examine marketing objectives in general, without worrying about targeting. Typical goals include for example, increasing brand awareness and sales (Cotter, 2002), generating leads and improving the findability of one's products or service (Ruffolo, 2018). Win Marketing, a British marketing consultancy, also lists growing one's market share and entering new markets as typical goals for its customers (Win Marketing, n.d.).

Different objectives call for different kinds of targeting. For example, if the marketer is interested in increasing brand awareness, targeting can be relatively loosely done. Instead of trying to turn individual consumers into customers, the goal is to reach a large amount of people and let them know about your products and services. P&G's recent decisions exemplify this. If the objective is instead to grow one's market share, it might be logical to target your marketing efforts towards competitors' customers. This can

potentially allow a company to steal market share from competitors, if they are able to convince consumers that their products and services are superior. When setting up campaigns to improve the findability of your products, one key method is search engine optimization (SEO). It essentially means developing one's website in a way which makes search engines serve it as a top suggestion when potential customers are searching with relevant keywords (Yalçın & Köse, 2010). Paid placement is another option, which, while being more expensive, ensures that your page is shown as a top result for specific queries (Sen, 2005).

Ruffolo (2018) emphasizes the need for well-defined campaign objectives. Instead of stating a general need for more website visitors, a more useful objective is to set a precise numerical goal; for example, attract 20 000 website visitors. With clearly defined goals, not only is the campaign's performance easier to evaluate but communicating the purpose of the campaign becomes easier and there is less room for misunderstandings. However, before defining one's objective too strictly, marketers need to understand the audience they are trying to reach (Dodson, 2016). There is little sense in setting a goal that is unfeasible based on the amount of people that could be regarded as potential customers.

Key tasks for understanding one's audience and the potential to grow it include for example analyzing one's existing customers, examining what competitors are doing and defining one's value proposition (Porta, 2010; Cohn, 2015). The next step is creating customer profiles and ranking audiences (Dodson, 2016). Customer profiles are descriptions of ideal customers and are meant to serve as a basis for targeting. The traits listed in these profiles, which can include for example demographics, location, professional roles, interests and values (Ward, 2019; Dodson, 2016), are later used as criteria for targeting. Once desirable features of the target audience are determined, audience ranking comes into play. Dodson (2016) suggests using at least value to the business and accessibility as factors for ranking. The goal is to identify the parts of your target audience that show the most promise in terms of purchasing power and likelihood of purchasing and to recognize that all desired traits might not be feasible to use for targeting, as finding appropriate consumer data would be difficult or impossible (ibid.). Although the internet offers companies numerous ways to collect very precise consumer data, data collection has its limits. Next, we discuss the different ways in which companies can collect consumer data online.

2.2.2   Collecting Data

There are three main methods of data collection for the purpose of targeting ads. Behavioral targeting is based on the consumers prior browsing, demographic targeting allows segmenting the market and reaching a certain target group and contextual targeting is done by collecting data about the website that is being viewed, to serve consumers related advertisements (Tupamäki, 2017).

Behavioral targeting is done by collecting clickstream data, which allows companies to follow web users' browsing habits (Tupamäki, 2017). This helps them profile consumers, typically anonymously, as identities are not necessary for targeting, and find out who are most likely to be interested in an advertised product. The strength of this method is that it can potentially create very detailed perceptions of specific web users' interests, thus being a very powerful tool in targeting. However, because clickstream data is typically gathered by using (third-party) tracking cookies placed on various websites, a company that wishes to gather meaningful amounts of clickstream data must be able to place advertisements on a large number of different websites. Realistically this means that most companies wouldn't be able to collect meaningful clickstream data on their own. Instead, large advertising networks (e.g. Google & Baidu), benefit greatly from their massive reach and are able to sell clickstream data-based targeting services (Gomer & al., 2013).

Retargeting is a special use case of behavioral targeting and clickstream data. It means using cookies embedded onto your own online store to show web users advertisements of the products they browsed but did not buy on your site, when they are later visiting other websites. Lambrecht and Tucker (2013) find retargeting to be especially effective when a consumer has both viewed a product on your site earlier and subsequently searched for more information about the product from external sources, such as reviews. This means using retargeting together with clickstream data gathered from other websites can be an effective way to increase conversion rates. It makes sense for companies to target consumers who have already shown interest in their products. Attracting new customers usually costs significantly more than retaining existing customers (Wertz, 2018). The potential downside of retargeting is that it can feel very intrusive for the consumers. If one is shown advertisements for products they have just been browsing, it can feel like they have no online privacy. This can lead to negative feelings towards the advertised product and the advertiser and might also negatively affect future sales.

While many of the data collection methods presented here are technically possible to use without ad networks, they benefit greatly from the extra reach of the networks. Conversely, purchase data gives companies a good opportunity to target advertising without necessarily buying services from ad agencies or ad networks. This is if the targeting is done within their own website based on data collected on the site. Using purchase data means analyzing the shopping baskets or earlier purchases of a web users and then recommending related products to them based on that (Tupamäki, 2017), making it a form of behavioral targeting. If done on a closed platform, such as a company's web store, targeting based on purchase data can be perceived as being less intrusive, as data doesn't need to be shared with third parties. Additionally, it can improve the shopping experience, if it helps consumers discover useful products. The downside of course is that this method cannot be used to bring in new customers, but only to get existing customers to buy more.

Demographical targeting can be based on for example profile data, which has the potential to offer more thorough information on web users than the previously mentioned data collection methods. The method can be used on any website where users create profiles that contain personal information. The most obvious example is social media sites, where users willingly input various personal details, such as hobbies and interests. Despite users being free to choose how much they are willing to share about themselves, privacy concerns are still a big factor for profile data usage, as can be seen from Facebook's recent scandals.

Utilizing profile data also allows advertising platforms to do what cookies struggle with – follow users between different devices. Understanding how ads seen on one device affect browsing behavior later on another device can be very useful in determining the effect advertising campaigns have. For online purchase decision tracking, the benefits are obvious, but some advertising networks, such as Facebook, make it possible for companies to connect in-store purchases to their online ad campaigns by using profile data (Facebook Ads Help Center (1), n.d.). On Facebook, this can be done by uploading hashed data to Facebook's services. This data, which could be for example collected through a membership card system, could include for example email addresses and phone numbers of people who have made in-store purchases. Facebook's algorithms then compare the uploaded data with Facebook's own databases to match purchases with known Facebook users. These matches can then be used to determine how Facebook advertising campaigns have affected purchases at physical stores (Facebook Ads Help Center (1), n.d.).

Location data is one of the most simple but useful tools for targeting. Typically, the general location of a consumer is very easy to determine based on their IP address. While this is already of great importance on desktops and laptops, the role of exact location data increases significantly on mobile devices. Smart phones can potentially be used for very timely advertising, for example, a consumer might receive notifications or text message advertisements from a company whose brick and mortar store they are about to walk by, if they have given the application access to location data on their phone (Andrews et al., 2015). This kind of advertising is a double-edged sword, as some consumers are more than happy to find coupons or deals at their fingertips while walking past a store, whereas others will see this as an infringement of their privacy. On mobile devices, location-based targeting can also be used together with temporal data, to ensure that ads are only shown when the nearby stores are open for example (Tupamäki, 2017).

Contextual advertising differs from all of the other methods, in that it allows targeting to be done without any user data. By analyzing keywords and phrases found on a website, advertisements can be targeted based on the content of the website, rather than consumer data. This method works particularly well in situations where the advertisements are shown on niche websites, like certain forums, or on a platform like YouTube, where advertisements can be contextually related to a video's topic (Baldassarre, 2015). The main benefit of this method is the lack of privacy concerns. However, the drawback is that it can result in poor targeting in some cases. An example could be having a pre-roll advertisement for a car company on a video that is about cars that happens to feature a celebrity. Many of the viewers might not necessarily be interested in cars at all and are instead watching the video only because of the visiting celebrity.

Using search data allows companies to advertise directly to web users who use specific terms when searching for information on a website. This is a combination of behavioral targeting and contextual targeting, as it can take into account consumer input and other information on the viewed website. Like clickstream data, search data is typically used by search engine companies, who tend to be large players in the online ad industry. However, to a smaller scale, it can be used on forums and company websites independently. Search data is often paired up with location data, for example provided by web users IP addresses to improve targeting. This makes search data targeting a very useful tool for smaller companies, who are only interested in attracting local consumers.

In conclusion, while there are many different data sources that can be used for targeting, they mostly seem to have at least one of two main disadvantages. Either they

can't be used effectively without buying them as services from ad networks or they can result in online privacy concerns and subsequently damage the brand. Now that we have discussed how to overcome the first two challenges listed in Section 2.1, we can turn our focus towards the third challenge, performance measurement, which is a major part of the main research question.

2.3 Measuring and Evaluating Marketing Campaigns

In this section, we go through the building blocks of measuring the performance of online advertising campaigns. Web analytics is discussed first, as it creates the basis for performance data collection. Once it is addressed, focus turns to important metrics and key performance indicators (KPI) that are calculated based on the data. Lastly in Section 2.3.3, incremental impact experiments are presented as a way to use the metrics and KPIs in a meaningful manner to determine the effect targeting has on online ad campaigns.

2.3.1 Web Analytics

Defined as utilizing web traffic data for the purpose of improving one's website and understanding how people find and use the site (Waisberg & Kaushik, 2009), web analytics play a major role in the evaluation of the success of online marketing campaigns and targeting. Data collection for web analytics can be done in numerous ways, the most common of which are page tagging and logfile analysis (Clifton, 2010). The main difference between the two is that page tags gather information from web users' browsers, sending data to external servers, whereas web logs collect data within the website browsed, which is then stored on the website owner's servers (Waisberg & Kaushik, 2009; Clifton, 2010). This means page tagging doesn't require as much in-house analytics expertise as the logfile analysis method, but instead comes with privacy concerns, as data is collected by a third-party.

As many companies are not prepared to complete data-intensive tasks in-house, third-party web analytics tools, which often use the page tagging method, are the typical solution. One of the most popular ones is Google's freemium software Google Analytics. While it appears to generally be regarded as reasonably accurate, some concerns exist especially about the way data samples are used to estimate some KPIs (Chase, 2013; Bowman, 2018). For companies using the free version, Google uses samples to generate

insights, when the amount of data reaches a certain threshold to limit the amount of computing power needed. While the threshold, 500 000 sessions (Google Analytics Help, n.d.) in a selected data range, is more than enough for smaller companies, large companies can end up with very inaccurate analytics results, if the sample represents only a small portion of their total website traffic – unless they are willing to pay a hefty $150 000 yearly fee for the non-sampling premium version (Bowman, 2018).

While numerous third-party tools exist, some of which might be more accurate than Google Analytics, web analytics can never really be accurate due to multiple reasons. Most of the reasons relate to the difficulty of tracking web users. For example, cookie-based page tagging doesn't enable tracking an individual's browsing habits across multiple devices. This is because cookies are stored on a device and are not necessarily tied to a user, but instead the device. In a targeting context, this means that advertisers will not be able to get accurate data in a situation, where a consumer has seen a targeted advertisement while browsing on their laptop and then later uses their desktop to make a purchase at the advertiser's site. In this case, while targeting is successful, the page tagging data can't be used to determine that the purchase has been made due to successful targeted marketing. Instead, the data would suggest that an ad was shown, but did not lead to a purchase, making matters worse.

Another issue caused by the same basic problem is that page tagging doesn't allow trackers to recognize when the same device is used by multiple users. While this issue causes difficulties for web analytics accuracy, it is even more critical in targeting, if it is based on cookies (clickstream data). This is because it can lead to situations where consumers are shown ads that have been targeted based on someone else's browsing habits and interests, essentially resulting in wasted advertising space and wasted money.

Other things causing tracking difficulties are for example spam bots and web users deleting or not accepting cookies. While more advanced web analytics solutions might recognize traffic caused by bots automatically, Google Analytics for example doesn't, at least not in all cases. Instead website owners have to manually filter out spam to create more accurate insights, which can be difficult and time-consuming (Zheng, 2018). When it comes to users deleting cookies, or not accepting them in the first place, tracking them becomes much more difficult. While some data can be collected, for example with the logfile analysis method (Clifton, 2010), things like tracking the same consumer over multiple page visits becomes impossible, which reduces the accuracy of the analytics (Zheng, 2018).

Finally, as the Google Analytics data sampling example suggests, getting real-time analytics would require a lot of processing power, even if the data itself was accurate. For example, on Google's advertising platform, Google Ads, statistics can be delayed anywhere from a few hours to even two weeks, depending on the metric (Google Ads Help, n.d.). This can make it difficult to accurately measure how marketing campaigns affect a website's traffic and sales.

2.3.2 Metrics and Key Performance Indicators

The accountability of marketing has long been one of the main challenges for marketers (Patterson, 2007). As targeting potentially leads to more efficient marketing, measuring the impact targeting has on a campaign is bound to be tricky if the performance of the campaign is difficult to evaluate on its own. In this section, we first discuss the best practices regarding marketing performance measurement. Then, we consider the best ways to measure the performance of targeting efforts in those marketing campaigns.

At the turn of the millennium, most companies were paying for marketing services (e.g. provided by advertising agencies) in ways that were not tied to the results of the campaigns (Ambler, 2000). Ambler (2000) expressed a need for marketers to be held more accountable for their efforts and predicted that in the future pricing would be more performance-driven. At least for online ad targeting, this is now often commonplace. For example, on YouTube, advertisers have the option of only being charged when consumers engage with their targeted ads. For pre-roll video advertisements 'engaging' means watching an ad for at least 30 seconds (or until it ends if shorter than 30 seconds), rather than skipping it and for banner ads it means clicking on the ad to visit the advertiser's website. In this case, it is in the service provider Google's best interest to make sure that targeting is optimized, to maximize the revenue made from the ad spots. As a side note, traditional cost per mille-priced advertising is still also an option on YouTube for those advertisers that want more certainty or are only interested in a large reach and are not too interested in targeting.

Despite this kind of progress, companies still struggle to understand the causal connection between marketing campaigns and their effects, as is evidenced by P&G advertising cuts. To mitigate this challenge, performance measurement and monitoring must already be taken into account when forming the campaign strategy (Dodson, 2016). By doing this, the metrics can be chosen based on the objective of the campaign. Selecting

the right metrics and KPIs allows marketers to connect their efforts with the intended business outcomes (Patterson, 2007). Patterson (2007) recognizes that some commonly used marketing metrics, for example click-through rate, are not necessarily outcome-based, but still serve an important role in the overall understanding of campaign performance. Hence, Patterson suggests that marketing metrics can be placed on a 5-point continuum based on their purpose.

The continuum begins with simple activity-based metrics. The purpose of such metrics is to calculate how many times something has occurred (Patterson, 2007). Click-through rate (CTR), which tracks the proportion of people who click an advertisement when they see it, is a good example of an activity-based metric. While it can provide some insight on the campaign, for example how appealing an ad is to consumers, it does not tell anything about the outcome of the campaign if the objective is to increase sales. However, it can be argued that if a campaign's main objective is to attract more visitors to a website, then CTR becomes more of an outcome-based metric.

This brings forth one caution about Patterson's metrics continuum. It is based on the notion that the main responsibility of marketing is to "achieve profitable revenue growth … by finding, keeping and growing the value of profitable customers" (Patterson, 2007). Even though this notion might very well be accurate, it means that the continuum doesn't take into account that sometimes marketing has different objectives. This does not take anything away from the usability of the continuum and the importance of its main message of understanding to measure based on the desired outcome, but it does emphasize the importance of understanding that each campaign is a unique situation and the continuum has to be considered from an appropriate perspective. There is no unambiguous way of pointing out the best metrics that would be useful in all situations.

The second step of the continuum includes operational metrics. They are used to track the efficiency of a campaign. A good example of an operational metric is the campaign return on investment (ROI). It allows marketers to justify their actions, by showing that marketing efforts resulted in more income than expenditure. However, the focus is on measuring the campaign's efficiency and not the business outcome, which still might not convince top management of its significance. (Patterson, 2007)

Outcome-based metrics are in the middle of the continuum. Customer lifetime value (CLV) is one commonly used outcome-based metrics (Patterson, 2007). Measuring CLV allows companies to identify and focus on the most profitable consumers. This is a clear strategic asset for top level management and can be used for long-term success.

Leading-indicator metrics are even more advanced and come after outcome-based metrics on the continuum. They not only measure whether an objective is reached but can also be used for determining strategic direction. A leading-indicator metric might be used to monitor how one's market share is changing in relation to competitors due to marketing campaigns (VisionEdge Marketing, n.d.). If campaigns are found to have a positive effect, this can provide insight on what the market wants, giving companies an idea regarding what direction it should take for future product development. The most advanced metrics suggested in the continuum are predictive by nature. The proposition is that once enough data has been gathered with other metrics, it can be used to create statistical models that are able to predict how campaigns will affect business outcomes (VisionEdge Marketing, n.d.).

Stewart (2009) proposes another key framework for marketing accountability. When tracking the performance of marketing efforts, advertisers need to realize that the effects of marketing campaigns are spread out over a longer timeframe. Short-term effects include for example incremental sales, as advertisements trigger people's buying decisions. Brand equity on the other hand, is an example of a longer term, more persistent effect, as people might base their purchasing decisions on the familiarity of a brand that they have seen advertisements for in the past. Stewart (2009) also recognizes a third type of effect marketing can have, which is referred to as real options. The term, which is more often used in finance, proposes that marketing can create potentially valuable opportunities in the future (Stewart, 2009). In the context of targeting, one real option would be to use the retargeting method, where data collection can help companies show advertisements to web users who have previously purchased or browsed certain items on their web stores. In this case, the web store is considered a marketing channel, which provides the company valuable data, which can be used to increase future sales when it is used for retargeting ads.

Regarding the marketing metrics continuum presented by Patterson (2007) and the temporal division of campaign effects suggested by Stewart (2009), some parallels can be found. The short-term effects described by Stewart are less likely to be connected to business outcomes and are therefore better to measure with simple activity-based and operational metrics. The long-term effects on the other hand resemble business outcomes, for example in the case of customer lifetime value. Lastly, Stewart's real options are reminiscent of the future business development choices that the continuum's leading-indicator metrics can suggest.

So far, the discussion about metrics has been focused on marketing in general. Shifting our focus towards targeting, we should begin by understanding why targeting is done. If the main objective of marketing is revenue growth, the main responsibility of targeting is to make marketing more effective. As proposed in the marketing metrics continuum (Patterson, 2007), the second tier's operational metrics are key in measuring the effectiveness of marketing campaigns. If the desired outcome of online ad targeting is to increase advertising effectiveness, the right metrics for determining its performance are the operational metrics. However, as noted earlier about marketing in general, it must also be taken into account that sometimes, targeting is used to pursue other goals. For example, as discussed earlier, retargeting is used to re-engage with customers who have previously browsed certain products. Rather than making existing campaigns more efficient, it is more like a unique advertising method. Hence, it must again be emphasized that each campaign is a unique entity and measurements have to be adjusted accordingly.

Measuring the performance of individual campaigns would be somewhat useless if the results could not be compared to any benchmarks. In marketing and especially targeting, campaigns are often optimized by comparing them to other campaigns. In the next section, we discuss common methods used for comparing campaigns.

2.3.3   Incremental Impact Experiments

While it would seem like an obvious way to measure the effectiveness of different marketing techniques, measuring key metrics between two campaigns is not always reliable. There are many uncontrollable variables at play, which affect how well any given campaign does. For example, as many ad spots are now bought via bidding systems, the prices of ad spots fluctuate, as bids placed by other companies can vary a lot over time. If others bid higher than you, your ads might not be shown at the most prolific time or position on a website, which can naturally affect the outcome of your campaign. In this section, we discuss a few methods that are used to alleviate some of the differences between campaigns, to make them more comparable and therefore more helpful in determining what kind of targeting techniques provide the best results.

2.3.3.1 Public Service Announcement Ad Campaigns

One popular method to measure the effectiveness of advertising campaigns is to run two campaigns at once. One campaign shows advertisements for your products and is shown to

the treatment group in your target audience, the other shows an unrelated public service announcement (PSA) advertisement, for example an ad for a charity, to the control group. With the aid of the ad network, companies can then measure how seeing the advertisement affected the behavior of the treatment group, by comparing it to actions made by control group members. The obvious drawback to this method is the cost. Running two campaigns at once, one of which is specifically designed not to attract customers, could be considered a waste of money. On top of the extra cost, the PSA method can produce very unreliable results. (Johnson et al., 2017)

While this method works relatively well with traditional, non-targeted online advertising, the results get flawed when it is used on current ad networks. This is because ad networks try to optimize their income by targeting ads. Traditionally ad spots were sold on a cost per mille (CPM) basis, which means that the ad publisher got paid regardless of how consumers reacted to the ads. Now, ad spots are more often sold based on cost per click (CPC) or cost per action (CPA), which forces the ad networks to take more responsibility of the success of targeting, as their income is dependent on its success.

In the case of PSA ad campaigns, this optimization results in ad networks showing the PSA ads to very different consumers than those who see the "regular" advertisement (Johnson et al., 2017). To maximize the rate of clicks or actions made by consumers, ad networks use data to figure out who in the target group are more likely to click on ads for charities and who are more likely to take action when seeing an ad from the advertising company. This means that the treatment group and control group are not randomized, which can skew the results drastically. Another similar problem with comparing a company's regular ads to PSAs, is that in a normal situation, consumers would see other ads, possibly from the company's competitors, instead of the PSAs. If consumers were instead exposed to competitors' ads, it might affect their purchasing decisions and change the results.

2.3.3.2 Intent-to-Treat Advertising

Some ad networks offer an ad effectiveness measurement service called Intent-to-Treat, which improves on the PSA by making the treatment and controls groups more comparable. It is a relatively simple concept, as it doesn't take into account exposure information. This means that unlike with PSA advertising, the control group does not have to be tracked as intricately. With the PSA method, the control group is divided into those

who would have (i.e. saw the PSA) and those who wouldn't have (i.e. didn't see the PSA) seen the real ad. With Intent-to-Treat, the target audience is simply divided into those who will be show your ad (treatment group) and those who will be not (control group). Then the actions taken by consumers in each group are compared to see the effect of your campaign. (Johnson et al., 2017)

The main benefit of this method is that the treatment and control groups can be randomized even on platforms that use targeting algorithms, which makes the results more reliable on modern ad networks. Furthermore, as the PSA is no longer needed for comparison, companies only have to pay for one ad campaign, allowing them to increase the size of the campaign or save money. However, there is also a drawback. Because exposure information is excluded, and the experiment results therefore include those consumers who were not reached (or would not have been reached) by the campaign, the results are less accurate.

2.3.3.3 Ghost Ads

Arguably, the most accurate way to conduct a controlled marketing experiment is to use ghost ads in what is sometimes referred to as a Conversion Lift Test. As with the previous methods, the advertising company (Company A) splits its target audience into the treatment group and the control group. Company A's ads are shown normally to the treatment group. For the control group, the advertising platform runs simulations to determine which consumers would have been showed company A's ad. Instead of seeing the ad, they are instead shown an advertisement for the next highest bidder, to simulate a situation where company A wasn't actually making a bid for the ad spot. The platform then creates a ghost ad impression log, which allows the advertiser and the platform to analyze how these consumers' actions differ from those who have seen company A's ad.

Because the ghost ads are dependent on the optimization and bidding system of an advertisement platform, as the platform's algorithms need to recognize the situations where company A's ad would be the optimal one to be shown, this method is an excellent alternative to the PSA method on CPC and CPA based advertising platforms. As with Intent-to-Treat advertising, using ghost ads also means that there is no need to run two simultaneous campaigns, which is required with PSAs. On the other hand, ghost ad experiments, which allow the inclusion of exposure information, have been shown to produce more accurate results than Intent-to-Treat experiments (Johnson & al., 2017).

However, due to the complexity of setting up the ghost ad service and updating it with the continually-changing optimization algorithms, very few ad networks offer it as a service. While some marketing professionals state that Google and Facebook are the only ones to provide ghost ad experiments (O'Neill, 2019), Facebook's own Ads Help Center suggests that they are using the Intent-to-Treat method instead (Facebook Ads Help Center (2), n.d.). Google on the other hand, which has been a trailblazer in this technique, are not offering ghost ad experiments to everyone, as they are only available for larger companies and their large campaigns (Morris, 2017). In a smaller market like Finland, this might mean that most companies don't have the option to conduct ghost ad experiments.

While ghost ad experiments are far more accurate than Intent-to-Treat and especially PSA testing in current ad networks, the latter have one key advantage. The PSA method can be used without the assistance of the advertising platform, as the two campaigns are run separately, with both providing measurement data to the advertiser. An Intent-to-Treat experiment also potentially allows advertisers to divide their target audience into the treatment and control groups and then compare the actions taken by the consumers in each group, without needing help from the ad platform. However, on some platforms, for example on Facebook, the randomization between the treatment and control groups can be done by the platform owner (Facebook Ads Help Center (2), n.d.). Ghost ads on the other hand, require active participation from the ad network, as the service provider is the only one that has access to the information needed to create the ghost ad impressions log. This means advertisers are dependent on the willingness, and the ability, of advertising platforms to help them run ghost ad experiments. This same reason also raises a familiar worry: as the results are provided exclusively by the company offering the service, can they be validated independently?

2.3.3.4 Split Testing

The methods discussed so far are useful for figuring out the incremental impact of advertising and targeting, as they compare the differences between the behavior of consumers who have seen and those who have not seen advertisements. Split testing on the other hand, is more useful in situations where advertisers are trying to finetune their advertisements and targeting criteria, to find the most effective solutions. Similarly to the other methods, the idea is based on A/B-testing.

The concept is to run two very similar campaigns, ideally only differing in one aspect. This allows the advertiser to see how this aspect affects the campaign's success. To get the most accurate results, the advertising platform has to be capable of splitting your target audience randomly and in a way that ensure minimal overlap between the groups. Randomization reduces the likelihood of skewed results. Making sure each consumer sees only one version of the ad ensures that their reactions are influenced by only that ad, making the results more accurate.

While this can be a great tool for optimizing marketing strategies, changing only one small variable at a time can be both costly and time-consuming. Furthermore, as consumers' reactions to advertising are constantly changing due to a number of external factors, the results should always be taken with a grain of salt and split testing should be done occasionally even after finding the most effective method. Even though one type of advertising campaign might lead to seemingly optimal results when advertising to your whole target audience, split testing can also be used to figure out if different types of campaigns are more successful with certain niches in your target audience. For example, for a grocery store, an ad containing current discounts might be optimal if targeting is done very loosely. But if the grocery store is targeting towards a niche group of consumers who are interested in organic foods, showcasing your selection of organic products to those consumers might lead to better results in that segment.

Social media is an excellent platform to conduct split testing, as the abundance of profile data allows for intricate changes in the target groups. Also, as social media platforms don't have to rely on cookie-based user recognition, the results can be tracked reliably even between devices.

For digital marketing to be financially justifiable, its effects must be measurable. In many situations marketing experts are not equipped with the required know-how in data handling and statistics. While universities and other education facilities are churning out increasingly data-savvy marketers, the current situation calls for more cross-functional collaboration in companies, as data scientists and finance experts can bring valuable insight and skills to marketing efforts.

2.4 Regulations

To avoid misuse of personal data, regulations have been put in place in many countries to protect the privacy of web users. For example, in the European Union, the General Data

Protection Regulation limits the collection of clickstream data, which has been one of the main methods of tracking users browsing habits – a critical element of successful targeting.

2.4.1 General Data Protection Regulation (GDPR)

Per GDPR, all websites must ask individuals in the EU for their permission for collecting and processing data in a transparent and clearly explained manner. When requested to do so, website owners are also required to present individuals records of their data. The same applies for deletion of data. The purpose of the regulation is to give web users more control over their online privacy and to create a unified set of rules for international businesses. (European Commission, n.d.)

While the regulation of data collection is essential to avoid misuse of data and issues like data breaches, many companies struggled to meet the new requirements (Mikkelsen et al., 2018), despite the two-year preparation period between the regulation's adoption on April 27, 2016 and its implementation on May 25, 2018. For example, multiple U.S. companies, whose main audience is outside of the EU, initially blocked European web users from their sites to avoid having to comply with the regulations (Lanxon, 2018). Although stricter data regulations can be difficult and costly to comply with, companies also stand to benefit from them in the long term, as avoiding data breaches and increasing the transparency of data processing are keys for maintaining a good brand image amongst increasingly tech-savvy consumers and investors (Bradford, 2019).

Many European consumers are surely willing to agree to data collection on certain websites, as it can be used for more than just ad targeting, to improve the overall usability of a website, for example in the case of using purchase data to improve the shopping experience, discussed in Section 2.2.2. Despite this, some researchers expect the GDPR, and subsequent regulations all around the world, to have a significant impact on the online advertising industry. Ghosh (2018) suggests that contextual targeting will become more prevalent, as it gives companies the chance to target advertisements without needing any user data.

From the perspective of the online advertising industry, when the GDPR was originally implemented in May 2018, it had an immediate and significant negative impact on the digital advertising marketplace. This is mostly due to companies not being ready for the new rules and either blocking European users or removing targeted advertisements

from their sites (Davies, 2018). While the situation has normalized over time, research by Reuters Institute found that the number of third-party cookies on European news sites declined by 22%, which suggests that website owners are pickier when it comes to partnering up with data collection companies (i.e. ad networks) (Benes (1), 2018). Instead of dealing with ad networks, advertisers are now more interested in private marketplace advertising, which allows advertisers and publishers to connect directly. This is despite it possible leading to higher prices for advertising, as it allows them to advertise directly without using services with questionable data sources in between (Benes (2), 2018). This could suggest that as Ghosh (2018) expected, advertisers are now looking to target contextually, by buying ad spots on websites that contain content that is somehow linked to their products or services.

Because GDPR essentially prohibits opt-out marketing, as web users must now be asked opt-in permission for data collection, email marketing became much less utilizable. Pre-GDPR, many web stores would start sending weekly newsletters and marketing emails to their customers without necessarily asking for permission, but instead by including a link in the email to opt-out from the newsletters. Per GDPR, these companies must now have their customers' consent before sending these emails (MacDonald, 2019). To get this consent, companies need to offer better incentives for consumers, for example in the form of coupons. Email marketing tends to be a powerful targeting tool, especially when used together with for example purchase data, as it can be used to attract existing customers back to your website.

2.4.2   Regulations in Other Regions

The GDPR has quickly become a global benchmark for data privacy and numerous multinational companies and countries alike are trying to adapt their ways to comply with it (Benady, 2018). Many companies that sell products to a global consumer base have decided to apply the new EU regulation to their worldwide business operations, as countries all around the world, inspired by the GDPR, are in the process of updating their privacy laws. The idea is to reduce the risk of losing out on business opportunities due to incompliant regulations.

For example, Japan has already pushed through new regulations that are comparable to the GDPR. In January 2019, the European Commission decided to add Japan to a list of countries it deems to have adequate data privacy regulations (Wessing,

2019). This decision is significant, as it means companies have more freedom concerning data transfer between the EU and Japan, as long as they comply with the respective laws. Data transfer between the EU and countries that are not deemed to have adequate data privacy regulations requires additional safeguards, which can be expensive and difficult to manage for companies (Wessing, 2019). Therefore, in many cases, it is in the best interest of the consumers, local companies and non-EU countries to push for GDPR compliant regulations. Some countries on the other hand, especially smaller emerging markets, are choosing to wait a while before adapting their regulations, to see what the effects of the GDPR are (Benady, 2018).

The United States are a bit of an exception when it comes to data privacy laws. Each state has its own interpretation of privacy laws and it seems unlikely that many of them would be interested in complying with the GDPR any time soon (Benady, 2018). According to Benady (2018), some states have shown more interest in updating privacy laws than others. Despite the ambiguous regulations, the US are considered to have adequate data privacy laws by the EU. This is due to an agreement called EU-US Privacy Shield, which protects the data privacy of EU consumers (European Commission, 2016). It remains to be seen if this agreement will be modified in the near future due to the GDPR, but for now it lowers the pressure on individual states to tighten their regulations.

2.5 Software-Based Issues

There are two types of commonly-used software that cause problems for the online advertising industry: ad blockers and virtual private networks (VPN). Ad blockers cause problems for online advertising in general, as their goal is to stop advertisements from appearing when their users browse the web. The usage of ad blockers was thoroughly discussed in my bachelor's thesis, therefore this section will instead focus on VPNs.

VPN software are specifically tricky for ad targeting, as they can be used to hide one's physical location, which would typically be easy to estimate based on one's IP address. As location-based targeting is almost always the basis for any online advertising campaign, because companies do not want to spend money on advertising to people who do not have access to their products or services, VPNs have a significant effect on the value of online advertising. Some VPNs also provide additional encryption services, which can also limit the tracking capabilities of ad networks, making behavioral targeting much more difficult to carry out.

VPNs can be described as tunnels that create private connections between two entities that are connecting on a public network (Microsoft, 2009). They were originally developed to secure connections between remote users and corporate networks, to ensure that confidential information did not leak out to the public internet (ibid). As the technology became more common, it became a useful tool not only for companies trying to keep their data confidential, but also for consumers that want to browse the web privately. VPN services usually offer a selection of servers from multiple different countries that a user can connect to. Users can then access the internet as if they were browsing from the location of the server. Because of this feature, VPN software proved to be a valuable asset in the empirical study conducted for this paper, which will be discussed in Section 4.

The popularity of consumer VPN services has increased significantly over the last few years. They are especially common in countries where the government is heavily regulating the content available via the internet, such as China and Turkey (Go-Globe, 2018). However, they are also increasingly used in countries that do not have internet censorship issues. Go-Globe's (2018) data lists accessing restricted entertainment content, such as Netflix libraries in other countries, as the most common reason (50%) for VPN usage. While the legality of such a use-case can be questioned, as it bypasses copyright agreements, 31% of VPN users note anonymity as an important reason for using the software, which is a more reasonable concern. Go-Globe (2018) also suggests that almost 25% of internet users use VPNs monthly and the trend points upwards. It is important to note that companies like Netflix, that offer different content in different locations, based on local copyright agreements, are fighting to block the use of VPNs and equivalent software on their service, due to fear of law suits from the content owners.

Seemingly very little research has been done on the negative effect VPNs have on the online advertising industry. To understand its impact better, we can consider instead the economic impact targeting has on the industry. According to Beales (2010), about 40% of ad networks' revenues came from behavioral targeting. Therefore, if a quarter of consumers use VPNs to hide their online browsing from trackers, this likely leads to a 10% decrease in the revenue of ad networks. However, this is just scratching the surface. The biggest impact is likely caused by VPNs masking the location of consumers. Even though this does not stop advertisements from being shown to consumers, the wrong ones are shown to the wrong people. A European consumer connected to a US VPN server would see advertisements from American companies for example, which in many cases would be pointless. Because of this, marketing campaigns are much less effective, which ultimately

affects companies' willingness to pay for ad spots, decreasing the revenue of ad networks and the overall value of online advertising.

Because many consumers are using VPN software to access geo-blocked content, it is unclear how regulations will adapt and deal with VPN usage. Even if they are eventually banned, which would be difficult to do globally, the underlying issue is that consumers want more privacy whilst browsing the web. One promising solution for this is a Web Identity Translator (WIT) (Papaodyssefs et al., 2015). This kind of new service would potentially increase the privacy of consumers while still enabling behavioral targeting to some extent. The idea is that while a WIT would allow trackers, such as ad networks, to follow consumers' browsing habits, the WIT would make them unidentifiable by making slight changes to the users' browsing patterns (Papaodyssefs et al., 2015). Especially in a data leak situation, this would provide consumers and extra layer of protection, as the data would be much more difficult to link to certain individuals (ibid). It remains to be seen if a service like this would provide enough value for consumers to create interest in it.

3   Research Model and Hypothesis Development

Earlier literature provides feasible solutions for the first three challenges presented in Section 2.1. However, the fourth challenge, understanding the online advertising platforms, is something that prior literature does not solve. In this paper, the participant observation research methodology is developed to suit the online environment and its use in online advertising research is demonstrated with an empirical study. The methodology is discussed in depth in Section 4. To create an interesting use case to showcase in this paper, a research model and some hypotheses were developed.

As doing a business use case demonstration was out of the question due to lack of opportunity, this model instead focuses on showing how participant observation can be used for both market research and marketing research. The idea is to collect data from four different markets and to observe whether ads are distributed to the consumer differently based on some market variables. The two variables used to differentiate markets are the size of the market and the level of online data privacy regulation. These are both typical examples of variables used in market research, when examining differences between two or more markets.

Treating them as binary variables (large market – small market; loose regulations – strict regulation) allows the creation the following matrix (see Figure 1), which helps with choosing suitable markets to study. If differences in the way campaigns are set up in different markets are found, they would emphasize the importance of measuring the effects targeting has on campaigns. There is no one-size-fits-all solution to targeting; the best solution depends on the situation at hand and without continually measuring how well targeted campaigns do, the results will be far from optimal.
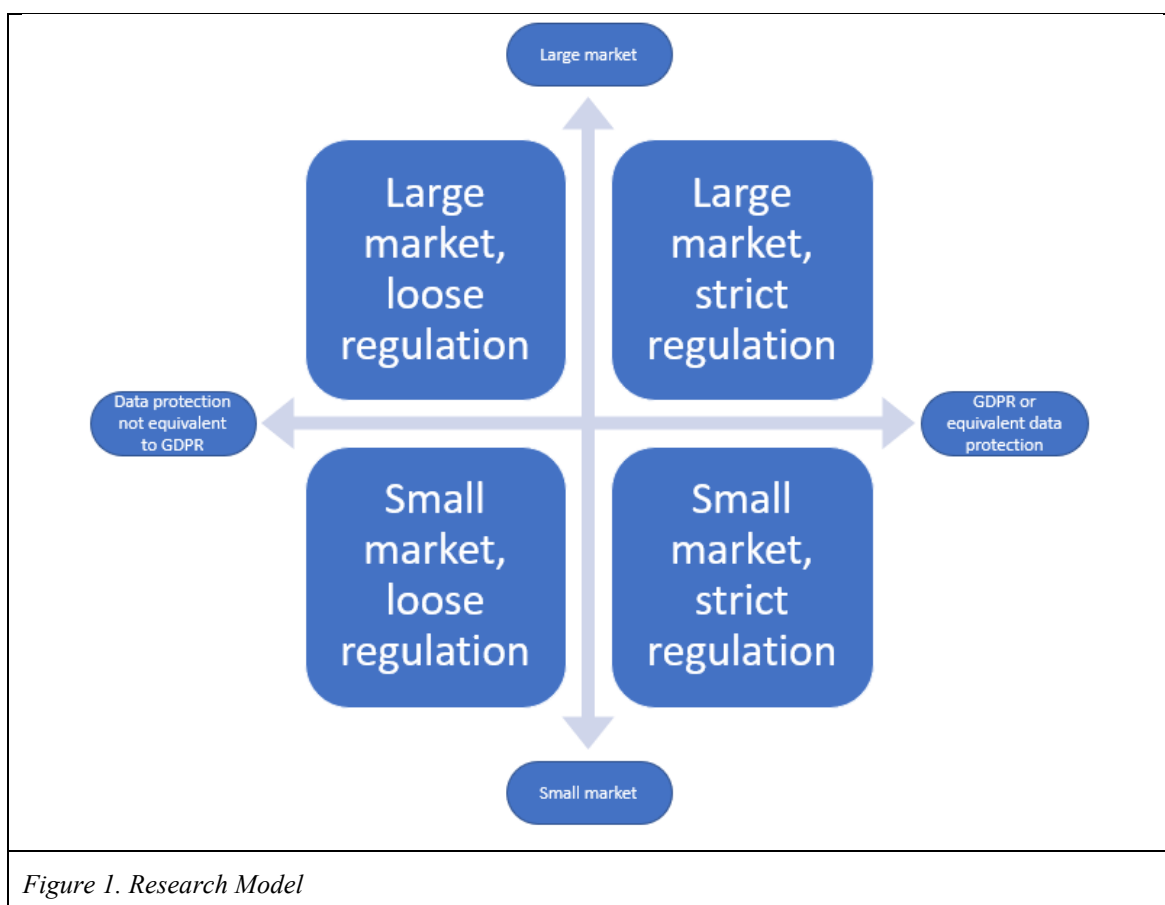
Figure 1. Research Model

Dividing markets with these two variables allows us to test two hypotheses. As pointed out in the literature review, some researchers (Ghosh, 2018) expect that strict data privacy regulations increase companies' interest in contextual advertising, to avoid infringing on consumers' privacies. This leads to Hypothesis 1:

H1: Stricter data privacy regulations result in more contextually targeted advertising

The second hypothesis proposes that in a smaller market, the amount of ad repetition (i.e. same consumer seeing the same ad multiple times) is greater than in larger markets. This is simply based on the rationale that in a smaller market, a smaller number of companies are advertising at any given moment; leading to more repetition. While ad repetition has been found to have positive effects to some extent, too much repetition can lead to a negative brand image, as consumers can find it annoying.

H2: The size of the market affects the amount of ad repetition

After collecting a sufficient amount of ad impression data on an account with personalized ads turned on and analyzing the data, it was decided that additional data would be collected on an account with no personal data and with personalized ads turned off for good measure. The idea is to see, if the repetition and contextuality of ads are different when using a Google account with ad personalization turned off. Motivation for the additional data collection comes from the low rate of contextually-relevant advertisements seen when personalization is turned on and personal experience of seeing seemingly more contextually-relevant ads when using a Google account with ad personalization turned off in the past.

As most of the advertising on YouTube is charged based on the CPA pricing model, where advertisers only pay, if consumers click or watch their advertisements long enough, the algorithms used for ad distribution should always aim to match ads with those consumers who are most likely to interact with them. When ad personalization is not an option (i.e. has been turned off by the consumer), this matching process becomes more difficult. The expectation, based on rational thinking, is that contextual ads will be more common when ad personalization is turned off, as it will be the next best thing for Google's algorithms to use when trying to maximize the platform's revenue. This formulates into Hypothesis 3a:

H_{3a}: Disabling ad personalization on a Google account, will increase the amount of contextually-relevant advertisements seen on YouTube

If this turns out to be true, then repetition might increase as well as a byproduct. For example, on a gaming video, consumers would then mostly see ads for other games (few unique gaming-related campaigns active at once), whereas with ad personalization they would see ads for both gaming (assuming it is one of their estimated interests) products and plenty of ads for other types of products as well, based on their other interests. Hypothesis 3b represents this supposition:

H_{3b}: If the amount of contextual targeting increases when ad personalization is disabled, ad repetition increases as well

To ensure that there is no personalization data available for Google (in case the settings do not do what they claim), for phase two, the test computer was formatted to erase all the data on its hard drives. Then a new Google account was created on the computer, immediately after which ad personalization was turned off and all data collection was paused in the account settings.

4 Methodology

Originally, the study plan was to interview marketing experts in large Finnish companies dealing in consumer goods. The purpose of these interviews would have been to understand how companies make decisions concerning online ad targeting. Secondly, the interviews could have been used to compare the different methods to target and to measure the effectiveness of targeting that have been uncovered in the literature review with the methods that companies in Finland are actively using. Thirdly, the interviews could have also helped identify trends regarding topical issues such as the growing emphasis on mobile marketing and the effects that stricter regulations, such as the GDPR, have had on online ad targeting.

Marketers from companies in various industries were approached via email. Unfortunately, the vast majority of the emails sent were never replied to, while the few responses received all turned down the request for an interview. Based on the negative replies, it seems that companies are not willing to discuss their targeting practices, even when anonymization of the interview data was an option. Two reasons come to mind for the somewhat surprising lack of participation interest – competitive advantage and the nature of the topic. Some companies might consider online ad targeting as a way to stand out from competitors and are not willing to discuss their successful methods, whereas others might be hesitant to discuss consumer data collection and data usage, as they are sensitive subjects that could be turned against the companies and lead to brand damage.

Realizing that interviews were not a feasible way to study this phenomenon, an alternative study plan was necessary. By applying a form of participant observation, we can examine how advertising campaigns appear to consumers. This allows us to research various different aspects of online advertising. For example, it can be used to understand the typically publicly unknown targeting algorithms that are used by ad networks to match consumers with advertisements. It can also be used to examine how changing the way an ad campaign is set up (e.g. targeting criteria and pricing) affects how it reaches a target audience. While participant observation cannot necessarily answer the questions that the interviews were supposed to, this alternative method makes it possible to research others element of marketing and targeting.

4.1 Research Methods

Participant observation can be defined as taking the role of someone within the studied environment and then examining the phenomena taking place in that environment from their perspective (Laurier, 2010). Participant observation is often used in ethnographic studies, where the researchers immerse themselves into foreign cultures and partake in their traditions and ways of life to understand cultural differences (Atkinson & Hammersley, 1998). In a business context, participant observation has perhaps most notably been used in the form of mystery shopping. In this method, researchers disguise themselves as customers while visiting physical stores and then observe the service they receive from employees (Wilson, 2001).

To suit the context of online advertising and the research model developed in Section 2.6, a slightly different approach has to be taken. To study each market, quantitative advertising data will be collected from the perspective of a consumer. As the goal is to understand trends in entire markets, rather than methods used by individual companies, collecting advertising data from the consumer side is more applicable, as consumers are exposed to ads from multiple companies. In practice, ad exposure data will be gathered by logging the ads seen on Google-owned YouTube over the duration of the experiment.

The research method used here has many similarities, but also some key differences, in comparison to mystery shopping. Both methods allow the researcher to observe a service from the consumer's point of view. They also make it possible to collect data discretely, as the service provider is unaware that they are being monitored. In mystery shopping, the service provider would be the employee dealing with the consumer/researcher and in the method used here, the service provider is Google, as the ad network. While the core concept is the same, one big difference is that mystery shopping is typically done within one's own store, whereas here we observe in an external environment that cannot be controlled as diversely.

Even though the observation environment is different, the phenomenon observed (i.e. ad distribution and targeting algorithms) is very much like any other service. For example, if mystery shopping is used to analyze the quality of service a customer gets when walking into a store, the store owners have the ability to set specifications concerning how employees should treat customers. Similarly, when setting up advertising campaigns on YouTube, companies can set certain criteria, for example they can choose

their target audience and the desired goal (e.g. lead generation, increased website traffic or increased brand awareness and reach) (Willson, 2018). However, in both mystery shopping and online ad targeting, the researcher cannot affect how the actual service is carried out. At the physical store, it is up to the employees' interpretations of the specifications set by the employer and in targeting, the service is done by an unknown algorithm that supposedly follows the set criteria. Because mystery shopping is an established method, the similarities between it and the method used for this paper support the feasibility of using such a method to observe and analyze online services, such as targeting.

Another thing worth mentioning is that participant observation generally requires an environment that is well defined. For example, in ethnographic studies a country or a specific town could be appropriate environments, and in mystery shopping the physical store provides a good study environment. For some online services, this type of "closed" environment might be difficult to define. YouTube however, as a platform, offers a relatively good basis. The targeting service is provided by the platform owner based on the platform owner's data, no external parties are involved. As discussed further later on, the amount of data available for the platform owner can also be somewhat controlled by the researcher, allowing the observation of different kinds of scenarios.

Participant observation is also typically used to gather qualitative data, for example in the form of fieldnotes (Laurier, 2010). However, in this paper, the data collected is mostly quantitative, which is better suited for testing the hypotheses presented in Section 2.6. Some qualitative notes are also taken to track unexpected events and key observations that would not be perceptible from the quantitative data alone. This can be considered triangulation, which means using multiple methodologies to study the same phenomenon (Jick, 1979). The purpose of triangulation is to increase the accuracy of one's findings, by allowing the researcher to see the phenomenon from multiple points of view. While the main focus is on quantitative data, collecting some qualitative data on the side can also deepen one's understanding of the observed matter (Jick, 1979). For example, while collecting the data to test Hypotheses 1 and 2, the observation was made that Google's algorithms appeared to favor personalized advertising over contextual advertising when user data was available. Due to this, additional data collection was necessary to see how the service changes when user data is not available. This is discussed more in Section 3.4.

To see how market size and strictness of data privacy regulations affect advertising campaigns, data needs to be collected in four different types of markets, listed in Figure 1 in Section 2.6. To gain access to four markets, a VPN software is used. It provides a cost-

efficient way to collect data in multiple markets, without having to travel to those countries. In most cases, advertising networks learn a user's location from their IP address, which is normally easily tracked by the website owner. As discussed earlier, using a VPN allows us to mask our IP address and hide our physical location, effectively making it seem like we are browsing from another location. Because the vast majority of ads on YouTube are localized, using a VPN allows us to observe and log ads from different countries.

Because VPN service providers have servers in a limited number of locations, careful consideration was necessary to find a service that would give us virtual access to four locations that suited our needs. The decision was made to use NordVPN, which is one of the largest global VPN service providers. For the markets with strict data privacy regulations, countries in the EU were selected, as the GDPR is widely recognized as world-leading, when it comes to data privacy regulations. Therefore, in the study, the large market with strict regulations is Germany and the small market with strict regulations is Finland.

When choosing markets with loose data privacy regulations, the primary focus was to avoid countries that the EU recognizes as having an adequate level of data protection. These countries are listed by the European Commission (European Commission, 2019). Based on this and the list of available servers, Australia was chosen as the large market and Bosnia & Herzegovina as the small market. While both countries are actively trying to improve their data protection regulations, they are still not considered to be near to equaling the GDPR (Davies, 2019; European Commission, 2018). Figure 2 presents the updated research model with the hypotheses and the locations used for testing.

It is important to emphasize that this research model is meant as a proof-of-concept for using participant observation in the online advertising context. The matrix is not an essential part of participant observation, but here it is used to exemplify how the method can be used in market and marketing research.
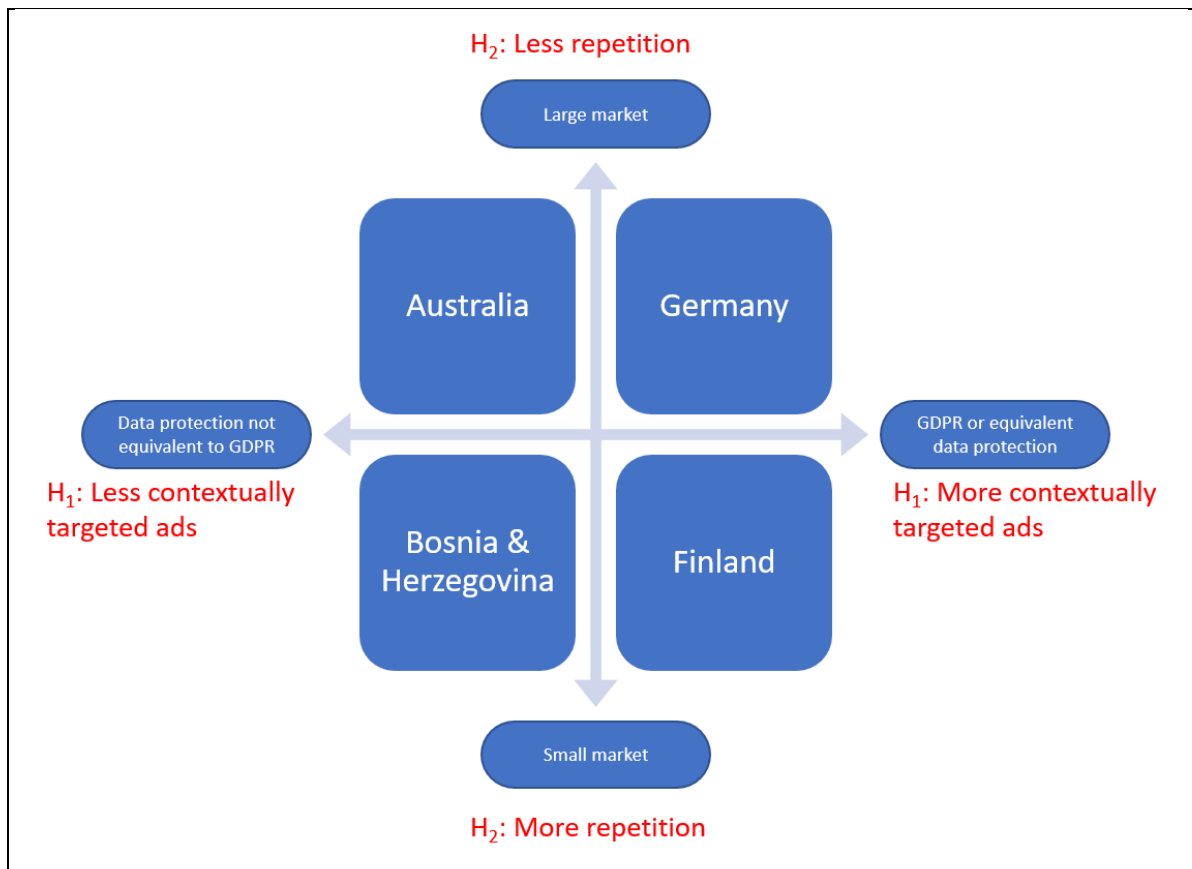
Figure 2. Research Model with Hypotheses and Researched Markets

4.2 Research Preparation

Before collecting any advertising data, sufficient testing was done to ensure that a VPN can be used to accurately and durably mask one's location from Google. Some VPN software are known to leak information from time to time, which would obviously compromise the data collection. While NordVPN is generally regarded as a secure service (Griffith, 2018), the advertisements seen were constantly monitored, to make sure the location stayed correct. On one occasion, when connecting to a supposedly German server, the virtual location ended up in Northern Africa. This anomaly was noticed immediately and data was not collected until the situation was resolved by connecting to another VPN server, which functioned correctly. Therefore, the anomaly did not affect the results.

To ensure that comparable results can be collected in all of the different markets, the data collection needs to happen in as controlled an environment as possible. For this purpose, several decisions need to be made. The first of which, is to use YouTube as the platform for data collection. This decision allows us to use the same platform for data collection in all four countries, as YouTube is one of the largest online advertising

platforms in the world, reaching a truly global audience. As the platform is owned by Google, it has both the resources and the expertise to have top-of-the-line targeting algorithms and processes in place. As Google accounts are also used as accounts on YouTube, this allows the platform to use personalization data gathered from all of Google's services to target advertising.

On the topic of accounts, the second decision is to use an existing personal account with plenty of consumer data available for Google. The same account is to be used for the experiment in each country. Choosing YouTube as the experiment platform allows taking advantage of Google's activity controls and ad personalization settings. The activity controls allow users to choose what data Google is monitoring at any given time. To keep the results comparable, all activity controls will be paused for the duration of the experiment. This means that the ad personalization (i.e. targeting) profile that Google has generated over the years of using the account will not be altered by the videos viewed during the experiment. Appendix A shows the activity controls setup on the Google account used for the experiment.

The third decision, is that the videos that will be watched to generate the ads, are all from a few selected topics. These topics are selected based on two things: the account's personalization data (available from Google's ad personalization settings) and topics that are feasible to use for contextual advertising. Based on this, the topics chosen are cars, consumer electronics, cooking, gaming, sports and TV shows & movies. Appendix B shows the personalization settings used for the experiment. Ad personalization, which can be turned off completely, is left on for the experiment, because even though personalization is not measured in this experiment, having it enabled represents a more realistic situation. It also ensures that Google's algorithms have a good chance to work as intended, when choosing which ad to show.

By limiting the video topics, it is easier to select similar content, when collecting data for each country. This is important, because a similar experiment for my bachelor's thesis suggested that contextual targeting is far more utilized by companies in certain industries compared to others (Tupamäki, 2017). Another decision that helps keep the content comparable is to only watch videos from North American content creators. By doing this, we avoid the likelihood of local channel-based targeting. When advertisers create their campaigns, they have the option to place their ads on specific content creators' channels or even specific videos. During preliminary testing for example, it was noticeable that when watching football videos on the Bild Fussball channel in Germany, an ad for a

certain beer company was always shown. By avoiding local channels, we minimize the chance that channel-based targeting could sway our results. Another way that advertisers can choose to target their campaigns is by time of day. To ensure comparability, all results are gathered during day time in each respective location.

The final preparatory decision is to log only specific advertising formats on YouTube. Existing formats include display ads (i.e. banner ads), overlay ads, three types of video ads and sponsored cards. Out of these, we exclude display ads and sponsored cards. Display ads typically correspond with either the overlay ads or video ads on videos and would therefore in many cases create a double entry for one ad impression. Sponsored cards on the other hand seem to be extremely rare, for example during the data collection for this study none were seen.

4.3 Data Collection

With the preliminary testing complete and everything working as intended, the VPN is connected to a server in one of the four countries of interest. After that, YouTube videos that fit the topics and other constraints mentioned in Section 3.2 are viewed to see what ads appear. Using Microsoft Excel, data is manually logged separately for each country for both video advertisements and overlay ads. Each advertisement seen is registered on a new row and the following pieces of data are gathered: video title, video URL, advertiser's name, localization of the advert, contextuality, repetition count, category and video tags. An extra column was used for other possible significant notes.

Of the two measurable phenomena, repetition is simpler to track than contextuality. In the repetition count column, each unique ad is numbered from 1 onwards. Every time the same ad is seen again, the number is increased by 1. When it comes to evaluating the contextuality of an advertisement, essentially all the other columns are considered. On YouTube, as throughout Google's Display Network, contextual advertising is done based on keywords and topics set by the advertiser. Google's algorithms then analyze a webpage's contents and finds suitable ad placements. This analyzable content includes at least the video title, the category (i.e. topic) of the video and the video tags, which are set by the content creator. By comparing these to both the advertiser and the advertisement, it was determined whether the advertisement was linked to the video's content contextually. Considering the contextuality of both the advertiser and the ad proved important, for example in a situation where Gigantti was advertising on a cooking video. Whereas

Gigantti is not explicitly known for cooking apparel, the ad was specifically for a Jamie Oliver branded barbecue set. This suggests that the advertisement was indeed contextually targeted. Because the contextuality is not always very evident, instead of a strict yes/no division, the contextuality was measured on three levels: Clear (1), mild (2) and none (0).

Unlike the other columns, the localization of the advert does not necessarily tie in to measuring the repetition or the contextuality of the ads. However, it is recorded in order to see if ads that have been localized (or are from local companies) appear differently compared to ads that are not localized. In this study, ads that are from local companies or the local marketing groups of international companies are marked as 'local'. Ads from international companies that are either dubbed or subtitled in the local language are marked as 'localized' and ads from international companies that are not customized to the countries in question are marked as 'non-local'.

5 Findings

In this Section, we will go over the results of the ad logging experiments. The results are presented as descriptive statistics. Statistical inference would be difficult to analyze, as the samples taken do not represent a static population. Instead the set of currently active ad campaigns changes constantly. First, we will discuss the results of phase one, which was conducted on a Google account with plenty of historical personalization data and ad personalization turned on. After this, we discuss how the situation changes when using an account with no personalized data and with ad personalization turned off for good measure.

Although it was concluded that using an account with ad personalization turned on is a more realistic scenario, as that is likely how Google's algorithms are intended to work, the results of phase one raise some questions about how the targeting process is done. In short, the results suggest that when ad personalization is available as an option, the algorithms might favor it over contextual advertising, regardless of what the advertiser has intended, when setting up the campaign.

5.1 Results of Phase One

For this phase, the intention was to collect ad impression data from each of the four countries until clear trends appeared. From very early on, it was clear that the overlay ads were not appearing frequently enough for any meaningful results. While data about them was collected throughout the study, it was not considered in the overall results. For example, in Germany, only 4 overlay ads were seen throughout the experiment. Closer inspection of Google's Ad Support explained the situation, as the way overlay ad campaigns are done has changed. Whereas they used to be created through Google Adwords, which is a widely used platform (now called Google Ads) for creating online ad campaigns, overlay ads are now only available for reservation sponsorship campaigns (Google Display Specs Help, n.d.). The difference is that video ad spots are bought via a bidding system and they are only charged for when users interact (i.e. click or watch long enough) with the ads, whereas reservation campaigns are charged based on impressions with the more traditional CPM cost structure (i.e. pay based on how many times the ad is shown, regardless of interaction). At the time of this experiment, in Q2 of 2019, advertisers seemed to be much less interested in reservation campaigns on YouTube.

Focus was therefore shifted towards only video advertisements. For Finland, Germany and Australia, 60 ad impressions were logged for each country. However, for Bosnia and Herzegovina, this proved unreasonably difficult, as YouTube appeared to run out of unique advertisements to show in the country (after about 40 impressions). Instead, only 50 ad impressions were collected for Bosnia and Herzegovina. The results were collected over a seven-day period. While gathering data over a longer period of time would mean that new campaigns would be made, and more unique advertisements would be available, seeing 60 advertisements in the span of one week is not unreasonable for a consumer. For example, on some longer (ca. 20-minute) videos, 8 advertisements were shown on a single video. When it comes to Bosnia and Herzegovina however, Google's algorithms do recognize the situation and in a smaller market with less active ad campaigns, less ads are shown to the viewers, presumably to avoid excessive repetition. Instead of showing two ads during each ad break, as in Finland, Germany and Australia, only one ad was shown at a time in Bosnia and Herzegovina.

5.1.1 Ad Repetition on YouTube

Before discussing the results, it is important to note, that when a video advertising campaign is set up on YouTube, advertisers have the option to limit how many times a specific user can see their ad. Based on the results, some advertisers ignore this option altogether, whereas others clearly use it to avoid excessive repetition, which can be annoying for consumers. It is also worthy to emphasize that the ad spots are bought via a bidding system, by determining how much you, as an advertiser, are willing to spend daily and by setting a target cost-per-action (CPA). The more an advertiser is willing to spend, the more often their ads will be shown, as Google's algorithms aim to maximize the platform's advertising revenues. In a situation where an advertiser set high spending limits and a high CPA target and also a limit for repetition, their ads will be among the first that a consumer who suits their targeting criteria sees. However, once the set repetition limit is reached, the ad will no longer be shown to that consumer.

As noted in the previous section, when using the VPN server in Bosnia and Herzegovina, unique ads quickly became harder to come by and in fact, from the 39th ad impression onwards only one ad was repeatedly shown until reaching the 50th impression. This suggests that in some small markets, the number of active campaigns at any one moment can be very small, which can lead to excessive repetition, if repetition limits are

not set when creating the campaign. The results from Bosnia and Herzegovina, a country with a population of about 3,5 million (United Nations, 2019), strongly support Hypothesis 2, which states that in a smaller market, ad repetition will be higher.

On the contrary, in Finland, the other smaller market that was researched, ad repetition was not a problem at all. Even after the 40th ad impression, multiple new unique ads still appeared. So, unlike with Bosnia and Herzegovina, in Finland, which can be considered a small market with its population of about 5,6 million (United Nations, 2019), the data does not support Hypothesis 2. Especially, when taking into account that in Australia, which is a much larger market (population: ca. 25,1 million (United Nations, 2019)) than Finland, repetition did cause some issues in this study. The biggest culprit was one ad campaign for an Australian political party rallying support for a nearby election. It is a clear case of an advertiser that had a high spending limit and a high CPA target, as out of the 60 ad impressions in Australia, the ad was shown 24 times! Even though the ad campaign goals for a political party might be very different to that of a business, it is hard to believe this level of ad repetition is either useful or intentional. Interestingly, the other large market, Germany, with its population of 82,4 million (United Nations, 2019), had very similar repetition results to Finland. Table 1 contains the repetition results for each country when all of the ad impression data is taken into account.

Table 1: Ad repetition (All data)

	Unique ads	Total ad impressions	Percentage of unique ads	Ads with 5 or more impressions
Finland	39	60	65 %	1
Australia	19	60	32 %	2
Germany	40	60	67 %	1
B&H	19	50	38 %	2

Both the percentage of unique ads and the number of unique ads that were seen five or more times suggest that ad repetition is more common in Australia and Bosnia & Herzegovina than in Finland and Germany. Hypothesis 2 is therefore not supported by these results. However, as each location, especially Australia, had one ad campaign that was clearly repeated more than the others, it is worthwhile to see how the results change, if we exclude the most repeated ad in each country. Table 2 shows the refined results. The

percentage of unique ads increases more significantly in Australia and Bosnia & Herzegovina, when the most repeated ads are excluded from the results. However, these markets still contain significantly higher ad repetition rates than Finland and Germany.

Table 2: Ad repetition (Excluding most repeated ad)

	Unique ads	Total ad impressions	Percentage of unique ads
Finland	38	54	70 %
Australia	18	36	50 %
Germany	39	50	78 %
B&H	18	37	49 %

5.1.2   Contextually-Targeted Advertising on YouTube

The premise of this experiment was to create as realistic a scenario as possible, so that the targeting mechanisms on YouTube would work as they are intended to. This is why an account with real user information was used with ad personalization turned on. The videos that were watched to generate the ad impressions were selected based on Google's estimate of the user's interests. It would have been unrealistic to suddenly jump to watching content that does not match prior browsing habits. However, it must be noted that doing this has very likely affected the results somewhat. This is because if one watches videos from a select few topics, which also happen to be their topics of interest, while having ad personalization turned on, just by chance some of the ads seen will be contextually related to the video's content, even though the advertisements have not been targeted contextually. This means that the number of contextually-relevant ads observed in this study could be slightly higher than the number of ads that were actually targeted contextually. While this is good to keep in mind when considering the results, it should not significantly affect the comparability of the results between two countries, as the effect is the same for all countries.

As noted in Section 3.3, the contextuality of the ads is rated on a three-step scale: clearly contextual, mildly contextual and non-contextual. Clearly contextual advertisements are those that are directly connected to the content of the video, the video tags and/or the video category. For example, when watching a video about computer hardware (Titled: Don't bottleneck your PC with your Monitor!), on a Finnish VPN server,

an ad from verkkokauppa.com was shown, which was specifically for AMD's computer hardware goods. This ad was clearly connected to both the video's content and the video's category (Science & Technology). Advertisements were determined to be mildly contextual if they were linked to the video category but did not necessarily have anything to do with the video's content or its tags. A good example of this is a Huawei phone advertisement that was shown on the same computer hardware video discussed above. While Huawei's phones certainly match the science and technology topic, they are less connected to the video's content on PC peripherals. Ads that could not be counted as clearly or mildly contextual were grouped up as non-contextual.

For each country, the percentage of clearly contextually-relevant ads and also the percentage of either clearly or mildly contextually-relevant ads were calculated. These numbers give us the lower and upper boundaries respectively, of a range between which the true rate of contextually targeted advertisements likely is. A range like this is helpful, since the division between clearly and mildly contextual ads cannot be perfectly objective, even if that is the intention. Table 3 presents the results for each country, when all of the ad impressions are considered.

Table 3: Contextually-relevant ads (All data)

	Clearly	Mildly	Total ad impressions	Clearly contextual	Clearly or mildly contextual
Finland	11	5	60	18,3 %	26,7 %
Australia	3	10	60	5,0 %	21,7 %
Germany	13	3	60	21,7 %	26,7 %
B&H	12	3	50	24,0 %	30,0 %

Most notably, all three European countries have a significant portion of clearly contextually-relevant advertisements, whereas the same cannot be said for Australia. On the other hand, the ads seen while browsing connected to an Australian VPN server, were mildly contextually-relevant more often than in the European countries. As with the results for ad repetition, these results are somewhat affected by the campaigns that suffered from excessive repetition. However, this time we cannot exclude the most repeated ads completely, as in some countries, notably Finland and Bosnia & Herzegovina, the most repeated ads were contextually-relevant on multiple occasions, whereas the most repeated

ads in Germany and Australia were not. Excluding those ads completely would therefore not be a fair treatment. Instead, the excessive repetition was controlled by only including the first and second, if applicable, impression of each advertisement. While this treatment might not be perfect, it reduces the effect that poorly managed ad campaigns have on the results. Table 4 shows the refined results.

Table 4: Contextually-relevant ads (Repetition controlled)

	Clearly	Mildly	Total ad impressions	Clearly contextual	Clearly or mildly contextual
Finland	11	3	51	21,6 %	27,5 %
Australia	3	6	30	10,0 %	30,0 %
Germany	12	3	48	25,0 %	31,3 %
B&H	6	2	29	20,7 %	27,6 %

Despite reducing the noise caused by excessive repetition, ads still appear to be least contextually-relevant in Australia. This supports Hypothesis 1, which suggested that contextual advertising would be used more in countries that have tighter data privacy regulations, like Finland and Germany. However, in Bosnia and Herzegovina, ads appeared to be contextually-relevant as often as in the two EU countries. Due to the smaller number of unique advertisements in Bosnia and Herzegovina, these contextuality results are arguably less accurate than in Finland and Germany. Because there are less advertisers at any given moment, the effect that one campaign has on the results intensifies. While this does not mean that the results are wrong, it does mean that they have to be taken with caution.

5.1.3 Results When Excluding Non-Local Advertisements

As mentioned in Section 3.2.2, the localization of the adverts was also logged. The purpose of this is to help determine whether ad campaigns that are local or localized to a specific country are run differently than international campaigns. Table 5 presents how many of the ads seen were local or localized in each country.

Table 5: Ad localization by country

	Local or localized ads	All ad impressions	Percentage of local or localized
Finland	52	60	86,7 %
Australia	47	60	78,3 %
Germany	55	60	91,7 %
B&H	32	50	64,0 %

An interesting thing to note is that the countries that had more ad repetition, also had less local or localized adverts. This means that Google's algorithms feed through more non-local advertisements, if there are not enough local or localized advertisements to go around. While this is not conclusive evidence, it does suggest that Australian advertisers might not be using YouTube as an advertising platform as eagerly as their Finnish and German counterparts. When determining the market size for the countries in this experiment, the main focus was on the size of the population. In hindsight, an equally important factor when it comes to online ad repetition however, is the online savviness of a market and its advertisers.

In any case, this does not stop us from comparing the total results to the results that only contain local or localized ads. Table 6 contains the repetition results of all local and localized ads and Table 7 shows how the results change when excluding the most repeated ad. In all four locations, the most repeated ad overall was local.

Table 6: Ad repetition (All local and localized ads)

	Unique ads	Total ad impressions	Percentage of unique ads	Ads with 5 or more impressions
Finland	34	52	65 %	1
Australia	11	47	23 %	2
Germany	35	55	64 %	1
B&H	10	32	31 %	2

Table 7: Ad repetition (Local and localized ads, excluding the most repeated ad)

	Unique ads	Total ad impressions	Percentage of unique ads
Finland	33	46	72 %
Australia	10	23	43 %
Germany	34	45	76 %
B&H	9	19	47 %

In comparison to the results presented in Tables 1 and 2, which include non-local ads, the local and localized ads have very similar amounts of repetition in Finland and Germany. The results in Australia and Bosnia & Herzegovina on the other hand show that a large part of the repetition observed in those countries are due to local campaigns, as the percentage of unique campaigns in those countries is lower in Tables 6 and 7 than in Table 1 and 2 respectively.

The results shown in Tables 8 and 9 are comparable to those presented in Tables 3 and 4, just with the non-local ads excluded from the data. The contextuality results for all three European countries don't seem to be affected significantly by the removal of non-local ads. For Australia on the other hand, none of the local or localized ads could be considered clearly contextually targeted. The excessively repeated political advertisement definitely skews the Australian results, but for none of the other 10 unique local or localized ads to be contextually-relevant is further evidence that Australian advertisers might not be using the YouTube advertising platform to its fullest potential.

Table 8: Contextually-relevant ads (All local and localized ads)

	Clearly	Mildly	Total ads	Clearly contextual	Clearly or mildly contextual
Finland	10	3	52	19,2 %	25,0 %
Australia	0	6	47	0,0 %	12,8 %
Germany	12	3	55	21,8 %	27,3 %
B&H	9	0	32	28,1 %	28,1 %

Table 9: Contextually-relevant ads (Local and localized ads, repetition controlled)

	Clearly	Mildly	Total ads	Clearly contextual	Clearly or mildly contextual
Finland	10	1	43	23,3 %	25,6 %
Australia	0	3	18	0,0 %	16,7 %
Germany	11	3	43	25,6 %	32,6 %
B&H	4	0	16	25,0 %	25,0 %

5.2 Results of Phase Two

For phase two, the VPN software is used to collect 40 ad impressions in Finland and in Germany on a Google account with ad personalization turned off. VPN servers in Australia and Bosnia and Herzegovina are not used for this phase, as comparing results from all four different markets is not necessary for testing Hypotheses 3a and 3b. Collecting data from two markets would help identify if the results are dependent on market conditions. Germany and Finland were logical choices for this phase, as the results from phase one were similar for both countries in both contextually-relevant ads and in repetition. Phase one also suggests that there are more advertisers currently (Q2/2019) on YouTube in Finland and Germany than in the other two countries. Collecting data from Finland and Germany therefore potentially gives us more accurate results.

According to the collected data, contextually-relevant advertisements appeared significantly more frequently when using an account with ad personalization turned off. This was the case in both Finland and Germany. Table 10 shows the contextuality results for Phase two and to limit the effect individual campaigns that appeared multiple times have on the results, Table 11 shows the results when only the first and second impression of each unique ad is taken into account.

Table 10: Contextually-relevant ads (Ad personalization turned off)

	Clearly	Mildly	Total ad impressions	Clearly contextual	Clearly or mildly contextual
Finland	13	8	40	32,5 %	52,5 %
Germany	11	3	40	27,5 %	35,0 %

Table 11: Contextually-relevant ads (Ad personalization turned off and repetition controlled)

	Clearly	Mildly	Total ad impressions	Clearly contextual	Clearly or mildly contextual
Finland	10	4	27	37,0 %	51,9 %
Germany	10	3	27	37,0 %	48,1 %

Compared to the contextuality results of Phase one, presented in Tables 3 and 4, it is clear that contextually-relevant ads are more frequent when ad personalization is turned off. These results strongly support Hypothesis 3a. The results of Phase two also support Hypothesis 3b, as ad repetition is far more common when ad personalization is turned off. Table 12 shows the ad repetition results for phase two and the results are finetuned in Table 13 in a similar way as in Table 2 for Phase one.

Table 12: Ad repetition (Ad personalization turned off)

	Unique ads	Total ad impressions	Percentage of unique ads	Ads with 4 or more impressions
Finland	18	40	45 %	3
Germany	20	40	50 %	3

Table 13: Ad repetition (Ad personalization turned off, exluding most repeated ad)

	Unique ads	Total ad impressions	Percentage
Finland	17	32	53 %
Germany	19	30	63 %

Figure 3 showcases whether the Hypotheses developed in Section 3 are supported by each country. It is clear that the first two hypothesis are not supported by all of the data, whereas Hypotheses 3a and 3b are strongly supported by the results.

	Finland	Germany	Australia	Bosnia & Herzegovina
H1	Supported	Supported	Supported	Not supported
H2	Not supported	Supported	Not supported	Supported
H3a	Supported	Supported	-	-
H3b	Supported	Supported	-	-

Figure 3. Summary of Support for the Hypotheses

6 Discussion

The results of the first phase of the quantitative research did not fully support either of the first two hypotheses. When it comes to the market size affecting the amount of ad repetition, two things can be noted. In some small markets, like Bosnia and Herzegovina in this case, ad repetition is definitely higher than in larger markets on a platform like YouTube. This happens if in a given market, there are not enough advertisers currently advertising on the platform. While the size of the market undeniably affects the number of available advertisers, the other thing that appears to matter is the digital maturity of the market's advertisers. This is why in Finland, which is also a small, but on the other hand more digitally mature market, ad repetition was not an issue in the experiment.

It is no surprise that Germany has low ad repetition results, as it is the country with the seventh most YouTube viewers in the world (Migiro, 2018). German companies have every reason to be interested in advertising on the platform, as it gives them access to a massive targetable audience. In many ways, the least expected result is that Australia had a relatively high amount of ad repetition. Even when excluding the poorly configured political campaign, ad repetition is much higher than in Finland and Germany, according to our findings. One possible explanation is the low population density of Australia. Even though the total population is almost five times that of Finland's, it is spread out over a large area. This could mean that Australia is effectively comprised of multiple smaller markets, when it comes to advertising, if companies are not operating nationwide. One marketing article does bring up another potential reason, which is the lack of competition (Parry, 2015). As an isolated market, Australian companies might face less pressure from external competitors, which has slowed down their transition into digital maturity.

Regarding the hypothesis that strict online data privacy regulations like the GDPR would entice advertisers into using more contextual targeting instead of personalized targeting, results are also inconclusive. While contextual advertising appeared to be used significantly more on YouTube in Finland and Germany than in Australia, which has less strict regulations, Bosnia and Herzegovina was right up there with the former two, despite looser regulations. One possible explanation could be that as Bosnia and Herzegovina is geographically close to the EU and is a potential candidate for EU membership, local companies are influenced by the GDPR, even if it is not yet part of local regulations. Why Australian companies are seemingly not very interested in contextual advertising, might be down to similar reasons as the high repetition.

Although it might seem counterintuitive, contextual targeting might actually lead to better results than personalized targeting on YouTube (Sweeney, 2018). Personalized targeting might allow advertisers to recognize individuals that are most likely interested in their products, but contextual advertising has an important advantage. If a consumer is shown an ad that is contextually related to the video they are about to watch, they are already likely in the mood for watching content about that kind of products. This can go a long way in making advertisements feel less annoying and more interesting.

Unlike the first two, Hypotheses 3a and 3b are strongly supported by the empirical part of this study. The rate of contextually-relevant advertising increased in both Finland and Germany, when using an account with personalized advertising turned off. This suggests that when ad personalization is not an option, Google's algorithms begin using more contextual data to improve the likelihood of consumers interacting with advertisements. While the increase in contextually-relevant advertisements is significant, at least half of the ads were still in no way related to the content viewed. This is likely due to two things. Firstly, when creating ad campaigns, advertisers have the option to choose and define the way they want targeting to be done. They can use personalized advertising, contextual advertising, both or neither. Advertisers who are more interested in a large reach and increasing consumer awareness of their products, rather than directly attracting customers, might opt to run non-targeted campaigns.

Secondly, as proposed in Hypothesis 3b, using contextual advertising can lead to increased ad repetition. The results from Phase two support this, as the rate of unique advertisements was lower in both Finland and Germany when using an account with ad personalization turned off. To avoid excessive repetition, some non-contextually-relevant ads are shown as well. If a consumer is only watching videos from one topic, for example autos & vehicles, on a Google account with ad personalization turned off, showing them only car-related ads would quickly lead to over-repetition, as there are a limited number of suitable advertisers with active campaigns at any given moment.

However, it must be noted that this kind of effect should be minimal in this study. This is because in Phase one, videos were watched from six video categories (Autos & Vehicles, Film & Animation, Entertainment, Gaming, Science & Technology and Sports), all of which were in one form or another present on the estimated list of interests (see Appendix B) on the Google account used. In Phase two, the new account used obviously did not have any interests set, but the videos watched were similarly distributed amongst those same six video categories. Therefore, if all ad campaigns used both personalized and

contextual targeting, we should have a very similar group of advertisers in both Phases. As Phase two, with personalized ads turned off, resulted in more repetition with less unique advertisements, it suggests that some advertisers are only interested in advertising on YouTube when ad personalization is an option. The ones that are still interested, are more willing to use contextual advertising as a method of targeting.

Interestingly, when collecting data for Phase two, with personalization turned off, it was noticeable that less ads were shown than when personalization is turned on. This is similar to what happened when collecting data in Phase one for Bosnia and Herzegovina, where ads are much scarcer than in the other markets, presumably because the number of advertisers is lower, and Google's algorithms recognize that it would lead to excessive repetition if ads were shown more frequently. As the amount of advertisements seems to be lower when ad personalization is turned off, even in markets like Finland and Germany, where there appear to be plenty of advertisers on the platform when ad personalization is an option, this suggests that there is an opening for companies to use more contextual advertising to easily reach consumers who have turned personalization off. Currently, it seems that few companies are interested in advertising on YouTube to consumers who have turned personalization off, at least in Finland and Germany. As a large part of the video ads are bought through a bidding system, having less demand for the ad spots likely means that they are cheaper. Even if contextual advertising would not be as effective in attracting customers as personalized advertising, if the price difference is substantial enough, it can be a useful method nonetheless.

When advertising on traditional (i.e. non-online) media (TV, newspapers, outdoor ads), an advertiser typically knows the context in which their ad will be shown, at least approximately. In online advertising this is often not the case. Due to this, and the somewhat surprising results found in this study, it is strongly recommended for advertisers to conduct research from a consumer's perspective on the platforms they intend to advertise on, as demonstrated in this paper. It is a bit like using a mystery shopper test to analyze the customer experience at a store. By taking on the role of a consumer, companies can analyze the way consumers experience their advertisements. This way, advertisers can better understand how their campaigns reach their audiences and possibly those who are not in their target audiences. It also helps advertisers understand how the platforms' advertising algorithms work, which will allow them to optimize their future campaigns.

6.1 Functionality of the Methodology

The customer perspective has long been used for studying ad effectiveness, for example in the form of surveys (Gluck, 2011). Surveys can be used to answer how well consumers recall the ads they have seen and to assess how consumers react to specific ads. However, finding suitable consumers to participate is both difficult and expensive (Gluck, 2011). While the participant observation method demonstrated in this paper cannot necessarily be used to evaluate the performance of ad campaigns, it offers companies and researchers alike a way to see the advertising process from a consumer's perspective, without needing to approach actual consumers, making the method easier to manage and likely cheaper than customers surveys. It can be used to learn about advertising platforms and the services provided on them, which is one of the key challenges that Procter & Gamble faced, when it decided to cut back on digital advertising.

This method also allows research to be done discretely. It does not require help from the ad network; in fact, the platform owner is likely unaware that research is being done. Therefore, it can be used to bring a well-needed dose of transparency to the online advertising industry. Similarly, it allows companies a way to observe how their competitors are using targeting on the platform. In doing this, companies can identify ways to differentiate their campaigns and the distribution of their campaigns and gain a competitive advantage over their competitors. For example, they can identify niches in their target audience that are not being pursued by competitors.

On the topic of niche audiences, YouTube is a very versatile platform. Conducting a study like this on YouTube allows the researcher to control quite a few variables. For example, the researcher can create Google accounts and train them to resemble various customer profiles. In this case, training the accounts means browsing websites, making Google searches and watching YouTube videos that reflect the interests the desired customer profile has. Once Google has established a list of estimated interests, which can be further modified to one's liking, the researcher can pause all of Google's data collection using the account settings options to keep the estimated interests constant throughout the duration of the research. This makes the platform very favorable for controlled experiments, such as the ones demonstrated in this paper.

Another use case for this methodology is monitoring how one's advertising campaign is being served to consumers. By collecting data with the Google accounts that resemble the target audience, advertisers can assess how changing the campaign setup

affects how their ads reach consumers. For example, the advertiser could tweak the targeting criteria or the bid they offer for advertising spots. Then, by collecting data with the methodology presented in this paper, they can determine how these changes affect how soon or perhaps how often a consumer sees an ad. Because the bids of other advertisers are unknown, this method allows one to gauge how high their bid needs to be, in order to make their advertisement appear before the ad of a competitor, for example. With a higher bid, excessive repetition becomes more of a risk however. Even if the advertiser sets a limit, for example 5 repetitions per individual consumer per day, it is likely useful to know how quickly a consumer that fits your target audience reaches the limit. With a high bid, it is possible that a consumer sees your ad five times in just 30 minutes of browsing through videos. It is then up to the advertiser to determine whether this is desirable, or if further tweaking of the campaign is preferable, to make the ad appear less frequently to avoid annoying the consumer.

One of the limitations of this methodology is that collecting data in the demonstrated manner is quite cumbersome. Manually recording ad impression data takes some time, but even worse is having to wait quite some time between seeing more ads. If one were to jump from one video to another directly after seeing an ad, as an attempt to increase the number of ads seen, the platform does not typically show you another ad. Instead, it appears that ads are only triggered if the consumer has viewed a video for some (currently unknown) amount of time. There are two probable reasons why the platform behaves like this. Firstly, it annoys consumers less if they are browsing through videos, especially on a playlist, and feel the need to skip some of the videos. Secondly, and perhaps more importantly, this limits the effect bots have on ad campaigns. View bots have been an issue on the platform for years, as some content creators are eager to increase their visibility by making it seem like their videos are more popular than they really are (Keller, 2018). If the bots are used to quickly move from one video to another, they might not trigger as many ads as consumers who watch videos from beginning to end. So, while this might be a useful function for advertisers regarding their campaign measurements, it does slow down the data collection for the methodology proposed in this paper.

Interestingly, while bots might be part of the reason data collection is slow, they could potentially be used to also significantly increase the amount of data collected. Robotic process automation (RPA), or perhaps more advanced artificial intelligence software, could be used to continuously gather some ad impression data on platforms like YouTube. RPA, which can mimic the actions made by humans controlling computers (e.g.

moving a cursor, copying and pasting), could be used to collect similar quantitative data as the kind collected for this paper. As RPA lacks cognitive abilities (Strömberg, 2018), it would not be able to analyze the data however. For example, in the case presented here, the software would not be able to recognize contextual-relevance of ads. The researcher would still have to manually go through the data, which could include the advertisers' names, video tags, video categories, video titles and possibly screenshots to determine contextuality. However, this would greatly increase the number of ad impressions that the researcher could analyze. As RPA software can run on virtual machines, one computer could be used to collect vast amounts of data. Datta et al. (2015) have developed a tool called AdFisher, which allows automated collection of text advertisements. While not applicable for YouTube's video advertisements, the tool allows researchers to collect data much quicker. In fact, Datta et al. (2015) note, that when collecting vast amounts of advertising data with an automated tool, it is important to consider the possible extra expenses the research is creating for advertisers. As text adverts with CPM pricing are usually quite cheap, it is not as big of a problem when using AdFisher. However, if a similar tool was developed for collecting data about video ads, which tend to be more expensive for advertisers, large scale data collection could create excessive costs for advertisers.

It must be noted that using bots to inflate the number of views on YouTube is against the site's terms of service (Keller, 2018). However, the terms of service only forbid using bots that send "more request messages to the YouTube servers in a given period of time than a human can reasonably produce in the same period by using a conventional on-line web browser" (YouTube, 2018). The legality of collecting ad impression data with RPA on YouTube is therefore questionable and it is strongly suggested to discuss the use of data collection bots with a legal expert in advance, especially regarding use for commercial purposes.

6.2 Limitations

The experiments are essentially snapshots, as they are conducted in a short time frame. Long term experiments could show trends more reliably. The scale of the study is also quite small, which make the results statistically less significant. Also comparing more countries would be interesting, even without any hypotheses, just to see if trends could be found based on geographical location for example.

This study is also limited to one platform. If ad impressions were collected on another platform, the results might be different. Although the intention is to test the hypotheses in a real-life situation, the experiments conducted for this paper also serve another purpose. Namely, showing one method of analyzing how a targeted ad campaign is carried out in practice, without having to trust data given by the targeting service provider. While this method is very cumbersome to do manually, with robotic process automation, or other types of bots, a lot of data could be gathered over time with minimal effort.

7 Conclusion

The online advertising industry suffers from a lack of transparency. A few large ad networks control most of the industry, acting as data aggregators and offering advertising services, including ad targeting. Individual advertisers have very little power, when it comes online advertising in general, and especially regarding the targeting process. They face numerous challenges when planning and setting up targeted advertising campaigns, the following of which are discussed in this paper:

1. Identifying the target audience
2. Choosing how to use targeting
3. Measuring the performance of targeting
4. Understanding the online advertising platforms

7.1 Research Summary

Earlier research suggests that having a comprehensive campaign strategy with well-defined objectives is key to overcome the first three challenges listed. Examining existing customers, monitoring competitors and creating desirable customer profiles are used to find suitable target audiences. To identify the optimal target audience(s), companies need to take into account what they want to accomplish with their campaigns. Targeting ads to a market niches allows companies to focus their campaigns towards those consumers who they see as potential customers or perhaps their most valuable customers. Loose targeting on the other hand retains the benefits of having a large reach.

Identifying target audiences plays a major role when choosing what kind of targeting methods one wants to use. To reach particular audiences, various data sources are often required to identify the correct consumers. In most cases, advertisers do not have to supply all of this data on their own. Instead, they buy targeting services from companies that specialize in data aggregation. In these situations, advertisers need to choose for example whether they want to use behavioral targeting, which can be used for intricate personalized advertising, or instead use less detailed methods like contextual advertising, which might appear less obtrusive for consumers.

The performance measurement of ad campaigns and especially targeting efforts, is arguably the toughest marketing challenge for companies. To justify spending thousands or

even millions of dollars annually on marketing, it must be measurable in a way which proves its impact and value. Knowing what to measure and how to gather the required data are key for advertisers.

Web analytics provide the technical capabilities of tracking how consumers react to ad campaigns. The data collected with web analytics methods, such as page tagging and logfile analysis, is then presented as metrics and KPIs. Using these methods requires a lot of technical expertise, which is why many companies buy it as a service. Some ad networks provide performance measurement services for free, as an incentive to run campaign on them. Data sampling, botting and the difficulty of following consumers' browsing between multiple devices (e.g. smartphones and computers) sometimes result in inaccurate performance results, which advertisers have to be aware of.

Previous marketing research suggests that advertising campaigns should be measured with outcome-based metrics. For marketing in general, this typically means measuring how campaigns are used to find valuable customers and to increase their value by enticing more purchases. However, the objective of targeting is not necessarily the same as for marketing. As an added service, the purpose of targeting is to make marketing campaigns more efficient. Therefore, this paper proposes that measuring the effects of targeting should focus on metrics that measure the efficiency of marketing, such as campaign ROI, as they are the outcome-based metrics in this context.

Measuring the performance of a single campaign would not tell an advertiser much about its efficiency, if there is nothing to compare the results to. Using previous campaigns' performance results as a baseline is a start, but is not particularly reliable, as campaigns can be very different from one another. Incremental impact experiments on the other hand allow advertisers to compare the performance of their campaigns to other active campaigns. Unfortunately, to use more accurate incremental impact experiments, such as ghost ads, advertisers require assistance from ad networks. Advertisers are therefore dependent on the ability and willingness of ad networks to supply them with ghost ad experiment services. This brings us to the research problem, which raises concerns about the ability of advertisers to independently evaluate the services provided to them by ad networks. The main research question specifically asks:

> *How can companies accurately, and preferably independently, measure and/or estimate the value created by targeting in their online ad campaigns?*

Regarding the performance measurements provided by ad networks, advertisers can often supplement them by using third-party campaign tracking tools, which can help increase the accuracy of the results. Some third-party tools are for example more capable of recognizing botting than the ad networks' freemium tools, such as Google Analytics. This allows advertisers to determine if they are reaching the consumers they want, or if their ads are instead shown to bots. Some of the less accurate methods for conducting incremental impact experiments (i.e. PSA campaigns and intent-to-treat experiments) are in some cases possible to be used independently. However, it is difficult to recommend using them, if the more accurate ghost ad experiments are an option.

Although comparing performance results given by the ad networks to those gathered with third-party tools gives advertisers some sense of how their advertisements reach consumers, the targeting process still suffers from a lack of transparency. It is this lack of transparency that leads to notable advertisers like P&G reducing their spending on digital advertising and to situations like the YouTube Adpocalypse. This problem is presented as one of the secondary research questions in this paper:

How can advertisers evaluate the services provided by ad networks, with the intention of making better marketing decisions?

In this paper, we develop and demonstrate the use of the participant observation methodology in the context of online advertising and targeting. By observing how ads are presented to consumers, companies can increase their understanding of advertising platforms like YouTube. This method also allows advertisers to examine how minor changes in the campaign setup affect how the advertisements reach consumers. It is important to note that this method does not necessarily give advertisers performance metrics for their campaigns. This is because the focus of the use case demonstrated in this paper is on determining how the targeting algorithms work on YouTube, rather than how consumers react to the advertisements seen. However, if the knowledge this method provides is combined with performance metrics provided by Google or third-party service providers, advertisers can further increase their understanding of how to optimize their targeted online ad campaigns. Because this method can be used independently by advertisers, without the help of ad networks or other parties, it provides at least some resolution to the secondary research question about independent campaign performance measurement.

Pressure from large advertisers and consumers' increasing concern over their online privacy are changing how online advertising and targeting are done. For example, in an effort to make YouTube easier to understand and more "brand safe" for advertisers, Google is offering advertisers in selected countries a service called Google Preferred (Think with Google, 2016). This service simplifies the process of placing ads on videos from popular content creators that have been manually verified as brand-suitable. Targeting is done very loosely, only based on video categories (i.e. contextual advertising). In some ways, it resembles a transition away from intricate personalized targeting, towards a model more familiar from television advertising. Tightening data privacy regulations are also playing a role in how consumer data is collected and used for the purposes of online ad targeting. This prompted the secondary research question:

Have tighter online privacy regulations affected how companies approach and use ad targeting?

In demonstrating the use of the participant observation methodology, to answer this question, data was collected from four markets with varying data privacy regulations. In Germany and Finland, both EU countries with relatively strict privacy regulations, contextual advertising was found to be used much more frequently that in Australia, which has less strict regulations for now. In Bosnia and Herzegovina, which also has less strict privacy regulations, contextual advertising was use about as much as in the EU countries. This can be partially explained by its close proximity to the EU however. Many companies that advertise in Bosnia and Herzegovina have to comply with EU regulations, as they also have customers from EU countries.

Earlier research on the effects of tightening privacy regulations suggests that advertising becomes much less effective on websites where contextual targeting is difficult to do, when new privacy regulations are implemented (Tucker, 2012). This means that when personalized targeting becomes more difficult to do, contextual targeting is a good way to maintain the efficiency of advertising. On a platform like YouTube, where a consumer sees an ad when viewing a specific video, there is often plenty of potential for contextual advertising, even if some content might be more suitable for it than others. Therefore, it is not surprising that in countries in and around the EU, advertisers have turned to using more contextual targeting on YouTube.

However, interestingly Goldfarb and Tucker (2010) found that contextual targeting can be used to increase the effectiveness of online advertising, as long as the ads are not

obtrusive. Video ads, the format that this study focused on, are arguably a rather obtrusive form of advertising. This is because video ads interrupt the browsing experience for the consumer. While not as bad as flashy pop-up ads, video ads are certainly more obtrusive than regular banner advertisements. Because Goldfarb and Tucker's (2010) results were based on only banner ads, it would be interesting to study if similar results would be found with video ads. On the other hand, according to Sweeney (2018), contextual advertising on YouTube is leading to better results than other targeting methods.

7.2 Research Implication and Contributions

The key methodological contribution of this research paper is the development and demonstration of using participant observation in the context of online advertising and targeting. Participant observation has been previously used for evaluating business services within in physical location with the mystery shopping methodology. A similar approach is taken here for evaluating online business services. In this paper, its usefulness is demonstrated in market research, by comparing how online marketing is carried out in different markets, and also in marketing research, by showing how it can be used to identify marketing opportunities and problems, for example caused by repetition.

Demonstrating the use of the methodology developed led to multiple theoretical implications. In addition to the effects of tightening regulations on advertisers' targeting decisions, which are already discussed in Section 6.1, important findings were made regarding ad repetition and the effect of consumer data availability on contemporary targeting algorithms.

Market size did not appear to consistently affect the amount of ad repetition, suggesting that there are other market variables that affect the amount of companies actively advertising on an online platform. In Chapter 5, the digital maturity of a market is proposed as a more significant indicator of the number of active advertisers on a platform. In each of the four markets observed, some advertisements felt excessively repeated. However, the vast majority of existing literature on ad repetition seems suggests that up to a reasonable limit, repetition is beneficial for the effectiveness of advertising. For example, Kirmani (1997) finds that consumers associate high repetition with good product quality. Moorthy and Hawkins (2005) on the other hand suggest that as long as an ad is of high enough quality, repetition will lead to a positive impact on the consumer's perceived brand image. No prior research was found on ad repetition on YouTube however. Again, it would

be interesting to study if repetition leads to as positive results with online video ads. Because video ads are more obtrusive than banner ads, having to watch the same video over and over might be perceived as more annoying than seeing the same banner ad multiple times, since banner ads are far easier to ignore.

Regarding the availability of consumer data, when ad personalization was turned on and plenty of consumer data was available on the Google account used, Google's algorithms seemed to favor personalized advertising over contextual targeting on YouTube. When ad personalization was turned off, more contextually-relevant ads were shown instead. Because Google's is attempting to maximize the ad revenue on the platform, it makes sense that its algorithms use the best data available to serve advertisements to those consumers who are most likely to watch or click on them. Earlier research on how ad networks use the data available to them are limited. One notable paper by Wills and Tatar (2012) observes how the list of interests that Google generates and makes available for consumer (see Appendix B as an example) correlates with the advertisements seen. Wills and Tatar (2012) focus on banner advertisements and find that typically when an ad seems to be targeted behaviorally, the topic is listed as an estimated interest for the user. They also find that ads can be targeted contextually, behaviorally and based on location, or by using a combination of these methods. The findings presented in this paper suggest that Google's targeting algorithms change the way targeting is done based on the data that is available. When personalized advertising is allowed, ads are more likely to be targeted behaviorally and when personalization is turned off, the platform favors contextual advertising.

Since the targeting algorithms are unknown to advertisers, this also leads to significant practical implications for advertisers. When setting up the campaigns, advertisers are usually allowed to decide what kind of targeting they want to use and who they want to target (Willson, 2018). However, before running the campaign and seeing the results, advertisers have a difficult time knowing how the decisions made actually affect the campaign in practice. Using the participant observation method demonstrated in this paper also requires running campaigns, but it can be used to better understand the effects the campaign setup decisions have on the outcome. Participant observation also allows advertisers to learn more about the platforms they are advertising on. One of the key issues noted by Procter & Gamble, when dramatically cutting their targeted online advertising, is the lack of understanding the platforms.

On a final note, to make the most of targeted online advertising, advertisers should use as many methods of making sure they are getting their money's worth as they can. Getting performance data from multiple providers is recommended whenever possible. Pushing ad networks to make the targeting process more transparent is also a key issue and needs the support of as many advertisers as possible. Similarly, advertisers should inquire about the possibility of accessing top-of-the-line incremental impact experiments like ghost ads, to ensure their campaigns are performing as expected. While individual advertisers do not have much power in terms of influencing the massive ad networks, these issues are critical for making online advertising and targeting more financially accountable.

References

Books and reports

Dodson, I. (2016). The art of digital marketing: the definitive guide to creating strategic, targeted, and measurable online campaigns. John Wiley & Sons, Hoboken. 387p.

European Commission. (2016). EU-U.S. Privacy Shield. Available at: https://ec.europa.eu/info/sites/info/files/factsheet_eu-us_privacy_shield_en.pdf [Accessed on July 22, 2019]

European Commission. (2018). Bosnia and Herzegovina 2018 Report. Available at: https://ec.europa.eu/neighbourhood-enlargement/sites/near/files/20180417-bosnia-and-herzegovina-report.pdf [Accessed on April 25, 2019]

Go-Globe. (2018). The State of VPN Usage – Statistics and Trends (Infographic). Available at: https://www.go-globe.com/blog/vpn-usage-statistics/ [Accessed on May 8, 2019]

Laurier, E. (2010). Participant Observation. in N Clifford, S French & G Valentine (eds). Key Methods in Geography. 2nd edn, SAGE Publications Ltd, London, pp. 116-130.

Statista (2018) Digital advertising spending worldwide from 2015 to 2020 (in billion U.S. dollars). Available at: https://www.statista.com/statistics/237974/online-advertising-spending-worldwide/ [Accessed on November 28, 2018]

United Nations. 2019. World Population Prospects 2017. Available at: https://population.un.org/wpp/ [Accessed on May 7, 2019]

Articles

Andrews, M., Luo, X., Fang, Z., & Ghose, A. (2015). Mobile ad effectiveness: Hyper-contextual targeting with crowdedness. Marketing Science, 35(2), 218-233.

Atkinson, P., & Hammersley, M. (1998). Ethnography and participant observation. Strategies of Qualitative Inquiry. Thousand Oaks: Sage, 248-261.

Beales, H. (2010). The value of behavioral targeting. Network Advertising Initiative, 1.

Cotter, S. (2002). Taking the measure of e-marketing success. The Journal of Business Strategy, 23(2), 30-37.

Cunningham, S. (2018). Creative destruction in the screen industries and implications for policy. Media International Australia, 169(1), 5-15.

Evans, D. S. (2009). The online advertising industry: Economics, evolution, and privacy. Journal of economic perspectives, 23(3), 37-60.

Jick, T. D. (1979). Mixing qualitative and quantitative methods: Triangulation in action. Administrative science quarterly, 24(4), 602-611.

Johnson, G. A., Lewis, R. A., & Nubbemeyer, E. I. (2017). Ghost ads: Improving the economics of measuring online ad effectiveness. Journal of Marketing Research, 54(6), 867-884.

Kirmani, A. (1997). Advertising repetition as a signal of quality: If it's advertised so much, something must be wrong. Journal of advertising, 26(3), 77-86.

Lambrecht, A., & Tucker, C. (2013). When does retargeting work? Information specificity in online advertising. Journal of Marketing Research, 50(5), 561-576.

Moorthy, S., & Hawkins, S. A. (2005). Advertising repetition and quality perception. Journal of Business Research, 58(3), 354-360.

Patterson, L. (2007). Taking on the metrics challenge. Journal of Targeting, Measurement and Analysis for Marketing, 15(4), 270-276.

Raynor, M. E. (1992). The pitfalls of niche marketing. Journal of Business Strategy, 13(2), 29-32.

Sen, R. (2005). Optimal search engine marketing strategy. International Journal of Electronic Commerce, 10(1), 9-25.

Stanford, S. (2018). YouTube and the Adpocalypse: How Have The New YouTube Advertising Friendly Guidelines Shaped Creator Participation and Audience Engagement? Thesis at Lund University.

Stewart, D. W. (2009). Marketing accountability: Linking marketing actions to financial results. Journal of Business Research, 62(6), 636-643.

Strömberg, K. (2018). Robotic Process Automation of office work: benefits, challenges and capability development – A cross-sectional field study of five large Finnish corporations. Master's Thesis at Aalto University.

Tucker, C. E. (2012). The economics of advertising and privacy. International journal of Industrial organization, 30(3), 326-329.

Tupamäki, K. (2017) Targeting of Online Advertising. Bachelor's thesis at Aalto University. Available at: https://aaltodoc.aalto.fi/handle/123456789/27250 [Accessed on November 28, 2018]

Waisberg, D., & Kaushik, A. (2009). Web Analytics 2.0: empowering customer centricity. The original Search Engine Marketing Journal, 2(1), 5-11.

Wilson, A. M. (2001). Mystery shopping: Using deception to measure service performance. Psychology & Marketing, 18(7), 721-734.

Yalçın, N., & Köse, U. (2010). What is search engine optimization: SEO?. Procedia-Social and Behavioral Sciences, 9, 487-493.

A separate part of a collection, handbook, or conference proceedings

Datta, A., Tschantz, M. C., & Datta, A. (2015). Automated experiments on ad privacy settings. Proceedings on privacy enhancing technologies, 2015(1), 92-112.

Gill, P., Erramilli, V., Chaintreau, A., Krishnamurthy, B., Papagiannaki, K., & Rodriguez, P. (2013, October). Best paper--Follow the money: understanding economics of online aggregation and advertising. In Proceedings of the 2013 conference on Internet measurement conference (pp. 141-148). ACM.

Gomer, R., Rodrigues, E. M., Milic-Frayling, N., & Schraefel, M. C. (2013). Network analysis of third party tracking: User exposure to tracking cookies through search. In Proceedings of the 2013 IEEE/WIC/ACM International Joint Conferences on Web Intelligence (WI) and Intelligent Agent Technologies (IAT)-Volume 01 (pp. 549-556). IEEE Computer Society.

Krishnamurthy, B., & Wills, C. (2009, April). Privacy diffusion on the web: a longitudinal perspective. In Proceedings of the 18th international conference on World wide web (pp. 541-550). ACM.

Papaodyssefs, F., Iordanou, C., Blackburn, J., Laoutaris, N., & Papagiannaki, K. (2015, November). Web identity translator: Behavioral advertising and identity privacy with wit. In Proceedings of the 14th ACM Workshop on Hot Topics in Networks (p. 3). ACM.

Toubiana, V., Narayanan, A., Boneh, D., Nissenbaum, H., & Barocas, S. (2010, March). Adnostic: Privacy preserving targeted advertising. In Proceedings Network and Distributed System Symposium.

Wills, C. E., & Tatar, C. (2012). Understanding what they do with what they know. In Proceedings of the 2012 ACM Workshop on Privacy in the Electronic Society (pp. 13-18). ACM.

Internet-references

Baldassarre, R. (2015). 7 Great Ways to Advertise on YouTube. Available at: https://www.searchenginejournal.com/7-great-ways-advertise-youtube/137472/#close [Accessed on March 22, 2019]

Benady, D. (2018). GDPR: Europe is taking the lead in data protection. Available at: https://www.raconteur.net/hr/gdpr-europe-lead-data-protection [Accessed on July 22, 2019]

Benes, R. (1). (2018). Five Charts: Yes, the GDPR Is Shaking Up Digital Marketing. Available at: https://www.emarketer.com/content/five-charts-the-gdpr-is-shaking-up-digital-marketing [Accessed on March 25, 2019]

Benes, R. (2). (2018). How GDPR Is Moving Ad Budgets Toward Private Marketplaces. Available at: https://www.emarketer.com/content/how-gdpr-is-moving-ad-budgets-toward-private-marketplaces [Accessed on March 25, 2019]

Bergen, M., De Vynck, G. & Palmeri, C. (2019). Nestle, Disney Pull YouTube Ads, Joining Furor Over Child Videos. Available at: https://www.bloomberg.com/news/articles/2019-02-20/disney-pulls-youtube-ads-amid-concerns-over-child-video-voyeurs [Accessed on May 20, 2019]

Bien, B. (n.d.). 3 Ways to Track Google AdWords Campaign Performance. Available at: https://imatrix.com/blog/3-ways-to-track-google-adwords-campaign-performance/ [Accessed on November 28, 2018]

Bowman, A. (2018). What is Google Analytics? The Pros, The Cons, and the Ugly Missing Link [2019] Available at: https://www.impactbnd.com/blog/google-analytics-shortcomings [Accessed on March 23, 2019]

Bradford, V. (2019). GDPR: 9 THINGS WE'VE LEARNT IN THE PAST 9 MONTHS. Available at: https://www.cognition24.com/gdpr-9-things-weve-learnt-in-the-past-9-months/ [Accessed on March 25, 2019]

Bruell, A. & Terlep, S. (2017). P&G Cuts More Than $100 Million in 'Largely Ineffective' Digital Ads. Available at: https://www.wsj.com/articles/p-g-cuts-more-than-100-million-in-largely-ineffective-digital-ads-1501191104 [Accessed on May 21, 2019]

Chase, R. (2013). Can You Trust Your Google Analytics Data? Available at: https://www.blastam.com/blog/can-you-trust-your-google-analytics-data [Accessed on March 23, 2019]

Cohn, C. (2015). Steps To Identify Your Target Market. Available at: https://www.forbes.com/sites/chuckcohn/2015/02/06/steps-to-identify-your-target-market/#7ba8e8e2229d [Accessed on May 24, 2019]

Davies, T. (2019). Australian data privacy laws "rushed". Available at: https://gdpr.report/news/2019/03/05/australian-data-privacy-laws-rushed/ [Accessed on April 25, 2019]

Dua, T. (2017). Two of the world's biggest advertisers are cutting back on their digital ad spend. Available at: https://www.businessinsider.com/two-of-the-worlds-biggest-brands-are-cutting-back-on-on-digital-ads-2017-6?r=US&IR=T [Accessed on May 20, 2019]

European Commission. (n.d.). 2018 reform of EU data protection rules. Available at: https://ec.europa.eu/commission/priorities/justice-and-fundamental-rights/data-protection/2018-reform-eu-data-protection-rules_en [Accessed on March 25, 2019]

European Commission. (2019). Adequacy decisions. Available at: https://ec.europa.eu/info/law/law-topic/data-protection/international-dimension-data-protection/adequacy-decisions_en [Accessed on April 25, 2019]

Facebook Ads Help Center (1). (n.d.). About Offline Events. Available at: https://www.facebook.com/business/help/1142103235885551 [Accessed on July 29, 2019]

Facebook Ads Help Center (2). (n.d.). About Facebook conversion lift studies. Available at: https://www.facebook.com/business/help/399737743699353 [Accessed on April 23, 2019]

Ghosh, D. (2018). How GDPR Will Transform Digital Marketing. Available at: https://hbr.org/2018/05/how-gdpr-will-transform-digital-marketing [Accessed on March 25, 2019]

Google (n.d.). Understanding how Google ads work. Available at: https://safety.google/privacy/ads-and-data/ [Accessed on November 28, 2018]

Google Analytics Help. (n.d.). About data sampling. Available at: https://support.google.com/analytics/answer/2637192?hl=en [Accessed on April 30, 2019]

Griffith, E. (2018). Is Your VPN Leaking? Available at: https://uk.pcmag.com/features/89987/is-your-vpn-leaking [Accessed on May 4, 2019]

Hale, J. L. (2019). YouTube Disables Comments On Tens Of Millions Of Videos Featuring Kids In Effort To Prevent Adpocalypse 2.0. Available at: https://www.tubefilter.com/2019/02/28/youtube-disables-comments-videos-minors-adpocalypse/ [Accessed on May 20, 2019]

Jenblat, O. (2018) Facebook Ad Targeting Is Shifting And Here's What You Can Do About It. Available at: https://www.forbes.com/sites/forbesagencycouncil/2018/10/08/facebook-ad-targeting-is-shifting-and-heres-what-you-can-do-about-it/#7a26a8e6189e [Accessed on November 28, 2018]

Keller, M. H. (2018). The Flourishing Business of Fake YouTube Views. Available at: https://www.nytimes.com/interactive/2018/08/11/technology/youtube-fake-view-sellers.html [Accessed on May 27, 2019]

Lanxon, N. (2018). Blocking 500 Million Users Is Easier Than Complying With Europe's New Rules. Available at: https://www.bloomberg.com/news/articles/2018-05-25/blocking-500-million-users-is-easier-than-complying-with-gdpr [Accessed on March 25, 2019]

MacDonald, S. (2019). GDPR for Marketing: The Definitive Guide for 2019. Available at: https://www.superoffice.com/blog/gdpr-marketing/ [Accessed on March 25, 2019]

Microsoft. (2009) Virtual Private Networking: An Overview. Available at: https://docs.microsoft.com/en-us/previous-versions/windows/it-pro/windows-2000-server/bb742566(v=technet.10) [Accessed on May 8, 2019]

Migiro, G. (2018). Which Countries Watch The Most YouTube? Available at: https://www.worldatlas.com/articles/which-countries-watch-the-most-youtube.html [Accessed on May 8, 2019]

Mikkelsen, D., Soller, H., Strandell-Jansson, M. & Wahlers, M. (2018). McKinsey & Company. GDPR compliance after May 2018: A continuing challenge. Available at: https://www.mckinsey.com/business-functions/risk/our-insights/gdpr-compliance-after-may-2018-a-continuing-challenge [Accessed on March 25, 2019]

Morris, J. (2017). Google's Ghost Ads: What They Are and How They Work. Available at: https://3qdigital.com/blog/googles-ghost-ads-work/ [Accessed on April 23, 2019]

Neff, J. (2017). P&G TELLS DIGITAL TO CLEAN UP, LAYS DOWN NEW RULES FOR AGENCIES AND AD TECH TO GET PAID. Available at: https://adage.com/article/media/p-g-s-pritchard-calls-digital-grow-up-new-rules/307742 [Accessed on May 20, 2019]

O'Neill, S. (2019). TIL: GHOST ADS. Available at: https://www.nanigans.com/2019/01/25/til-ghost-ads/ [Accessed on April 23, 2019]

Parry, L. (2015). Why The Lack Of Competition In Australia Is Basically Bad News For Everybody. Available at: https://www.bandt.com.au/marketing/why-the-lack-of-competition-in-australia-is-basically-bad-news-for-everybody [Accessed on May 8, 2019]

Porta, M. (2010). How to Define Your Target Market. Available at: https://www.inc.com/guides/2010/06/defining-your-target-market.html [Accessed on May 23, 2019]

Ruffolo, B. (2018). How to Come Up With 2019 Marketing Goals Based on Business Goals. Available at: https://www.impactbnd.com/blog/how-to-set-marketing-goals-based-on-business-goals [Accessed on May 23, 2019]

Sweeney, E. (2018). Ad targeting around content most effective at reaching YouTube audiences, study suggests. Available at: https://www.marketingdive.com/news/ad-targeting-around-content-most-effective-at-reaching-youtube-audiences-s/522601/ [Accessed on May 8, 2019]

Terlep, S. & Seetharaman, D. (2016). P&G to Scale Back Targeted Facebook Ads. Available at: https://www.wsj.com/articles/p-g-to-scale-back-targeted-facebook-ads-1470760949 [Accessed on May 21, 2019]

Think with Google. (2016). Google Preferred. Available at: https://www.thinkwithgoogle.com/products/google-preferred/ [Accessed on May 30, 2019]

VisionEdge Marketing. (n.d.). Making Your Way Along the Marketing Metrics Continuum. Available at: https://visionedgemarketing.com/images/stories/PDFs/metricscontinuum.pdf [Accessed on May 24, 2019]

Vranica, S. (2018). P&G Contends Too Much Digital Ad Spending Is a Waste. Available at: https://www.wsj.com/articles/p-g-slashed-digital-ad-spending-by-another-100-million-1519915621 [Accessed on May 20, 2019]

Ward, S. (2019). Target Marketing and Market Segmentation. Available at: https://www.thebalancesmb.com/target-marketing-2948355 [Accessed on May 24, 2019]

Wertz, J. (2018). Don't Spend 5 Times More Attracting New Customers, Nurture The Existing Ones. Available at: https://www.forbes.com/sites/jiawertz/2018/09/12/dont-spend-5-times-more-attracting-new-customers-nurture-the-existing-ones/#3ba0f4f05a8e [Accessed on March 22, 2019]

Wessing, T. (2019). EU Japan Adequacy Decision in force. Available at: https://www.lexology.com/library/detail.aspx?g=1d8e4223-b88d-485d-9e49-4787bbd11995 [Accessed on July 22, 2019]

Willson, B. (2018). YouTube Ads Beginner to Expert in 1 Video | 2019 YouTube AdsTutorial [Video file]. Available at: https://www.youtube.com/watch?v=jgV5Jxr8Va4 [Accessed on May 21, 2019]

Win Marketing. (n.d.). marketing objectives - where do you want to be? Available at: https://www.winmarketing.co.uk/marketing-strategy/marketing-objectives/ [Accessed on May 23, 2019]
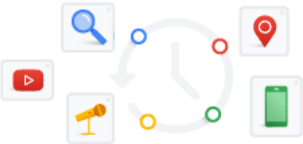
YouTube. (2018). Terms of Service. Dated: May 25, 2018. Available at: https://www.youtube.com/static?template=terms [Accessed on May 28, 2019]

Zheng, D. (2018). Why Is Google Analytics Inaccurate? Available on: https://www.crazyegg.com/blog/why-is-google-analytics-inaccurate/ [Accessed on March 23, 2019]

Appendix A: Activity controls for Google account used in experiment

Appendix B: Ad personalization settings on account used for experiment

Ad personalization

Google makes your ads more useful on Google services (such as Search or YouTube), and on websites & apps that partner with Google to show ads. Learn more

Ad personalization is ON

MORE OPTIONS

How your ads are personalized

Ads are based on personal info you've added to your Google Account, data from advertisers that partner with Google, and Google's estimation of your interests. Choose any factor to learn more or update your preferences. Learn more

18 44 years old	Male
Action & Adventure Films	Adventure Games
Browser Games	Computer & Video Games
Computer Components	Computer Hardware
Computer Peripherals	Computers & Electronics
Cooking & Recipes	Custom & Performance Vehicles
Family Films	Food & Grocery Retailers
Hockey	Laptops & Notebooks
Luxury Vehicles	Mobile Phones
Restaurants	Roleplaying Games
Science Fiction & Fantasy Films	TV Sci-Fi & Fantasy Shows
Word Games	