Privacy Dashboards: The Impact of the Type of Personal Data and User Control on Trust and Perceived Risk

Eelco Herder
eelcoherder@acm.org
Radboud Universiteit
Institute for Computing and Information Sciences
Nijmegen, The Netherlands

Olaf van Maaren
Radboud Universiteit
Institute for Computing and Information Sciences
Nijmegen, The Netherlands

ABSTRACT

Website owners often collect personal data to, among other things, advertise more efficiently and to analyze and increase sales. They can inform users in various ways about what data they collect and how they process it. This study focuses on the use of privacy dashboards, which are increasingly present on websites, but not very well studied yet. In this study, making use of an experimental webshop and various interviews, we investigate which elements of privacy dashboards increase customers' trust and reduce the perceived privacy risks. The results indicate that the presence of derived data, such as average values or profile classifications, had a more negative impact on trust and perceived privacy risk than inferred data, such as interest probabilities – and both categories had a larger impact than provided and observed, unprocessed user data. Further, it emerges that a greater degree of control leads to a somewhat greater trust. The study confirms the benefits of a privacy dashboard in addition to a privacy policy and provides guidelines on the level of abstraction on which user data should be presented.

CCS CONCEPTS

• **Information systems** → **Personalization**; **Recommender systems**.

KEYWORDS

privacy dashboards, scrutability, gdpr, user control, user trust, perceived risk

ACM Reference Format:
Eelco Herder and Olaf van Maaren. 2020. Privacy Dashboards: The Impact of the Type of Personal Data and User Control on Trust and Perceived Risk. In *Adjunct Proceedings of the 28th ACM Conference on User Modeling, Adaptation and Personalization (UMAP '20 Adjunct), July 14–17, 2020, Genoa, Italy*. ACM, New York, NY, USA, 6 pages. https://doi.org/10.1145/3386392.3399557

1 INTRODUCTION

Many websites have recently implemented so-called privacy dashboards in addition to their existing privacy notices. Well-known examples include Google Dashboard[1] and the Microsoft Privacy Dashboard[2]), but – arguably motivated by the European General Data Protection Regulation (GDPR) – also smaller websites, such as online supermarkets[3] increasingly have privacy dashboards.

According to Privacy Patterns[4], privacy dashboards "provide successive summaries of collected or otherwise processed personal data for a particular user, representing this data in a meaningful way. This can be through demonstrative examples, predictive models, visualizations, or statistics."

Implementing a privacy dashboard does not only provides transparency to the user, but it seems to have more advantages for a company. For example, it shows that a company takes the privacy concerns of a consumer seriously. Addressing these concerns can provide a competitive advantage for the company with regard to the willingness of users to use the service of that company instead of services by providers that do not [18, 26]. In addition, the data seems more reliable when it is checked or modified by the user through a privacy dashboard [27]. These are all arguments that can be used for further take-up of privacy-enhancing technologies, provided that they are backed up by empirical data.

Privacy dashboards can, however, differ considerably from each other. For example, different data types can be displayed, the degree of interaction can differ, and there can be different ways users can access the information. The purpose of this study is to determine what aspects privacy dashboards lead to increased trust, reduced perceived risk and more purchases from an online retail organisation. Particularly, we will investigate the impact of the type of data – directly provided or observed vs derived or inferred data – and the level of control.

The paper is structured as follows. In Section 2, related work is presented. The methodology and materials of the study are explained in Section 3, followed by the results in Section 4. We end the paper with discussion and conclusions in Section 5.

2 BACKGROUND AND RELATED WORK

According to Zimmerman et al [27], "Privacy dashboards are tools to provide data subjects with a clear and easily understandable overview over data a data controller has accumulated about them, and empower data subjects to control processing or usage of that data, as well as future collection of data by the data controller". The degree of control that a privacy dashboard provides can vary. Some privacy dashboards only have a read-only function. Interactive

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
UMAP '20 Adjunct, July 14–17, 2020, Genoa, Italy
© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-7950-2/20/07…$15.00
https://doi.org/10.1145/3386392.3399557

[1] https://myaccount.google.com/dashboard?pli=1
[2] https://account.microsoft.com/privacy/
[3] e.g. https://www.ah.nl/mijn/dashboard
[4] https://privacypatterns.org/patterns/Privacy-dashboard

privacy dashboards offer the user mechanisms of control, but large differences in functionality can be observed.

In this section, we discuss the different types of user data involved, principles of scrutability and the European GDPR, and the relation between perceived risk, trust and intention to buy.

2.1 Provided, Observed, Derived and Inferred Data in E-Commerce

When a consumer makes a purchase online, data is collected from this consumer, which is used for various purposes [3]. An important purpose is identification for service purposes: for product purchases, data such as name, address and payment data is needed [5]. Data is also used for making strategic decisions. For example, an organisation can find out which products are sold at what time and can make a decision whether they will continue to sell these products [6, 20]. In addition, an online retailer can create customer profiles and make choices which customer groups to focus on [7]. Many online retailers also use data to personalize their services [8].

Several studies have shown that web users have different attitudes and concerns regarding data disclosure, largely depending on the nature of the personal data that is disclosed [17]. In this study, we make use of the categories as drawn up by Abrams [1]: provided data, observed data, derived data and inferred data. *Provided data* is data explicitly provided by the user, *observed data* includes (ininterpreted) click-behavior data from which user interests and buying behavior may be derived or inferred.

Derived data is data that is derived from other (e.g. observed) data and thus becomes a new data element related to the individual. As an example, the average amount of money a customer has spent in a webshop might trigger classification into a customer profile such as “City Budget" or “Premium"[5].

Inferred data is the result of a probability-based analytic process aimed at finding correlations between different data points. The result is typically a probability percentage and is therefore less certain than derived data [14]. For example, when 40% of all users who buy meat substitutes turn out to be vegetarian, the probability is 40% that a user is vegetarian when they buy a meat substitute.

2.2 Scrutability, the GDPR and Privacy Dashboards

Privacy dashboards can be seen as an implementation of the concept of *scrutability*, introduced by Judy Kay [15]. Scrutable user models provide insight on which user data is collected, how it is interpreted, how it contributes to adaptation decisions and with parties the data is shared; scrutable user models should also provide users with means to control each step in this process.

Several principles of scrutability have become mandatory in the European Union by means of the GDPR that gives internet users (“data subjects"), among others, the right to know[6] (a) the purposes of the processing, (b) the categories of personal data concerned and (c) the recipients or categories of recipient to whom the personal data have been or will be disclosed. Note that the GDPR seems not to specifically cover derived and inferred data.

[5]https://nos.nl/artikel/2211640-albert-heijn-stigmatiseert-bij-instructie-personeel.html

[6]https://gdpr-info.eu/art-15-gdpr/

Current privacy dashboards typically mainly cover provided and observed data. However, if users expect that organizations do work with derived or inferred data, they expect these data categories to be included as well, as recently shown in a “GDPR reality check" performed on the German loyalty program Payback [4]. Furthermore, which data has been shared with other parties appears to be included very seldomly: for example, only very recently, Facebook included “Off-Facebook Activity" in their privacy settings[7].

2.3 Perceived Risk, Trust and Intention to Buy

Consumers usually base their choices on incomplete information: they simply cannot know everything about an organization. As a result, there is often a degree of risk and uncertainty in an exchange between consumer and retailer. The consumers’ perceived risk largely influences the decision to purchase a product online [13].

Various studies have shown that trust in an (online) retailer can decrease the perceived risks [13, 22]. However, it is difficult to give a good general definition of the concept of trust [10]. Within an economic framework, trust as behaviour is described as a risky advance concession in the expectation of a positive outcome without any explicit contractual security or control measure against opportunistic behaviour [24].

There is also a third variable that can influence the decision to buy a product from a retailer positively: perceived benefits. Kim et al. [16] define perceived benefit as a consumer’s belief about the extent to which he or she will become better off from the online transaction with a certain website.

2.4 Privacy Nudges, Privacy By Design, and Dark Patterns

Acquisti et al [2] investigated user choices regarding privacy and security, as a result of available interaction mechanisms. Among others, they observed that users often acted upon incomplete and asymmetric information. Further, users often make use of heuristics and shortcuts when balancing possible privacy risks and concrete (financial) benefits. They suggested that ‘soft paternalism’ and privacy nudges – such as privacy-friendly default choices, but also education and feedback – would help users in making informed decisions.

Hoepman et al [12] aimed to translate the concept of privacy nudges into concrete privacy design strategies, based on the following principles: minimise, hide, separate, aggregate, inform, control, enforce and demonstrate. They stated the expectation that by providing users direct control over their personal data, they are more likely to correct errors. Particularly, they required data controllers – such as online commerce – to demonstrate compliance with the privacy policy and any applicable legal requirements.

In contrast to nudges for privacy-friendly choices, website owners can also decide to employ so-called *dark patterns* that encourage users to choose less privacy-friendly options. In a survey of 11.000 shopping websites, Mathur et al [19] found several examples on patterns that are covert, deceptive and information hiding in nature. Further, many patterns exploited default patterns and framing effects. Concrete examples include sneaking additional products into

[7]https://thenextweb.com/basics/2020/01/28/how-use-facebooks-off-facebook-activity-tool/

shopping baskets, using language, emotion and visual presentation to steer users away from certain choices.

In summary, the body of related work confirms that in e-commerce several types of personal data are collected, varying from provided and observed data, to derived and inferred information. Privacy dashboards provide users means to obtain insight and control how the data is collected and used, which is required by the European GDPR and is also expected to increase trust and intention to buy. However, privacy dashboards can also be designed in such a way that users unknowingly make less privacy-friendly options. In this study, we investigate user responses to various dashboards that differ in terms of privacy friendliness and user control.

3 METHODOLOGY

A controlled experiment was chosen to find out what data types and what degree of control are related to perceived risk, perceived trust and the intention to purchase. An online webshop has been built for this purpose (called "Jud Fashion"). This fictional webshop sells everyday clothing for women and men.

The online webshop and the survey are made with the help of Wix.com, an online platform for web design, with which websites can be created. The data that has been collected is stored in the Wix databases. Only the researcher had access to these databases. After the research, the data was exported from these databases and the databases were removed from the Wix server. The experiment has been distributed among the researcher's network (friends, social media and an owned website).

A total of 97 participants have been recruited for the online experiment, drawn based on accessibility. Between seven and fifteen participants were randomly associated with each of the nine scenarios described below. In addition, five of the 97 participants have been interviewed. All participants were Dutch and the study was conducted in the Dutch language.

The experiment started with a visit to the experimental website, upon which a pop-up opens with a short explanation about the experiment and a first small survey. During this first survey, participants are asked to enter both their name and e-mail and to make a number of choices between clothing items. The underlying idea is that this *provided data* can be used for determining clothing preferences in order to create a user profile.

Subsequently, participants were asked to take a look at some clothing items that are available in the webshop. This activity delivered some *observed data*: all pages visited on jud-fashion.nl, the browser language, the location of the device and the website where the user comes from.

In the background, the provided and observed data were used for building a user profile – i.e. interest in suits would lead to a 'business' profile, serving as *Derived data. Inferred data* involved probability calculations of the user's gender, income level and to what extent the participant has a conservative clothing style).

Then, the users were asked to visit the "privacy" page. On this page, a participant will find either a privacy dashboard (eight randomly assigned scenarios) or just a privacy policy (one scenario, used as a baseline). The eight scenarios in which the privacy dashboard is displayed differ in the degree of control and in the data types for which insight/control is provided (see Section 2.1) and the previous paragraphs. An overview of these scenarios can be found in Table 1. As can be seen, we measure the effect of *omitting* one data type one at a time, leaving all other data types in the dashboard.

Scenario	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**
Control	*Yes*	*Yes*	*Yes*	*Yes*	*No*	*No*	*No*	*No*
Provided		x	x	x		x	x	x
Observed	x		x	x	x		x	x
Derived	x	x		x	x	x		x
Inferred	x	x	x		x	x	x	

Table 1: The 8 different configurations of privacy dashboards in terms of types of data and level of user control

Participants in the group with *no control* can only inspect their data via the privacy dashboard. Conversely, participants in the group *user control* can modify and delete provided, observed, derived and/or inferred data, depending on the condition that they were assigned to.

In order to ensure that the design of the privacy dashboard is sufficiently realistic and user-friendly as possible, our dashboard – as displayed in Figure 1 is based on a prototype that originated from a study by Raschke et al.[23]. The personal data are displayed vertically in a list. The category of each item is indicated by an icon that provides information about how the item has been processed. In addition, the source is stated, as well as the data on which the item was processed. In the user control condition, three dots are displayed in the top right corner that allows the participant to revise permissions, modify or delete the item.

To prevent the design of the privacy dashboard from having too much influence on the research results, a few people were asked to assess the design of both the webshop and the privacy dashboard prior to the start of the experiment. These persons did not participate in the study.

To measure *online trust*, an item scale from McKnight et al [21] will be used to guarantee construct validity. Featherman and Pavlou [9] conceptualized *perceived privacy risk* as an important part of perceived risk and they defined it as the potential loss of control over personal information. During this study we will use the items from a research done by Kim et al [16].

In this study we look at the *intention to buy* of and do not measure the actual purchases. The items of Kim et al. [16] were chosen, because these items do not presume that an interviewee has already purchased a product on the website or was already familiar with the retailer.

Finally, five semi-structured interviews of about 20 minutes in length were conducted, to find out participants' knowledge and awareness data processing practices, their perceived importance of privacy, and their experiences with privacy dashboards. Atlas.TI v8[8] was used to analyze the interviews. The recorded interviews were first transcribed, codes were then assigned to parts of the interviews and compared with each other, in order to discover patterns and gain an overview of the users' perspective on the use of privacy dashboards and the importance of privacy.

[8]https://atlasti.com/

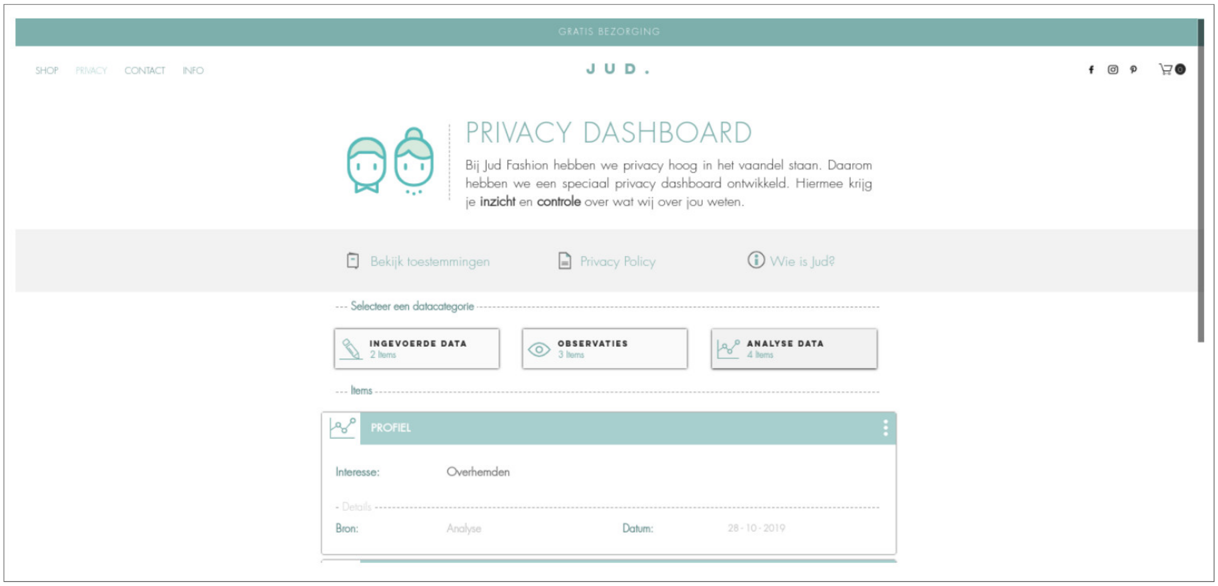

Figure 1: Part of the privacy dashboard used in the study.

4 RESULTS

In this section, we present the results of the experiment. First, we present the results of the surveys filled out by all 97 participants. The reliability of the study was measured with Cronbach's alpha, in order to verify that the internal consistency of each items was equal to or higher than 0.7. Interview results are presented in the form of observations, illustrated by relevant quotes from partiicpants.

4.1 Survey Results

First, we are interested in the effect of the *presence of a privacy dashboard* on consumers' trust. The results show that the variables trust and perceived risk differ significantly between the groups: the perceived privacy risk is larger for the baseline group that has only seen a privacy policy ($\mu = 4.28$ vs $\mu = 3.09$, $p \leq .01$) and the trust is slightly lower for this group as well ($\mu = 3.45$ vs $\mu = 3.80$, $p \leq .01$). However, the results do not show a significant difference in the intention to purchase between the groups.

Further, we investigated the effect of the *degree of control* on the data through a Privacy Dashboard on the the perceived privacy risk. The results indicate that the online intention to purchase and the perceived privacy risk do not differ significantly between the groups with user control and without control. There is, however, a weakly significant difference in trust ($\mu = 3.71$ for the group without control vs $\mu = 3.92$ for the group with control, $p = 0.084$).

Regarding the *user data types* (provided, observed, derived, inferred), we analyze the effect of omitting one of these categories one at a time. An Anova test shows a significant difference in user's trust between the different groups ($F(3.75) = 3.621, p \leq .05$). The Bonferroni post hoc test shows that the only significant differences can be found between the groups where the *provided* was omitted ($\mu = 3.61$) and where the *derived* was omitted ($\mu = 4.111$), the overal mean was $\mu = 3.824$. Although not sufficiently significant ($p = .069$), the positive effect on trust is higher for derived data than for inferred data.

Similarly, also for *perceived risk* a significant difference between groups was found ($F(3.75) = 3.812, p \leq .05$). Again, when the provided data is omitted, the perceived privacy risk is larger ($\mu = 3.222$) than when the derived data is omitted ($\mu = 2.48$), compared to the overall mean $\mu = 3.07$. As with trust, also here the perceived risk is lower when derived data is omitted than when inferred data is omitted. For *online purchase intention*, only the omission of derived leads to a significant difference.

As can be expected, a significant relationship was found between the perceived privacy risk and trust ($F(1.77) = 62.21, p \leq .01$), with perceived privacy risk an explanatory variable for consumers' trust. No significant relationship between privacy risk and online purchase intention was found, in contrast to the relation between trust and online purchase intention ($F(2.76) = 21.37, p \leq .01$).

In summary, the results show that the presence of a privacy dashboard leads to a greater degree of trust and a smaller perceived privacy risk compared to the display of a privacy policy only. A greater degree of control leads to a significantly greater trust in the webshop than the lack of this control. Finally, regarding the data types, the omission of derived data generally has a positive effect on trust, perceived risk and online purchase intention, while omitting provided data (thus leaving all other data types in the dashboard) has a negative effect on trust and perceived risk on average.

4.2 Interview Results

Prior to the experiment, the participants were asked how they view data processing. Four out of five respondents regularly reflect on the fact that data is collected when they are online. According to the respondents, the main purpose of the organizations for collecting data is that the companies themselves benefit from it. They all admit that they actually don't know exactly what they collect:

[If they estimate] that someone has a below average income and he/she buys [an expensive] product, then they are not going to warn him that he should not do that. It is therefore self-interest. (Respondent 3 – Translated)

Only one of the respondents had previously tried to gain insight into his/her data. It concerned data from Facebook and it was perceived as very unclear to this test subject. None of them had seen a privacy dashboard before.

Inferred data is not mentioned by any of the respondents before the start of the experiment . Derived data were partially mentioned: only examples of computational derived data, e.g. averages, were given. This corresponds to the article by Abrams [1], which states that the awareness of inferred and notational derived data is low.

During and after the experiment the respondents were asked what they think of the data that has been collected. Almost all of them responded with surprise to the inferred data. Some responded that they could have known this, while others were rather surprised why the webshops wants to find out something like ‘income’ and ‘gender’. So, in some cases, the privacy dashboard raised rather than answered questions among respondents.

Respondents appear to find privacy not the most important factor when deciding to buy a product somewhere. In some cases, store reputation or reliable delivery may be more important. Some say that online privacy and does not differ much from the privacy that you have when walking into a physical store:

When people on the street see me walking into the Livera, they will also see that I am going to buy underpants. I can’t really lie awake about it. (Respondent 4 – Translated)

Although the data collection practices did surprise respondents, four out of five people said that they would not change their behavior as a result. Four out of five also indicate that the presence of a privacy dashboard is not necessarily a reason to buy their products from a webshop; other factors such as convenience play a more important role. This largely corresponds to the results from the experiment, where the differences in intention to buy are not significant in most comparisons.

5 DISCUSSION

Previous studies have shown that when organisations provide tools that provide insight into the collected data, the customers’ trust will increase. This study confirmed that the participants who were offered a privacy dashboard indeed had a slightly higher trust in the webshop and a lower perceived privacy risk compared to those who were only provided a privacy policy.

In this study, the effect of *omitting* data categories in a privacy dashboard was measured. This may appear counter-intuitive from a user modeling point of view, but is arguably more in line with the way online data collection and derivation practices take place: as a first step, as much data is collected and analyzed as appears reasonable and needed. Without incentives for data minimization, companies often have the tendency to keep as much as possible. Analysis of user interaction with and responses to privacy dashboards, as done in this study, provides valuable information on which data collection practices are more acceptable than others.

A limitation of this study is the large number of dashboard designs compared to the number of participants in the controlled studies. As a results, the effect sizes of our findings are quite small. However, our intention was not to investigate differences between two concrete designs, but to explore a continuous design space in terms of data types (provided, observed, derived, inferred) and user control, taking availability of participants and time constraints into account.

Our results showed that omitting derived data has a positive effect on trust and perceived risk; conversely, it can be concluded that introducing derived data into a privacy dashboard will have a negative effect. By contrast and as expected, users were not very concerned regarding data they provided themselves.

The negative effect of derived and inferred data on trust and perceived risk is likely caused by users perceiving such inferrals as unexpected and unpredicatable: the ‘system’ is not expected or supposed to ‘know’ them [11] or the inferrals might be considered wrong. However, as discussed in the related work, if users have reason to believe that data is being used for profiling or inferring interests, they might consider the dashboard incomplete and perhaps ‘not fair enough’ if these categories are not available [4].

The observation that inferred data (such as the probability that one is a vegetarian) is considered less intrusive than derived data (e.g. being categorized as a vegetarian) may have interesting implications for the explanation of recommendations [25]. Current recommendation techniques typically make use of inferred data after all. Instead of ‘translating them back’ into comprehensive and comprehensible user profiles, user acceptance may be higher if they are just being kept inferences.

In the light of the GDPR, a natural consequence of *not* including data from a certain category in a privacy dashboard would be to stop collecting or inferring these data. Naturally, other - arguably less ethical - solutions exist, including collecting and inferring the data nevertheless while pretending that the privacy dashboard provides full transparency; or to derive or to infer certain characteristics or probabilities on the fly during the recommendation process, assuming that this eliminates the need for including it in the privacy dashboard.

To return to the concept of scrutability: results from this study suggest that insight in data collection practices has a larger effect on trust and perceived risk than actual control. No significant effect on the intention to buy was measured, which may be interpreted as that websites owners would have no financial incentives for investing in a privacy dashboard. However, we have also shown that indirect effects via trust and perceived risk are present.

6 CONCLUSION AND FUTURE DIRECTIONS

This study has shown the positive impact of privacy dashboards on trust and perceived risk: already providing insight in data collection practices increases both, adding user control increases both values even more. We particularly focused on the different levels of abstraction of user data, and found that derived and inferred data was considered as more problematic than provided and observed data.

Future research directions include the comprehensibility and user acceptance of privacy dashboards with an increasing amount and diversity of data and data types. Particular attention should be paid to user acceptance of the presentation of inferred and derived data.

Another dimension to investigate is the delivery mode of privacy dashboards: currently, users often must actively search and visit these dashboards; It would be interesting to investigate whether pro-actively inviting users to visit the dashboard on a regular basis has a positive or negative impact on perceived risk or trust.

REFERENCES

[1] Martin Abrams. 2014. The origins of personal data and its implications for governance. *Available at SSRN 2510927* (2014).

[2] Alessandro Acquisti, Laura Brandimarte, and George Loewenstein. 2015. Privacy and human behavior in the age of information. *Science* 347, 6221 (2015), 509–514.

[3] Shahriar Akter and Samuel Fosso Wamba. 2016. Big data analytics in E-commerce: a systematic review and agenda for future research. *Electronic Markets* 26, 2 (2016), 173–194.

[4] Fatemeh Alizadeh, Timo Jakobi, Jens Boldt, and Gunnar Stevens. 2019. GDPR-Reality Check on the Right to Access Data: Claiming and Investigating Personally Identifiable Data from Companies. In *Proceedings of Mensch und Computer 2019*. 811–814.

[5] Suhail Ansari, Ron Kohavi, Llew Mason, and Zijian Zheng. 2001. Integrating e-commerce and data mining: Architecture and challenges. In *Proceedings 2001 IEEE International Conference on Data Mining*. 27–34.

[6] Brad Brown, Michael Chui, and James Manyika. 2011. Are you ready for the era of 'big data'. *McKinsey Quarterly* 4, 1 (2011), 24–35.

[7] Sirish Chandrasekaran, Robert Levin, Harry Patel, and Roger Roberts. 2013. Winning with IT in consumer packaged goods: Seven trends transforming the role of the CIO. *McKinsey & Company* (2013), 1–8.

[8] Thomas H Davenport et al. 2006. Competing on analytics. *Harvard business review* 84, 1 (2006), 98.

[9] Mauricio S Featherman and Paul A Pavlou. 2003. Predicting e-services adoption: a perceived risk facets perspective. *International journal of human-computer studies* 59, 4 (2003), 451–474.

[10] Sonja Grabner-Kraeuter. 2002. The role of consumers' trust in online-shopping. *Journal of Business Ethics* 39, 1-2 (2002), 43–50.

[11] Eelco Herder and Boping Zhang. 2019. Unexpected and Unpredictable: Factors That Make Personalized Advertisements Creepy. In *Proceedings of the 23rd International Workshop on Personalization and Recommendation on the Web and Beyond*. 1–6.

[12] Jaap-Henk Hoepman. 2014. Privacy design strategies. In *IFIP International Information Security Conference*. Springer, 446–459.

[13] Ilyoo B Hong and Hoon S Cha. 2013. The mediating role of consumer trust in an online merchant in predicting purchase intention. *International Journal of Information Management* 33, 6 (2013), 927–939.

[14] Information Commissioner's Office (ICO). 2017. *Big data, artificial intelligence, machine learning and data protection.* https://ico.org.uk/media/for-organisations/documents/2013559/big-data-ai-ml-and-data-protection.pdf

[15] Judy Kay. 2006. Scrutable adaptation: Because we can and must. In *International Conference on Adaptive Hypermedia and Adaptive Web-Based Systems*. 11–19.

[16] Dan J Kim, Donald L Ferrin, and H Raghav Rao. 2008. A trust-based consumer decision-making model in electronic commerce: The role of trust, perceived risk, and their antecedents. *Decision support systems* 44, 2 (2008), 544–564.

[17] Alfred Kobsa. 2007. Privacy-enhanced web personalization. In *The adaptive web*. 628–670.

[18] Naresh K Malhotra, Sung S Kim, and James Agarwal. 2004. Internet users' information privacy concerns (IUIPC): The construct, the scale, and a causal model. *Information systems research* 15, 4 (2004), 336–355.

[19] Arunesh Mathur, Gunes Acar, Michael J Friedman, Elena Lucherini, Jonathan Mayer, Marshini Chetty, and Arvind Narayanan. 2019. Dark patterns at scale: Findings from a crawl of 11K shopping websites. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019), 1–32.

[20] Andrew McAfee, Erik Brynjolfsson, Thomas H Davenport, DJ Patil, and Dominic Barton. 2012. Big data: the management revolution. *Harvard business review* 90, 10 (2012), 60–68.

[21] D Harrison McKnight, Vivek Choudhury, and Charles Kacmar. 2002. Developing and validating trust measures for e-commerce: An integrative typology. *Information systems research* 13, 3 (2002), 334–359.

[22] Paul A Pavlou. 2003. Consumer acceptance of electronic commerce: Integrating trust and risk with the technology acceptance model. *International journal of electronic commerce* 7, 3 (2003), 101–134.

[23] Philip Raschke, Axel Küpper, Olha Drozd, and Sabrina Kirrane. 2017. Designing a GDPR-compliant and usable privacy dashboard. In *IFIP International Summer School on Privacy and Identity Management*. 221–236.

[24] Tanja Ripperger. 2003. *Ökonomik des Vertrauens: Analyse eines Organisationsprinzips*. Vol. 101. Mohr Siebeck.

[25] Nava Tintarev and Judith Masthoff. 2007. A survey of explanations in recommender systems. In *2007 IEEE 23rd international conference on data engineering workshop*. IEEE, 801–810.

[26] Heng Xu, Tamara Dinev, H Jeff Smith, and Paul Hart. 2008. Examining the formation of individual's privacy concerns: Toward an integrative view. *ICIS 2008 proceedings* (2008), 6.

[27] Christian Zimmermann, Rafael Accorsi, and Günter Müller. 2014. Privacy dashboards: reconciling data-driven business models and privacy. In *2014 Ninth International Conference on Availability, Reliability and Security*. 152–157.