TRACK AND TREAT – USAGE OF E-MAIL TRACKING FOR NEWSLETTER INDIVIDUALIZATION

Research paper

Bender, Benedict, Chair of Business Informatics, esp. Processes and Systems, University of Potsdam, Potsdam, Germany, bbender@wi.uni-potsdam.de

Fabian, Benjamin, Chair of Business Intelligence and Data Science, Hochschule für Telekommunikation Leipzig (HfTL), Leipzig, Germany, fabian@hft-leipzig.de

Haupt, Johannes, Chair of Information Systems, Humboldt University of Berlin, Berlin, Germany, johannes.haupt@hu-berlin.de

Lessmann, Stefan, Chair of Information Systems, Humboldt University of Berlin, Berlin, Germany, stefan.lessmann@hu-berlin.de

Neumann, Tom, Humboldt University of Berlin, Berlin, Germany, tom.neumann.1@student.hu-berlin.de

Thim, Christof, Chair of Business Informatics, esp. Processes and Systems, University of Potsdam, Potsdam, Germany, cthim@wi.uni-potsdam.de

Abstract

E-Mail tracking mechanisms gather information on individual recipients' reading behavior. Previous studies show that e-mail newsletters commonly include tracking elements. However, prior work does not examine the degree to which e-mail senders actually employ gathered user information. The paper closes this research gap by means of an experimental study to clarify the use of tracking-based information. To that end, twelve mail accounts are created, each of which subscribes to a pre-defined set of newsletters from companies based in Germany, the UK, and the USA. Systematically varying e-mail reading patterns across accounts, each account simulates a different type of user with individual reading behavior. Assuming senders to track e-mail reading habits, we expect changes in mailer behavior. The analysis confirms the prominence of tracking in that over 92% of the newsletter e-mails contain tracking images. For 13 out of 44 senders an adjustment of communication policy in response to user reading behavior is observed. Observed effects include sending newsletters at different times, adapting advertised products to match the users' IT environment, increased or decreased mailing frequency, and mobile-specific adjustments. Regarding legal issues, not all companies that adapt the mail-sending behavior state the usage of such mechanisms in their privacy policy.

Keywords: E-Mail Tracking, Newsletter, Individualization, Personalization, Privacy

1   Introduction

E-mail tracking encompasses methods for gathering information regarding an individual user's reading behavior. Previous studies show that professional e-mail senders routinely embed tracking elements in newsletters and other marketing communication (Fabian et al. 2015). Since tracking is often conducted without consent of the tracked individual, such practices raise ethical and privacy concerns,

especially because the majority of users is unaware of the possibility to track e-mail reading behavior (Mandl et al. 2015).

E-mail tracking approaches split into tracking links and tracking images. The former use embedded references to collect information once a user opens the link in an e-mail. In this sense, the tracking link approaches requires active participation from the user in the form of clicking a link. Tracking images are images embedded in HTML-based e-mails, which e-mail clients fetch from a (tracking) server once a user opens an e-mail. They facilitate data collection regarding the user reading behavior *without the recipient's permission* (Bender et al. 2016), thus exacerbating their threat to data privacy and justification from an ethical point of view.

Previous research focuses on the prevalence of e-mail tracking (Fabian et al. 2015) and the detection of potential tracking images within e-mail communication (Bender et al. 2016). A limitation of prior work lies in its focus on the detection of elements that *potentially* facilitate tracking. For example, embedding a tracking image in an e-mail fulfills the technological prerequisites to track whether a user opens an e-mail. However, confirming the presence of tracking elements in e-mails does not clarify the extent to which senders actually process and employ the information they can potentially gather. Examining the actual use of tracked information is the goal of this paper. In particular, this study clarifies whether e-mail senders adjust their communication policies in response to user data gathered through tracking. In line with prior work, we focus on professional e-mail newsletters because such communication serves a marketing goal and thus incentivizes senders to individualize e-mail messages.

To examine the use of tracking data by commercial mail senders, we design an experiment with twelve e-mail accounts, each of which simulates a specific type of user with individual e-mail reading behavior. We ensure that behavioral differences across user accounts are easy to track by means of tracking images (Suneetha & Krishnamoorthi 2009). Each account subscribes to the same set of professional newsletters, which we gather from companies of various industries. We concentrate on German, British and US companies to evaluate cross-country differences related to different regulations and legal restrictions.

To the best of our knowledge, this is the first experimental study that provides evidence that companies actually use the data they collect through e-mail tracking to adjust marketing communication on an individual level. Our analysis also reveals that a fraction of companies employ personal response data to individualize the frequency, timing and content of marketing communication. These results confirm that data collection, storage and analysis on the personal level takes place and emphasizes the need for additional research regarding the extent of identified privacy risks and the development of efficient protection strategies.

We organize the paper as follows. The next section discusses prior work on e-mail tracking and related tracking technologies. Section 3 elaborates on our experimental design. We then present and discuss results in Section 4. Section 5 concludes the paper.

2 E-Mail Tracking Fundamentals

This section discusses the process of e-mail tracking and its technological fundamentals in the form of tracking links and images. *Tracking links* are hyperlinks in an e-mail that are augmented with identifiers, which are not part of the reference but convey information about interaction with the link. In particular, tracking links can include a unique identifier that allows to detect and log whether an individual e-mail recipient follows the link (Fabian et al. 2015). Technically, this is typically realized using an individual link for every recipient to be able to detect any website request using web server analytics (Agosti & Di Nunzio 2007) or a redirection service (Nikiforakis et al. 2014). The latter also facilitates matching the browsing behavior on the target page with an e-mail recipient through the identifier transmitted via the referrer URL in the specialized link (Jin et al. 2010).

Tracking images are external image references within HTML e-mails that contain identifying information. Figure 1 depicts the tracking process for e-mails that reference external image resources. The

sender prepares an HTML e-mail including an image reference augmented by information on the identity of the receiver and the content of the e-mail. After the e-mail is sent, it passes several mail transfer agents (MTAs) until it reaches the receiver's MTA. Next, the recipient opens a mail client, which synchronizes the local mail repository with the newest version of the recipient's MTA. When the recipient opens the e-mail with a tracking image, the mail client requests the image from the referenced destination. The web server logs this request and provides the image to the recipient's mail client. Finally, analysis of the server log files provides detailed insights on the recipient's e-mail reading behavior.
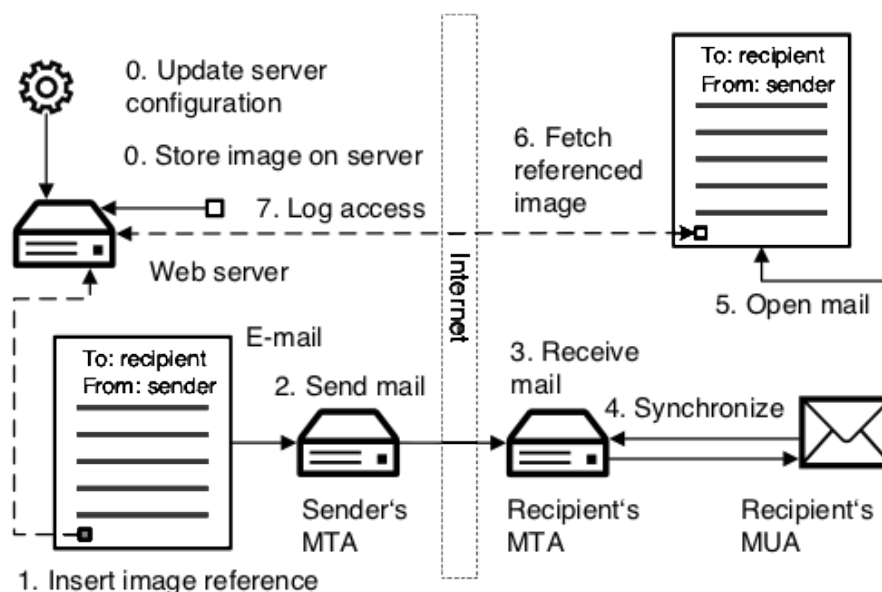

Figure 1.   E-Mail Tracking Process (see Bender et al. 2016)

2.1   Related Literature

E-mail tracking can be interpreted as the application of common web tracking mechanisms in the e-mail context. In the following, we discuss 1) relevant web tracking studies and 2) studies specifically related to e-mail tracking.

The tracking of web users has been an activate topic of research, see the recent surveys by Bujlow et al. (2017) and Ermakova et al. (2018) as well as individual studies (Evans et al. 2003; Bouguettaya and Eltoweissy 2003; Han et al., 2012; Roesner et al., 2012; Gomer et al., 2013; Hamed et al., 2013; Acar et al., 2014; Libert, 2015; Schelter & Kunegis, 2016a, 2016b; Englehardt & Narayanan, 2016; Ermakova et al., 2017). The web-browsing behavior of online users is considered a worthwhile source for detailed profiling (Mitchell, 2012; Falahrastegar et al., 2016) to improve commercial activities such as targeted advertising (Roesner et al., 2012; O'Connell, 2014). Enabled by a variety of techniques (Besson et al., 2014; Sanchez-Rola et al., 2016; Bujlow et al. 2017), web tracking has become ubiquitous (Roesner et al., 2012; Schelter & Kunegis, 2016a, 2016b; Ermakova et al., 2017) on single sites, but also across websites and even across devices (Mayer & Mitchell, 2012; Gomer et al., 2013; Falahrastegar et al., 2014, Brookman et al., 2017). Some articles have analyzed the methods and extent by which relevant information can be extracted from tracking data (Suneetha et al. 2009, Bujlow et al. 2017). Besides targeted advertising (Sanchez-Rola et al., 2016; Parra-Arnau, 2017), web tracking can be applied for personalization, advanced web-site analytics, and social network integration (Sanchez-Rola et al., 2016; Mayer & Mitchell, 2012; Roesner et al., 2012).

For online users, web tracking practices can result in increased online privacy risks (Moscato et al. 2009; Jin et al. 2010; Mayer & Mitchell, 2012; Roesner et al., 2012), including price discrimination,

government surveillance, and identity theft (Bujlow et al., 2017). The extent to which tracking is made transparent within the privacy policies of business to consumer companies depends on user expectations (Moscato et al. 2013).

E-mail tracking, i.e. the use of web tracking methods in e-mail communication, has become a growing concern in scientific literature as well as in the public press. A description of techniques for extracting user information from e-mails is given by Chipperfield et al. (2006) and Cselle et al. (2007). As discussed by Fabian et al. (2015), e-mail tracking allows the collection of detailed information on individual reading behavior without explicit consent of the user. In this regard, tracking images represent a more severe privacy issue since information is collected automatically when an e-mail is opened, whereas tracking links require active clicking on the referenced content. Bender et al. (2016) provide a first international study regarding the use of e-mail tracking in commercial newsletters and focus on the conceptualization of potential countermeasures.

There are studies that highlight some functional advantages of e-mail tracking, e.g., that the basic structure of the e-mail service does not allow a sender to be certain that a message is really delivered to the right receiver (Oppliger 2007). Schmidt (2013) discusses the usage of tracking images and evaluates current protection through commonly used e-mail software for personal use. The information that can be collected comprises primary information that can be gathered directly from the tracking server logs, and secondary information, based on additional resources to enhance and combine with the primary information. Examples of primary information include the time or the client's user-agent string that was used to request the image. Examples for secondary information are the location from which the e-mail is retrieved as well as potentially a user's affiliation, or if an e-mail has been printed or forwarded (Bender et al. 2016). The combination of information allows building a profile of the individual user's behavior.

An important aspect that distinguishes e-mail tracking from general web tracking techniques is that the collected data is not anonymous, since it can be directly attributed to an e-mail recipient identified by a unique e-mail address (Jin et al. 2010). Since email addresses often contain the name of the individual and the name of an affiliated institution in the domain and are often used to sign in on several websites, some of which may require personal information, they facilitate the identification of individuals to a larger extent than web tracking.

3   Study Design

We conduct a controlled experiment by simulating user interaction with marketing newsletters in order to evaluate whether e-mail senders vary their communication and sending policies depending on the recipient's reading behavior. Using a set of artificial user accounts allows us to minimize confounding factors by standardizing user characteristics. This section describes the experimental setup used for data collection and the user behavior profiles.

To collect data, we set up twelve up e-mail accounts on Gmail. Ten accounts simulate a specific, consistent user behavior. The remaining two accounts do not conduct any activity to allow comparison and validation of the results. We create all user identities to be older than 21 years to eliminate potential restrictions in the offerings and choose user birthdays to be outside the data-gathering period to eliminate bias from potential birthday related offerings. Given our focus on tracking, all identities share the same gender (male) to avoid gender-specific offerings in view of the content comparison. Other personal information required during newsletter registration is held constant over identities and matched to characteristics we expect for subscribers of each company. For example, we use country-specific addresses to prevent a potential relocation to another subsidiary of the company.

For each newsletter subscription, we ensure a 'clean' browser environment to prevent potential linking between the accounts and the deduction of preferences from the browser history, which is a common practice in web tracking (Nikiforakis et al. 2014). For example, we delete cookies, history and form entries, etc. from the web browser cache and begin each registration process from a new browser session. In addition, we ensure location-specific IP addresses within the subscription process. Finally, we

limit the information provided to companies during newsletter signup to mandatory entries. This helps preventing content variation based on preferences or attributes given during the subscription. In case where such information was mandatory, we provided the same data for all accounts.

We register each mail account for the same set of commercial newsletters selected from the largest e-commerce companies based in Germany, the United Kingdom, and the United States of America. Large companies are likely to have knowledge as well as the resources to employ individual targeting and complex analytic solutions. This study focuses on companies from one industry, online retail, for several reasons. First, online retail specializes in digital business. Therefore, we expect companies to be well developed with regard to technological possibilities in general and technologies to enhance customer-centric processes in particular. Second, a large product portfolio simplifies segmentation and individualization of offerings compared to other industries (e.g., public transport, manufacturers, etc.). Third, personalization and customer targeting are established success factors in online retail (Golrezaei et al. 2014). It is thus plausible to expect e-tailors to be pioneers in personalization.

Within retail, we consider six areas: clothing, electronics, general retails, home goods, supermarket and tourism. Within these areas, we select on successful and large business-to-consumer retailers as determined by means of country-specific rankings based on revenue or sales (Germany: EHI Retail Institute & Statista (2017), USA: eMarketer, cited in Zaczkiewicz (2016), UK: Internet Retailing Media UK (2016)). The rationale for this selection is that large retailers are more likely to have the resources and know-how to engage in tracking, targeting, and personalization. The listed companies were assigned to each of the six retail categories, if applicable. Trading companies without a retail focus were excluded from the study in order to ensure a defined and comparable sample of companies with an incentive for newsletter personalization. Furthermore, globally active trading companies, e.g. Amazon, were excluded from the study, since the attribution to a single region is imprecise and might distort the country-specific results.

In total, each of the 12 e-mail accounts subscribes to 52 company newsletters. E-mails are collected for a ten-week period from the 12th to the 21st calendar week of 2017. We argue this a reasonable time span for newsletter senders to collect user information and to adjust or individualize e-mails.

Simulating different user behaviors and retrieving e-mails to access referenced images requires a dedicated and customizable software. We have developed a corresponding system using the Java programming language. Java is a suitable choice because it features various easy-to-use components such as JavaMail for mail access and JSoup for parsing XML-based files like HTML-based mails that jointly provide the required functionality. All information gathered for the experiment is stored in a relational database. Importantly, to simulate different scenarios, all images within one mail are fetched according to the individual user/account profile.

Since the study aims to evaluate the use of information gathered through tracking images during e-mail reading, the experiments need to simulate relevant user behavior. For the e-mail tracking process (Figure 1) it is essential to fetch external referenced content during the reading process. Thereby, the tracker can use all the information available during the image request to build a profile. From a conceptual point, we divide the information in infrastructural and behavioral aspects that might influence the newsletter targeting. Infrastructural information such as the devices used are typically static thus making it easier to conduct corresponding targeting activities. The behavioral aspects are more dynamic and it is therefore more complex to deduce corresponding targeting activities.

To implement the behavioral user profiles, we develop a separate request function for every test account. Each account can have its own settings for the IP address, user agent string and predefined execution time. To allow for simultaneous image requests of different test accounts, a multithreading procedure has been employed to comply with concurrency requirements. We use predefined timers to start the respective threads, which helps steering the exact time sequences for image requests.

Table 1 gives the experimental factors for each account. To simulate different user behavior, we vary the time and frequency of e-mail access and the device type and location.

| Account | Factor | | Device type | Location |
	Frequency	Reading time		
1	1/day	random	Windows	Germany
2	3/day	random	Windows	Germany
3	1/day	fixed time (1pm)	Windows	Germany
4	1/day	fixed time (10pm)	Windows	Germany
5	1/day	3 minutes after reception	Windows	Germany
6	random	random	Windows	Germany
7	1/day	random	Windows	USA
8	1/day	random	OSX	Germany
9	1/day	random	Android	Germany
10	1/day	random	iOS	Germany
11	never	none	-	-
12	never	none	-	-

Table 1.   Simulated Behavior of E-Mail Accounts in the Experiment

Accounts open e-mails with a reading frequency fixed at once a day with the exception of account #2, which opens e-mails three times a day, and account #6, which opens each e-mail at a random time and frequency, but at least once. The time at which e-mails are opened is randomly drawn from a uniform distribution over the minutes of the day per account and day for most accounts. Accounts #3 and #4 open all new e-mails at a fixed time of the day at 1 a.m. and 10 p.m., respectively. Account #5 opens e-mails three minutes after they are received.

We fix the device type for each account through manipulation of the user-agent string to the desktop (Windows/OSX) or mobile (Android/iOS) operating system of the most common vendors Microsoft and Apple, respectively. Each account accesses e-mails and external content through one of two proxy servers to fix the location derivable from the IP address. We use the same proxy server with a static IP address for the duration of the experiment. Locations are either a German university or a school located in Hanford, California, for account #7.

4 Study Results

We begin the presentation and discussion of empirical results with reporting descriptive statistics related to the newsletter e-mails gathered through the user accounts. During the collection period, we receive 12,404 valid e-mails in total, of which 12,346 are in HTML. The HTML e-mails can include tracking images. Not all companies started delivering newsletters. In total, 44 out of the 52 companies sent newsletters and the sending behavior differs across newsletters. Most of the e-mails come from German companies (45.6 %) whereas the share of newsletters from the USA and UK are 34.4 % and 20 %, respectively. Newsletter shares across industries per country show that in Germany the dominating industry is supermarkets, in contrast to both other countries, where the general retail and home goods companies use e-mail marketing much more often. The clothing sector is similarly prominent in all three countries.

We employ the detection model of Bender et al. (2016) to identify tracking images. In all countries, the prevalence of tracking is comparable and at a high level. Newsletters from German companies have the lowest amount with 85.68 % of all e-mails containing at least one tracking image. 93.76% e-mails from the UK contain tracking images and 99.48% from the US. Examining the share of tracking e-mails across industries reveals that all newsletters from the General Retail and Clothing and 99% from the Home Goods sector contain tracking images (see Figure 2). To a lesser degree, 88% of elec-

tronics newsletters and 76% of supermarket e-mails contained tracking images, while touristic newsletters showed the lowest tracking rate with 60%.
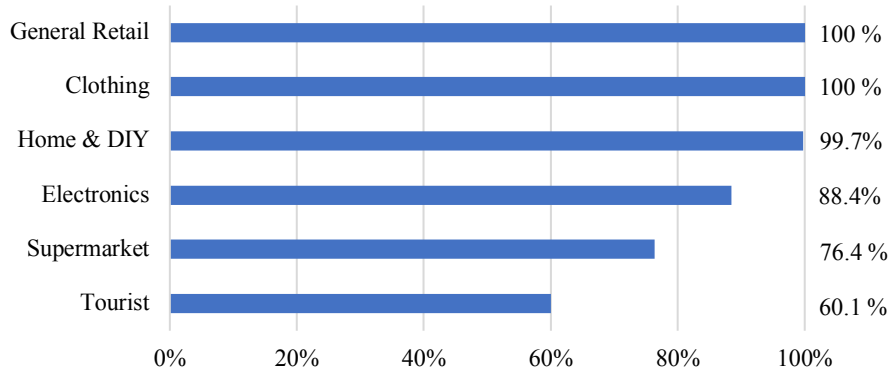

Figure 2.   Tracking rate for different trading industries

In the following subsections, we evaluate each experimental factor varied in the experimental design. We begin with an analysis of the overall number of mails received per account. Afterwards, location specific adjustments based on the offsite account are evaluated. We then evaluate content variation between the newsletters in the different accounts. Finally, results regarding varying sending behavior for the individual simulated accounts are discussed.

4.1   Amount of E-Mails Received

Analyzing the number of received e-mails for each account provides a first indication of differential sender behavior. The most remarkable aspect is that both validation accounts received substantially less e-mails than all other test accounts. This is a clear indication that companies observe the opening rates of subscribers and adjust the sending behavior accordingly.
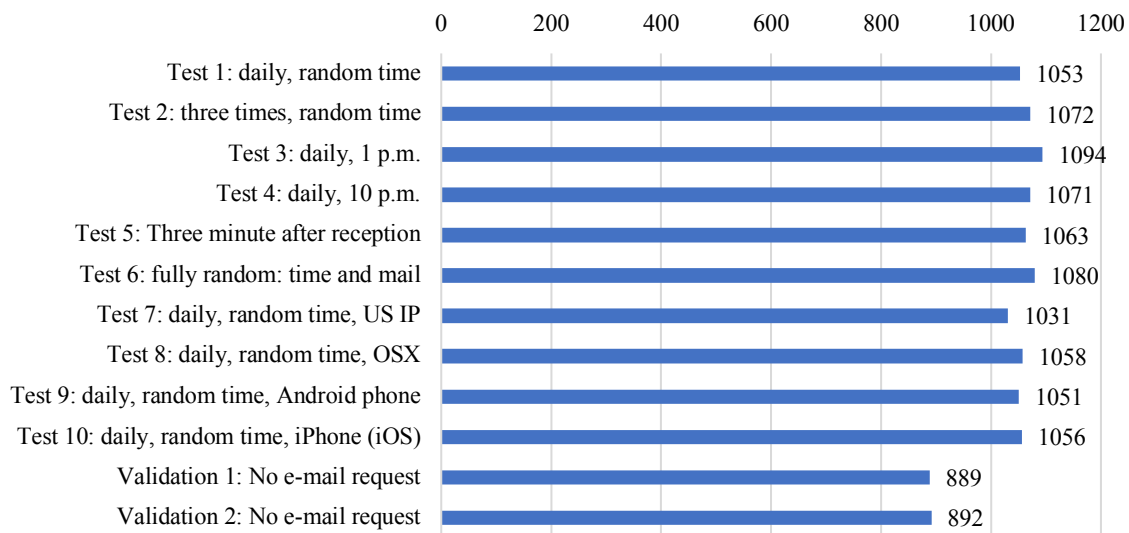

Figure 3.   Received E-Mails per Account

Within the other accounts, the number of received e-mails ranges from 1,026 to 1,094 messages, with an average of 1,063 mails per account. The noticeably smaller number for test account #7 is due to one newsletter that, for unknown reasons, has not been delivered to test account #7 while being active in all other accounts. To use a consistent and comparable data basis, we exclude this newsletter, i.e. 21 e-mails per account, from the subsequent analysis. The other differences in the number of received e-

mails between accounts can be attributed to small divergences in the number of mails sent across all companies (e.g., as opposed to a large deviation in the sending behavior of a small number of companies). During the data-gathering period, we observe a slight increase in the number of received mails for all test accounts, but not for the validation accounts.

We observe a significant difference in the number of received e-mails between the ten treatment and two control accounts (Welsh t-test, df = 9.9998, p < 0.001). The difference provides evidence that companies use e-mail tracking information to adjust their communication policy. Further analysis of sending patterns reveals that companies stopped sending e-mails completely after no e-mail openings were tracked, with two companies stopping after only one unopened e-mail and two companies stopping after two to three e-mails. One company explicitly acknowledged the observed behavior after sending four unopened newsletters to the validation accounts by sending an e-mail with the message "we miss you" and special promotions. None of the test accounts receives a comparable message. We take the retention offer as strong evidence that the confirmation of e-mail openings provided by tracking images is used to target customer individually.

Interestingly, we also observe a novel newsletter in the data. In particular, one company not within the experiment selection started sending messages to the test accounts but not the validation accounts, without explicit newsletter subscription by any account. An analysis of the company's affiliation revealed that the company is affiliated with a company that the accounts subscribed to. We interpret these findings to confirm that i) e-mail addresses are transferred to the subsidiary and ii) that the addresses are further qualified with information collected through e-mail tracking. We interpret the selective behavior to show that the company uses observed reading behavior to select only active accounts for transfer. Further research is necessary to establish if the data transferred to the subsidiary includes information on the individual reading behavior in addition to the e-mail address.

We observe weakly significant variation in the amount of e-mails received between mobile and desktop users (Welch t-test, df = 7.67, p-value = 0.038), with the mobile accounts receiving less e-mails.

4.2 Location-Specific Adjustments

We simulate one user to open newsletters from a different country to test for location-based targeting. Since global companies have local subsidiaries that could target customers directly, we expect to observe adjustment of the sender or localized communication content. However, the data do not indicate major differences in the sending behavior. None of the issuing companies changes its top-level domain or the address from which newsletters are sent in response to test account #7 opening each e-mail from a U.S. IP address. Possible reasons for companies to ignore the IP location are that IP addresses can convey false information, e.g. if a mail proxy server or VPN is in use, and that location information may be temporary, e.g. when a user opens mail on holiday.

4.3 E-Mail Content Adjustment

Beyond the adjustment of the sending schedule, e-mails can be personalized by changes in e-mail phrasing, formatting and content. We therefore go on to compare the corresponding e-mails across the different accounts for their body length, textual content and image URLs (excluding the tracking images).

The length of the text in e-mails received by each account could provide a first indication that systematic adjustment of e-mail content takes place. We conduct an analysis-of-variance (ANOVA) on e-mail length in characters but are unable to reject the null hypothesis of e-mail length with equal mean for all accounts ($F(11, 12132) = 0.84$, $p = 0.60$). Nevertheless, we often find substantial variation between accounts for a single e-mail. After inspection of these differences, we propose that A/B testing to be the main cause for the observed variance, where senders are using different design versions of the same e-mail to test the effectiveness of design choices. In some cases, several accounts received the exact same version of one e-mail while the other accounts received a different version.

In addition to the text of e-mails, marketing practice suggests to prefer short subject lines when targeting mobile customers due to their limited screen size. The difference in the average length of the e-

mail subject for mobile compared to desktop devices is small at 0.8 characters and statistically insignificant (Welsh t-test, df = 9.05, p-value = 0.14).

Similarly, we frequently observe the use of different image versions or different icons but are unable to determine any structure within the deviations. Other differences in the e-mails are variations of personalized promotion codes or the recipients' e-mail address mentioned in the fine print.

However, one electronic retail company specialized on computer, notebooks, mobile devices and peripherals adjusts marketing content based on the user's device. We observe that test account #10, which simulates an iPhone receives mails with substantially more Apple-related products than the other accounts over the full observation period. An example of the typical product offering for a comparison account (left) and account #10, which simulates the iPhone client, is presented in Figure 4. To test these observations, we identify the keywords Apple, iPhone, and MacBook, which occur substantially more often for this account than for the comparison accounts. While the iPhone account receives 576 mails containing Apple keywords, the account simulating the Apple laptop (#8) receives 517, which is slightly above the average of comparison accounts at 509.1. The difference in the average count of keywords between the accounts simulating an Apple system (#8 and #10) and all other accounts is not significant (Welsh t-test, df = 1.7729, p-value = 0.40).

We conclude that, even though variation could be observed, we are unable to identify statistically significant patterns or systematic variation and find no evidence that companies personalize the content of the newsletter based on information collected through e-mail tracking.
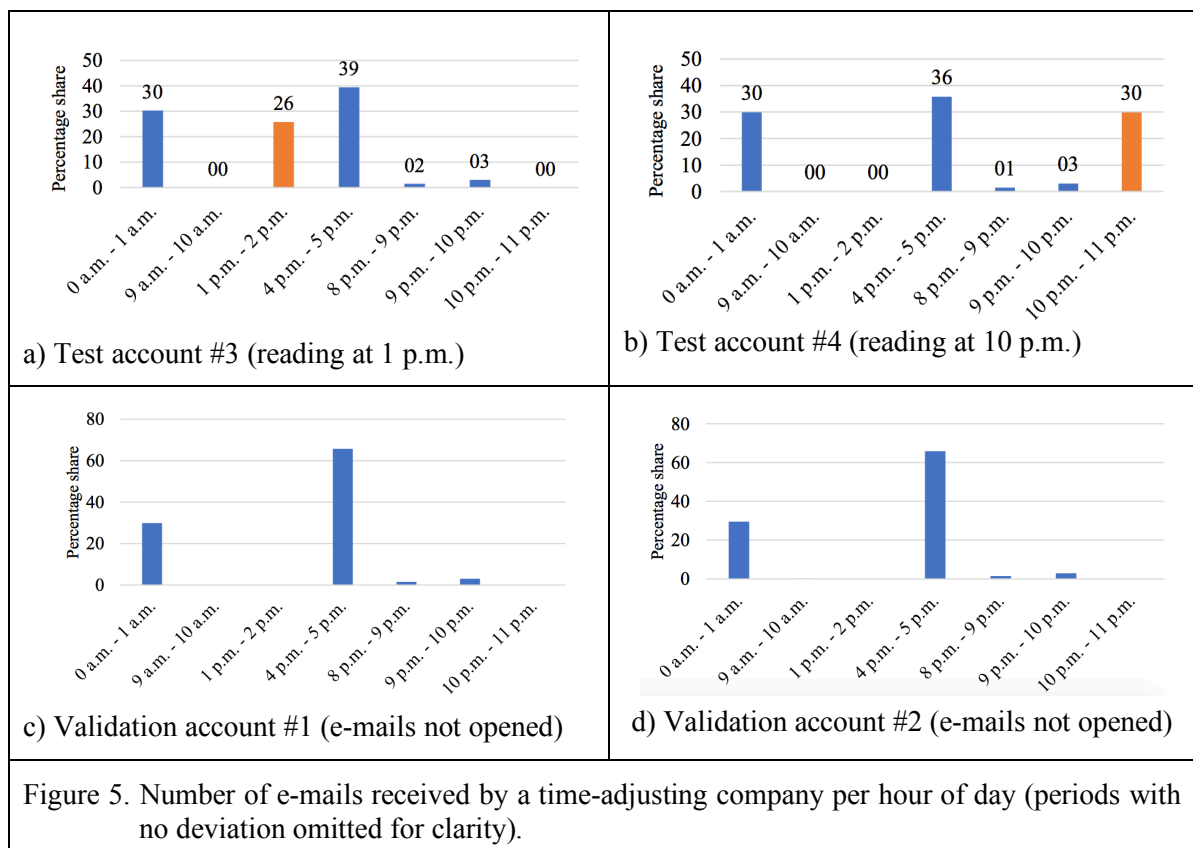

Figure 4.   Example Mail from Electronic Retailer. Test account #10 (right side) receives information on considerably more Apple products than the other accounts (example left)

4.4   Sending-Time Adjustments

E-Mail users read their mails at different times. Some may read mails occasionally within usual business hours only, whereas others check mails more frequently. Data on a user's reading behavior may convey information regarding her digital media usage and daily routine, which are valuable insights for marketing. To examine whether e-mail senders adjust their communication according to the reading times of recipients, test accounts #3-5 simulate different reading styles. Account #3 read mails at 1 p.m., while test account #4 reads at 10 p.m. Test account #5 simulates frequent e-mail checking upon notification that an e-mail has been received.

Taking the whole set of mails into account, we observe that only a single, very active company adjusts sending times in response to recipient behavior (Figure 5). Figure 5c and 5d show the e-mails from this company received by the validation accounts without e-mail access behavior. The mails that these receive as well as their time distribution are identical. The time distribution of test account #3 (Figure 5a) and test account #4 (Figure 5b) received differs substantially from the validation accounts. Although both accounts receive the same amount of 100 e-mails, the time at which these are sent differs and matches the different reading (time) preferences simulated by the accounts. While both validation accounts and test account #4 receive no mails between 1 and 2 p.m., test account #3 that read mails at 1 p.m. receives over a quarter of its mails in this timeframe. Test account #4 shows a similar result for its reading time at 10 p.m. In view of the magnitude of the effect, Figure 5 provides strong evidence in favor of a systematic variation in the communication style of the e-mail sender. On the other hand, it has to be noted that only a single sender in our sample adapts sending times to users' reading time preferences.

a) Test account #3 (reading at 1 p.m.)

b) Test account #4 (reading at 10 p.m.)

c) Validation account #1 (e-mails not opened)

d) Validation account #2 (e-mails not opened)

Figure 5. Number of e-mails received by a time-adjusting company per hour of day (periods with no deviation omitted for clarity).

5 Discussion

5.1 Discussion of the Results

The analysis of the e-mails gathered during the 10-week period shows that over 92% of the e-mails from the UK, US and Germany contain tracking images. The prevalence of tracking within the online retail area supports previous studies on the wide application of e-mail tracking mechanisms. During the newsletter signup, companies typically present their privacy terms and conditions. We investigated the statements for the newsletters used in this study. Only 21 out of 52 mention the possibility of using the data gathered for personalization and individualization. Especially newsletters from the UK and US-firms use tracking images to target customers individually while not stating this in their privacy

statements. On the other hand, German companies that employ tracking images consistently state this explicitly in their privacy statements. Some German newsletters also offer the option to choose whether companies may use the gathered information for personalization. Given the ubiquity of e-mail tracking and the lack of transparency regarding its use, this study extends the literature by providing an analysis on the reaction of trackers to observed recipient behavior.

Overall, we observe adjustments in sending behavior for 13 out of 44 marketing newsletters, with adjustments affecting the sending behavior. We find that senders respond most often to e-mail opening actions or the lack thereof. Several companies adjust their sending behavior upon realizing that receivers do not open newsletters. Considering users' reading time patterns, we find evidence that a single company within the sample adjusts their communication to accommodate the simulated reading times. We find no evidence that opening a mail multiple times or the location of e-mail access impacts sending behavior.

Surprisingly, we are unable to identify systematic and statistically significant personalization of e-mail content based on information collected through e-mail tracking. While we suspect personalization of product offerings based on the simulated device type for one company, further research is necessary to confirm our findings. While targeting customer individually according to their preferences and interests can be expected to increase click-through and product sales (e.g., Golrezaei et al. 2014), content personalization seems to be based on other data, such as previous purchases or similar user interaction.

5.2   Limitations

This study exhibits some limitations that give rise to future research. First, the ten-week time frame for data collection assumes companies to react relatively swiftly. Although some of the newsletter senders adjusted their behavior in this time span, especially companies with a less frequent newsletter delivery may have not had enough user data to systematically react to the user. A longer observation period may be necessary to identify adjustments for non-frequent newsletter.

Another limitation, which simultaneously is the explicit focus of this study, is the concentration on tracking images. Since the data gathered through tracking images needs to be analyzed and interpreted, it could be more complex to employ this as a basis for individualization than it might be to use other techniques (e.g., tracking links). To enhance the results and ideas of this study, it would be useful to include further types tracking mechanisms in further analyses. To consider tracking links in a follow-up study would add another dimension to the behavioral aspects. Another aspect would be selective reading of mails (e.g., select which mail to open based on their title).

The simulated user behavior was restricted to reading behavior of the e-mail. In practice, it is likely that user behavior on the website of the company, matched via the user account or link tracking, provides additional data to be used for personalization of marketing messages. Simulation of actual browsing and purchasing behavior, while complex to conduct in an automated fashion, has the potential to uncover additional personalization strategies, e.g., retargeting of abandoned products, which could be considered in further studies.

Furthermore, we cannot entirely rule out the possibility that some senders might have recognized our test accounts as artificial. The software components cause the simulated users to behave in a very consistent manner, which is unlikely for real e-mail users. Some newsletter senders might have realized this unusual pattern and might have reacted to it. On the other hand, use-cases for fake e-mail accounts that only read e-mails seem rather limited (e.g., compared to fake accounts for spamming, which would send a huge amount e-mails). In this regard, it is questionable whether senders have implemented sophisticated detection strategies for this kind of suspicious behavior we rely on in this study. For future studies, it would be useful to integrate more random behavior into the experiment to prevent detection mechanisms from uncovering our mail-reading engine.

Finally, a statistical limitation comes from the fact that we have only twelve accounts available. This comes from the vast effort to manually subscribe, for each account, to a large number of newsletters. However, empirical evidence in the form of descriptives and mean comparison across groups clearly suffers from the number of accounts, which, although substantially larger than what has been consid-

ered in prior work, is relatively small. We argue that this issue is inherent to the research problem and cannot be overcome easily. Captchas prevented an automation of newsletter subscriptions. On the other hand, crowdsourcing supporters of such research to assist with manual labor would inevitable raise awareness of the research, which might carry over to mailers and thus introduce bias. In this regard, we consider the results presented here as a valuable first evidence into a sparsely researched phenomenon but also strongly encourage further research to expand the scale of the analysis.

The study focused on the trading industry, even though many other industries are equally relevant. Future studies should therefore incorporate other industries to gain a more diverse picture on the application of e-mail tracking mechanisms.

6   Conclusion

E-mail tracking facilitates gathering information regarding the individual recipient reading behavior. Former studies reveal that professional e-mail newsletters commonly include tracking elements (Fabian et al. 2015). However, former studies do not check whether information that can be gathered through the tracking images is actually used by the e-mail senders.

Our experiment strives to close this gap and validates the usage of e-mail tracking information by e-mail senders. To that end, the study uses twelve mail accounts each of which simulates an individual user behavior to gather newsletters over a 10-week period from retail companies across Germany, the UK, and the USA. We find that 92% of the e-mails from the UK, US and Germany contain tracking images.

The experimental data shows that the tracking images detected are in fact used to assess individual behavior and to adjust marketing communication on the individual level. This confirms the relevancy of the potential threats to user privacy resulting from e-mail tracking. Even though e-mail clients typically allow to block external referenced content, such as images, and thereby counteract e-mail tracking images, studies revealed this blocking approach to be impractical for image-rich e-mails, such as newsletters. Blocking images completely is not assumed to be an effective strategy, which is why selectively blocking content is suggested (Bender et al. 2016). The experiment further reveals that companies employ personal response data to individualize marketing communication for newsletters. We observe individualization in 13 out of 44 (30%) newsletters. We find statistically significant response of senders to e-mail opening. In reverse, several companies stop delivering newsletters upon realizing that receivers do not open newsletters. With regard to time patterns, we find a single company to adjusts their newsletter mailings to accommodate the different simulated reading times. No evidence was found, that multiple openings as well as location related aspects impact sending behavior. We find a statistically weak significance for device-category specific adjustment of mail frequency, with mobile devices receiving slightly less e-mails than the desktop accounts. With regard to content adjustments, we are unable to show systematic variations as expected for online retailers. Nonetheless, a single company adjusted the products offered to infrastructure characteristics.

Observed adjustment patterns can be considered easy-to-implement options to individualize communication, especially for e-commerce retailers due to their typically advanced IT-systems regarding analytics and wide product portfolio. Targeting users individually is also important for such companies to succeed in customer acquisition, growth, and retention. Regarding legal issues, we find that that not all companies which adapt the mail sending behavior inform subscribers of such mechanisms in their privacy policy. These results confirm that data collection, storage and analysis on the personal level takes place and emphasizes the need for additional research regarding the extent of identified privacy risks and the development of efficient protection strategies.

References

Acar, G., Eubank, C., Englehardt, S., Juaerez, M., Narayanan, A. and Diaz, C (2014). *The Web Never Forgets: Persistent Tracking Mechanisms in the Wild*. Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security (CCS '14), USA, ACM, pp. 674-689.

Agosti, M., & Di Nunzio, G. M. (2007). *Gathering and Mining Information from Web Log Files*. First International DELOS Conference, Pisa, Italy, Revised Selected Papers, Springer Berlin Heidelberg, pp. 104-113.

Bender, B., Fabian, B., Lessmann, S., and Haupt, J. (2016). *E-Mail Tracking: Status Quo and Novel Countermeasures*. Proceedings of the International Conference on Information Systems (ICIS) 2016, Dublin, Ireland.

Besson, F., Bielova, N., and Jensen, T. (2014). *Hybrid Information Flow Monitoring Against Web Tracking*. Proceedings of the 26th Computer Security Foundations Symposium, New Orleans, USA, IEEE Computer Society, pp. 240-254.

Bonfrer, A., and Drèze, X. (2009). *Real-Time Evaluation of E-Mail Campaign Performance*. Marketing Science, 28(2), 251-263.

Brookman, J., Rouge, P., Alva, A., and Yeung, C. (2017). *Cross-Device Tracking: Measurement and Disclosures* Proc. Privacy Enhancing Technologies, 2, pp. 133–148.

Bujlow, T., V. Carela-Español, J. Solé-Pareta, and P. Barlet-Ros (2017). *A Survey on Web Tracking: Mechanisms, Implications, and Defenses*, Proceedings of the IEEE, 105, 8, pp. 1476-1510.

Chipperfield, T.R., M.G. Foulger, J.S. Cooper, and A.C. Storms (2006). *System and method related to generating and tracking an E-mail campaign*. USA Patent US Patent 7,065,555.

Cselle, Gabor, Keno Albrecht, and Rogert Wattenhofer (2007). *BuzzTrack: Topic Detection and Tracking in Email*. In: Proceedings of the 12th International Conference on Intelligent User Interfaces. New York, USA

EHI Retail Institute, and Statista. (2017). *B2C-E-Commerce: Ranking der Top100 größten Online-Shops nach Umsatz in Deutschland im Jahr 2016 (in Millionen Euro)*. In: Statista - Das Statistik-Portal. https://de.statista.com/statistik/daten/studie/170530/umfrage/umsatz-der-groessten-online-shops-in-deutschland/.

Englehardt, S., and A. Narayanan (2016). *Online Tracking: A 1-million-site Measurement and Analysis*, 23rd CCS, Vienna, Austria, ACM, pp. 1388-1401.

Ermakova, T., A. Hohensee, I. Orlamünde, and B. Fabian (2017). *Privacy-Invading Mechanisms in E-Commerce – A Case Study on German Tourism Websites*, International Journal of Networking and Virtual Organisations.

Ermakova, T., Fabian, B., Bender, B., and Klimek, K. (2018). *Web Tracking – A Literature Review on the State of Research*. Proceedings of the 51th Hawaii International Conference on System Sciences (HICSS 51), Hawaii Big Island, USA.

Fabian, B., Bender, B., and Weimann, L. (2015). *E-Mail Tracking in Online Marketing - Methods, Detection, and Usage*. Proceedings of the 12. International Conference on Wirtschaftsinformatik (WI 2015), Osnabrück, Germany.

Falahrastegar, M., H. Haddadi, S. Uhlig, and R. Mortier (2016). *Tracking Personal Identifiers Across the Web*, Int. Conf. on Passive and Active Network Measurement, Heraklion, Greece, LNCS, Springer, pp. 30-41.

Golrezaei, N., Nazerzadeh, H., and Rusmevichientong, P. (2014). *Real-Time Optimization of Personalized Assortments*. Management Science, 60(6), 1532-1551.

Gomer, R., E. M. Rodrigues, N. Milic-Frayling, and M. C. Schraefel (2013). *Network Analysis of Third Party Tracking: User Exposure to Tracking Cookies through Search*, IEEE/WIC/ACM Int. Joint Conf. on Web Intelligence (WI) and Intelligent Agent Technologies (IAT).

Han, S., J. Jung, and D. Wetherall, (2012). *A Study of Third-Party Tracking by Mobile Apps in the Wild*, Technical Report, University of Washington.

Hasouneh, A. B. I., and Alqeed, M. A. (2010). *Measuring the Effectiveness of E-Mail Direct Marketing in Building Customer Relationship*. International Journal of Marketing Studies, 2(1), 48-64.

Internet Retailing Media UK (2016). *The 2015 IRUK Top500 by Performance Cluster.* http://internetretailing.net/iruk/iruk-top500-2015/

Jin, L., Takabi, H., and Joshi, J. B. D. (2010). *Security and Privacy Risks of Using E-mail Address as an Identity*. Proceedings of the 2010 IEEE Second International Conference on Social Computing, Minneapolis, USA.

Libert, T. (2015). *Exposing the Hidden Web: An Analysis of Third-Party HTTP Requests on 1 Million Websites,* International Journal of Communication, 9, pp. 3544-3561.

Mayer, J.R. and J.C. Mitchell (2012). *Third-Party Web Tracking: Policy and Technology*, IEEE Symposium on Security and Privacy, San Francisco, USA, IEEE, pp. 413-427.

Mitchell, I.D. (2012). *Third-Party Tracking Cookies and Data Privacy*, SSRN Electronic Journal.

Moscato, Donald R., Shoshana Altschuller, and Eric D. Moscato (2013). *Privacy Policies on Global Banks' Websites: Does Culture Matter?* In: Communications of the IIMA 13.4, pp. 91–109.

Nikiforakis, N., Maggi, F., Stringhini, G., Rafique, M. Z., Joosen, W., Kruegel, C., Zanero, S. (2014). *Stranger Danger: Exploring the Ecosystem of Ad-Based URL Shortening Services*. Proceedings of the 23rd International Conference on World Wide Web, Seoul, Korea.

O'Connell, J. (2014). *The State of Programmatic Media*. https://research.adexchanger.com/the-state-of-programmatic-media.html

Parra-Arnau, J. (2017). *Pay-per-Tracking: A Collaborative Masking Model for Web Browsing*, Information Sciences (2017), 385–386, pp. 96–124.

Roesner, F., T. Kohna, and D. Wetherall (2012). *Detecting and Defending Against Third-Party Tracking on the Web*, 10th Int. Conf. on Web and Social Media, San Jose, USA, pp. 155-168.

Sanchez-Rola, I., X. Ugarte-Pedrerp, I. Santos, and P.G. Bringas (2016). *The Web is Watching You: A Comprehensive Review of Web-Tracking Techniques and Countermeasures*, Logic Journal of the IGPL, 25, 1, pp. 18-29.

Schelter, S., and Kunegis, J. (2016a). *Tracking the Trackers: A Large-Scale Analysis of Embedded Web Trackers*. 10th Int. Conf. on Web and Social Media (2016), Cologne, Germany, AAAI Press, pp. 679-682.

Schelter, S., and Kunegis, J. (2016b). *On the Ubiquity of Web Tracking: Insights from a Billion-Page Web Crawl*. arXiv preprint arXiv:1607.07403.

Schmidt, J. (2013). *E-Mail im Visier - Tracking im Alltag aufspüren und abstellen*. c't – Magazin für Computertechnik, 22, 130-135.

Suneetha, K. R., and Krishnamoorthy, K. R. (2005). *Identifying User Behavior by Analyzing Web Server Access Log File*. International Journal of Computer Science and Network Security, 9(4), 327-332.

Thode, W., Griesbaum, J., and Mandl, T. (2015). *"I would have never allowed it": User Perception of Third-party Tracking and Implications for Display Advertising*. Proceedings of the 14th International Symposium on Information Science (ISI 2015), Zadar, Croatia.

Zaczkiewicz, A. (2016). *Amazon, Wal-Mart Lead Top 25 E-commerce Retail List,* WWD. http://wwd.com/business-news/financial/amazon-walmart-top-ecommerce-retailers-10383750/