

You ‘Might’ Be Affected: An Empirical Analysis of Readability and Usability Issues in Data Breach Notifications

Yixin Zou, Shawn Danino, Kaiwen Sun, Florian Schaub

School of Information
University of Michigan
Ann Arbor, MI, USA

{yixinz,danino,kwsun,fschaub}@umich.edu

ABSTRACT

Data breaches place affected individuals at significant risk of identity theft. Yet, prior studies have shown that many consumers do not take protective actions after receiving a data breach notification from a company. We analyzed 161 data breach notifications sent to consumers with respect to their readability, structure, risk communication, and presentation of potential actions. We find that notifications are long and require advanced reading skills. Many companies downplay or obscure the likelihood of the receiver being affected by the breach and associated risks. Moreover, potential actions and offered compensations are frequently described in lengthy paragraphs instead of clearly listed. Little information is provided regarding an action’s urgency and effectiveness; little guidance is provided on which actions to prioritize. Based on our findings, we provide recommendations for designing more usable and informative data breach notifications that could help consumers better mitigate the consequences of being affected by a data breach.

CCS CONCEPTS

• **Security and privacy** → **Usability in security and privacy**; *Social aspects of security and privacy*; • **Human-centered computing** → *Empirical studies in HCI*.

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 the author(s) 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.
CHI 2019, May 4–9, 2019, Glasgow, Scotland UK

© 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-5970-2/19/05...\$15.00

<https://doi.org/10.1145/3290605.3300424>

KEYWORDS

Security; Privacy; Data breach; Usability; Notice design; HCI design; Content analysis.

ACM Reference Format:

Yixin Zou, Shawn Danino, Kaiwen Sun, Florian Schaub. 2019. You ‘Might’ Be Affected: An Empirical Analysis of Readability and Usability Issues in Data Breach Notifications. In *CHI Conference on Human Factors in Computing Systems Proceedings (CHI 2019), May 4–9, 2019, Glasgow, Scotland UK*. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3290605.3300424>

1 INTRODUCTION

Data breaches – security violations that compromise sensitive, protected, or confidential data of individuals [62] – have become increasingly common in recent years with significant repercussions for consumers. In 2017, 853 data breaches occurred and compromised 2.05 billion records in total, including consumers’ names, contact information, account numbers, credit card details, social security numbers, shopping and purchasing records, social media posts and messages, or even health records [63]. One of the major consequences of data breaches, identify theft, results in average financial damages of over \$1,000 per victim [32], not to mention the psychological trauma many victims experience during the identity recovery process [35].

To mitigate the severe consequences of data breaches, many countries have passed data breach notification laws, requiring companies to notify affected consumers. The key purpose is to inform consumers of the risks and motivate them to take protective actions, as well as urge affected companies to pursue better data security practices [1]. In the United States, data breach notification laws are industry and state-specific. All 50 U.S. states have mandated data breach notifications be sent if consumers’ personally identifiable information (PII) is involved [45]. Yet requirements vary significantly regarding how many state residents are affected before a notification has to be sent, and how soon the notification should be delivered to consumers [81].

While Romanosky et al.'s analysis in 2011 [70] showed that data breach notification laws reduced identity theft by 6.1% in the U.S., more recent studies reported consumer inaction following data breaches, suggesting notifications as an ineffective mechanism. In Ponemon Institute's 2014 national survey, 32% of respondents reported their reaction to a data breach notification is to "ignore it and do nothing" [61]. In 2017, the Equifax data breach compromised the records of almost half of the US population [26]; however, the adoption rate of credit freezes, a strong method to prevent new lines of credit being opened, was lower than 1% 10 days after the breach [16]. Thus, while data breach notifications are required by law, they are seemingly inadequate in motivating consumers to make use of available protective measures.

To shed light on the effectiveness of data breach notifications and potential directions for improvements, we conducted a content analysis of 161 notifications sent by companies to U.S. consumers between January and June 2018. We analyzed their readability, structure, risk communication, and presentation of recommended actions. Most analyzed notifications were lengthy and would be difficult to understand for the general public. They varied significantly in the format of headings and the incident description's specificity. Consequences and risks of the data breach were usually obfuscated by hedge terms such as 'potentially' and 'may,' as well as a 'no evidence' statement (e.g., "we have found no evidence indicating that your breached personal data has been misused"). Although most notifications provided a detailed explanation of recommended actions, those actions are typically buried in long paragraphs with little to no guidance regarding their effectiveness or urgency, making it difficult for the reader to navigate and prioritize listed actions. Based on our findings, we provide design and public policy recommendations for improving data breach notifications.

2 BACKGROUND

Data breach notification requirements vary widely across jurisdictions. In the European Union (EU), the General Data Protection Regulation (GDPR) requires data breach notifications to both the supervisory authority within 72 hours and affected European consumers 'without undue delay' [17]. The notification has to include a clear and plain description of the breach's nature and recommended protective measures [17]. Substantial fines for non-compliance with GDPR [48] pose incentives for companies to disclose mandated information and establish stronger security incident procedures and training [69].

In the United States, no equivalent federal data breach notification law exists [53]. Instead, a patchwork of sector-specific federal laws outline various requirements. For example, the Graham-Leach-Bliley Act (GLBA) [89] regulates data

breach notifications for financial institutions, prescribing several mandatory elements to be included [8, 84]. The Health Insurance Portability and Accountability Act (HIPAA) [86] establishes a 60-day notification deadline for data breaches that compromised consumers' health information, via a mailed letter written in plain language [36]. In addition to these sectoral laws, all 50 U.S. states have enacted their own data breach notification laws. These state laws vary substantially in stringency, resulting in inconsistent notification requirements among states, as well as different definitions of Personally Identifiable Information (PII) [42, 45], which if breached requires a notification. For example, California and Maryland consider medical information PII; other states, like New Hampshire and Iowa, do not. California is one of the few states that sets clear expectations about the structure and formatting of breach notifications in their law with a template [83]. The template not only includes specific wording of the title ("Notice of Data Breach") and headings, but also requires them to be conspicuously displayed. Additionally, the California law has a "plain language" requirement similar to GDPR. Arizona, Illinois, Oregon, New York, Vermont, and West Virginia also provide templates for companies to refer to when drafting data breach notifications, but with less strict and detailed structure and formatting requirements. Moreover, California and Connecticut are the two states requiring companies to offer identity theft protection services to help consumers deal with potential harms; similar legislation is pending in the State of New York [55, 94].

In our study, we analyzed data breach notifications from the State of Maryland, due to the comprehensive records in the Maryland Attorney General's public database of data breach notifications¹. Maryland's Personal Information Protection Act (PIPA) [29] defines that PII encompasses traditional types of PII (e.g., name and Social Security number), government-issued IDs, and health information (effective since January 2018) [49]. PIPA requires data breach notifications to be sent to individual consumers 'as soon as possible,' and within 45 days upon discovery of a data breach [49]. According to PIPA, a breach notification must specify the types of breached information, as well as offer the contact information for different entities such as major credit reporting agencies (for placing credit freeze etc.), the Federal Trade Commission (FTC) (for obtaining more information about identity theft protection), and the Maryland Attorney General (for reporting identity theft incidents) [50].

¹<http://www.marylandattorneygeneral.gov/Pages/IdentityTheft/breachnotices.aspx>

3 RELATED WORK

Our research builds on prior work on data breaches; in particular, consumers' response to data breaches, empirical analysis on data breach notifications, and the design and effectiveness of security and privacy notices.

Consumer Reactions to Data Breaches

Prior work suggests that consumers do not take adequate protective action when affected by a data breach. In Ponemon Institute's 2014 national survey, the concern of being a victim of identity theft increased by 21% following a breach, yet 32% of respondents reported their reaction to a data breach notification is to "ignore it and do nothing" [61]. Similarly, Gemalto's 2017 worldwide survey showed identity theft concern from two thirds of respondents; nevertheless, 56% continued using the same password for multiple accounts, and 41% did not adopt two-factor authentication when provided [28]. An exception is RAND's 2016 U.S. national survey, in which 62% reported accepting offers of free credit monitoring — a higher but still not satisfactory number [1]. For specific cases, surveys following Target's 2013 data breach showed that there was no significant decline in debit card usage in the year after, even though the perceived security of personal information associated with debit cards dropped [33]. In a qualitative study on Equifax's 2017 data breach [98], most participants expressed concerns about identity theft and privacy invasion, yet more than half did not take any protective measures. Together, these studies point to a dissonance between consumers' concerns and behavior following data breaches, which is reminiscent of the privacy paradox [57]: worries about identity theft or privacy are not reflected in people's behaviors.

Few studies have examined the reasons behind consumers' inaction to data breaches. Zou et al. [98] found that optimism bias [78] (i.e., underestimating one's likelihood of victimization), insufficient knowledge of available protective measures, and a general tendency to delay action until harm has occurred, dissuaded consumers to take actions after the Equifax breach. Mikhed and Vogan [52] found that clear evidence of being affected by a breach encouraged consumers to sign up for fraud protection services. Furthermore, Kude et al. [43], studying Target's data breach, found that whether a provided compensation (i.e., discount and free credit monitoring) is perceived as adequate was largely shaped by consumers' personality traits, as well as social influence (i.e., discussion with friends and peers).

Data Breach Notification Analysis

Companies ubiquitously send data breach notifications to affected consumers [61]. However, notifications from companies are usually not the first thing making consumers aware

of large-scale breaches, possibly due to the fact that most companies have to notify consumers via mailed letters as required by laws. In RAND's survey [1], 44% of respondents had already heard about a breach through other channels before receiving a notification from the company. Das et al. [20] further revealed that these channels were primarily news articles, television news, social media, and personal contacts [20].

So far, there has been little research on the actual content, language, and structure of consumer data breach notifications. Jenkins et al. [38] found that visual elements (e.g., special formatting such as italics, bold, underlining of subject lines), while used rarely (in less than 30% of their analyzed notifications), contributed to the restoration of the affected company's reputation in a follow-up experiment. In Veltsos' analysis of notification templates [92], the claim 'lost data might not be used at all' appeared in 2 out of 13 templates to 'soften the bad news', which, according to the author, does not help affected consumers overcome optimism bias (i.e., "I am the lucky one not affected when a crisis comes and affects so many people") [78] and rational ignorance (i.e., "I do not have the time and effort to look into this since the perceived benefit is so small") [22]. In Zou et al's study [98], participants complained that the use of ambiguous hedge terms (e.g., "your personal information may have been impacted") in Equifax's notification confused them and made them wonder whether they were truly affected. Golla et al. [31] examined real-world notifications for password breaches particularly, and found that most notifications, while successfully raising participants' concerns, did not lead to intentions to change compromised passwords and other secure practices that would protect them from future password-reuse attacks.

Perhaps the most recent and relevant study to our work is Bisogni's analysis of 445 data breach notifications issued in 2014 [10]. This paper assessed the presence of mandatory elements, clarity of breach description, communication tone in depicting possible consequences, and the affected company's openness to interact with consumers [10]. Due to the identified inconsistencies, the paper concluded that a U.S. federal data breach notification law is needed to standardize the timing and mandatory elements of notifications, which are both crucial for informing consumers of potential risks [10]. We expand on and complement Bisogni's study by (1) analyzing more recent data breach notifications and (2) focusing on readability and usability issues.

Security and Privacy Notice Design

Insights on how data breach notifications should be designed can be drawn from research on the design and effectiveness of privacy and security notices. "Notice and choice," as the predominant mechanism to protect consumers' privacy [18],

takes the approach to present privacy practices and associated risks to end-users (e.g., in privacy policies and terms of conditions), and offer the choice of acceptance or denial (e.g., choose whether or not to click on an “I agree” button) [79]. Efforts have been made to design more readable and transparent privacy notices as well as more salient and easy-to-use choices [18, 30, 74, 75].

Nevertheless, the “notice and choice” paradigm results in poor technical solutions to communicate risks and provide privacy protection in practice [14, 18, 79]. Privacy policies, while being the primary tools to inform users about companies’ data practices [68], are time-consuming to read [51] and complicated with a lot of jargon [39, 47, 64], leading to poor comprehension [91] and little changes in users’ actual privacy practices [71]. These policies vary significantly across companies, do not offer users sufficient choice, and sometimes make self-contradictory statements [19]. Mismatches also occur between users’ expectations and the company’s actual practices [65], leaving users exposed to unanticipated risks such as not knowing certain types of data being collected and shared. Numerous problems also exist for security warnings: users often ignore security indicators like the HTTPS icon in a browser’s address bar [21], or develop incomplete and inaccurate mental models of risks [11]. Users’ adherence to a security warning may be further exacerbated by poor visual designs and inadequate consideration of fatigue effects regarding the warning, as well as users’ prior experiences with the site [5, 6, 23].

To address aforementioned design issues and account for the complexity and nuance in decision-making processes [3], nudges, a concept from behavioral economics [82], have been integrated into privacy and security notices in a variety of contexts ranging from mobile application permissions [7, 97], online disclosure [59, 72, 95], online shopping [73, 88], to password management [25, 90] and computer security warning design [12, 93]. A daily nudge on mobile apps’ access of location data, for example, significantly raised users’ awareness and incentivized them to reassess and restrict their Android app permissions [7]. Providing reminders about the audience of disclosed content [95], or adjusting the framing of the notice [4, 72], could prompt users to be more careful about their online disclosure. Having clear and compact privacy information and notices on shopping interfaces [41, 88], or emphasizing that this product is targeted to the particular user [73], could shape users’ purchase intentions profoundly. Visual elements such as attractors in software installation dialogues [12, 23] and highlights of domain URLs in phishing warnings [60, 93] have been used to navigate users towards more cautious decisions in computer security.

Ultimately, these studies demonstrate the impact of design on individuals’ privacy and security decision making,

with potential implications for designing data breach notifications: a good notice should not only be concise and easy to understand with little ambiguity, but also clearly communicate risks and create strong incentives for recommended protective actions to be taken [2, 74, 75].

4 METHOD

Our discussion of related work shows that data breach notifications do not appear to be effective at spurring protective actions. We focus on examining the readability and usability issues of these notifications, as respective issues may affect consumers’ comprehension of risks associated with a data breach and available protective measures, which is a fundamental step for taking actions. More specifically, we analyzed 161 data breach notifications sent to consumers in the first half of 2018, which we obtained from the Maryland Attorney General’s public database of consumer data breach notifications.

Data Collection

Many U.S. state data breach notification laws require companies to submit data breach notifications sent to consumers to a state’s Attorney General office when the breach affects the state’s residents. To date, California, Iowa, Maryland, New Hampshire, and Vermont make these notifications public. Of the five states, Maryland’s database includes the largest number of breach notifications, indicating the possibility that their records may cover the widest range of data breaches. It also provides additional useful metadata such as the cause of a breach and types of information compromised.

From Maryland’s database, we downloaded all data breach notifications on record from January 1st to June 30th, 2018, in order to narrow the scope of our analysis while ensuring data recency. In total, we obtained 548 data breach notifications. We cleaned the dataset by removing duplicates and entries that did not include notifications to consumers (i.e., some files only included letters reporting the data breach to the Maryland AG, or a website announcement). After filtering, 326 data breach notifications remained. We then randomly selected 161 (~50%) from them to analyze. Appendix A in auxiliary materials provides a full list of our analyzed notifications.

Sample

The 161 notifications in our sample came from 159 unique companies. 154 of them (96%) were mailed letters; 4 were delivered via email, 1 company also sent in-app messages and push notifications to consumers using their app. 14 companies included multiple versions of notifications in their uploaded files. We analyzed the version with a potentially larger audience, e.g., we analyzed notifications delivered to adults instead of guardians of affected minors, and general

consumers instead of company employees. For versions differing in the types of breached information, we analyzed the first one by default.

Personal information, such as name (96%), social security number (52%), and address (51%), was the most commonly breached type in our sample. Financial (e.g., bank account numbers, credit card details) and health-related information (e.g., medical history, medications, health insurance information) were affected in 30 (31%) and 17 (10%) breaches, respectively. We also cross-referenced our sample with the Privacy Rights Clearinghouse’s data breach database (PRC) [85] to understand the overall magnitude of our analyzed breaches. PRC recorded 606 data breaches between January and June 2018, 56 (9%) of which appeared in our sample. According to the PRC data, our sample exposed 151.93 million records across the United States, which constituted 18.5% of the total exposed records coming from all data breaches listed by PRC (820.93 million) for this time frame. Of the 151.93 million records, 150 million were exposed in one breach (Under Armour); 9 breaches (16%) exposed over 10,000 records, and 46 (84%) breaches exposed over 100 records.

Data Analysis

We analyzed data breach notifications with respect to their readability and structure, risk communication, and presentation of recommended actions, using both quantitative and qualitative methods.

Quantitative analysis. Our quantitative analysis focuses on readability, which show to what extent a particular breach notification is comprehensible by the general public, who are typically the recipient. Two metrics we used, the Flesch Reading Ease Score (FRES) and the Flesch Grade Level (FGL) [24], are calculated based on the sentence length and word length of the text. We also looked into the Gunning Fog index (FOG) [34], another grade-level-based metric that factors in complex words (those containing three or more syllables). Additionally, we included statistics related to text characteristics, such as word count and sentence count, to assess notification length and estimate reading time. We used readable.io, a professional online text analysis service, for this qualitative analysis.

Qualitative analysis. We iteratively developed a codebook to assess (1) structure and formatting, (2) risk communication and (3) presentation of recommended actions. One researcher went through all notifications and developed an initial codebook using thematic coding and affinity diagramming [44]. Three members of the research team then independently analyzed a subset of 20 notifications (12.4%) randomly sampled from the dataset, reconciled codes and revised the codebook, eventually reaching good inter-coder reliability (Fleiss’

$\kappa=.75$). The final codebook (see Appendix B in auxiliary materials) has 9 categories (e.g., risk communication), 38 codes (e.g., “whether breached information was misused”), and 136 sub-codes (e.g., “absolutely”, “maybe”, “no”, “no evidence”, and “other”). The researchers then split the 161 data breach notifications and coded them independently using the final codebook.

5 RESULTS

Our analysis shows that current data breach notifications suffer from severe readability issues. Most of the analyzed notifications followed a similar structure, yet their style, length, and content specificity varied considerably. Furthermore, companies downplayed the severity and consequences resulting from a data breach. Even though all notifications recommended protective measures, there was little guidance on how consumers should prioritize among them.

Readability and Structure

A data breach notification should use clear and conspicuous language and an accessible format, in order to help the recipient quickly determine risks stemming from the breach and what actions to take. We analyzed readability, estimated reading time, and use of structural headings for each notification.

Advanced reading skills required. Severe readability issues surfaced from the analysis. The Flesch Reading Ease Score (FRES) [24] evaluates texts on a 0-100 point scale, with higher scores indicating more easy-to-read texts. The median of our sample’s FRES was 46.70 (Mean = 46.88, SD = 6.46). Figure 1 shows our sample’s FRES distribution mapped onto Flesch’s 7-level ranking system [37] from “very difficult” to “very easy.” 115 (72%) notifications fell into the difficult range (30-49); 43 (25%) were ranked as “fairly difficult” (50-59). This means that 97% of the notifications were fairly difficult or difficult to read. Conversely, only one notification was ranked “easy” and one “standard”, and there was no notification rated as “fairly easy” or “very easy.”

Converting the FRES to the Flesch-Kincaid Grade Level (FGL) [24], our sample’s median FGL was 10.0 (Mean=10.02, SD=1.18), i.e., reading abilities of at least a 10th grader are required to be able to understand half of the analyzed notifications. The FGL scores ranged from 6.4 (first-grade level) to 16 (graduate degree required); 75% had a FGL score higher than 9.4. As reference, prior literacy research suggests that materials addressed to the general public should aim for a junior-high reading level (i.e., 7 to 9) [37]. Using Gunning’s Fog index [34], the result was even poorer with a median of 11.6 (Mean=11.55, SD=1.33). This indicates that the existence of long words and jargon aggravated the readability of these notifications. Essentially, most analyzed notifications would

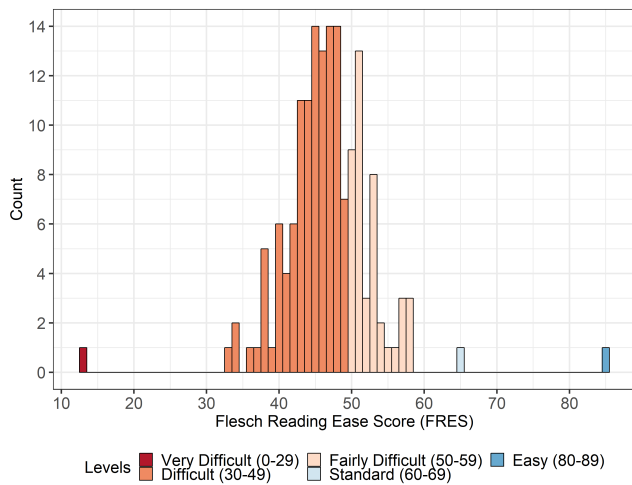


Figure 1: Distribution of the Flesch Reading Ease Score (FRES), mapped onto Flesch’s 7-level ranking system.

likely not meet “plain language” requirements in California’s data breach law or the GDPR – average readers will struggle to understand these breach notifications.

High estimated reading time. We further counted the words and sentences of each notification and used them to estimate the required reading time. The median word count was 1,575 words (Mean=1,539, SD=644), ranging from 213 to 3,414 words (or 7 pages of text). Most notices fell into the 1,000–2,000 word range (highest first quartile value: 1,130; highest third quartile value 1,845). The sentence count distribution also showed wide variance, with a median of 115.0 sentences (Mean=116.39, SD=48.83).

Following McDonald and Cranor’s [51] methodology for estimating the reading time of privacy policies, we assumed a reading speed of 250 words per minute, which is the average reading rate for people with a high school education [13]. The estimated required time to skim a data breach notification thus ranged from .85 to 13.66 minutes (Median=6.3, Mean=6.16, SD=2.57). Although this is a substantially shorter read than privacy policies (which can take upwards of 18 minutes, sometimes hours, to read [46, 51]), in today’s context where people are faced with an overwhelming amount of information, a 6-7 minute anticipated reading time, paired with the need for advanced reading skills, creates a considerable burden for consumers.

Structural headings are common and consistent. Headings structure text and guide a reader’s attention, helping them to quickly identify key information in long text. 106 (67%) notifications used headings to separate their main text into sections. Among them, 72 (68%) put the heading in a separate

What Happened?

On or about March 14, 2018, Mumm Napa and Kenwood Vineyards became the victim of a malicious cyberattack, which compromised a single employee’s account credentials. The account included a number of emails that contained personal information. We currently have no reason to believe that any Mumm Napa systems were compromised beyond limited information contained in the employee’s email account. Upon discovery of the compromise, we promptly acted to secure the compromised account and investigate the incident.

(a) Heading in separate line

What Happened?	A former employee may have accessed consumer data during her employment other than for the purposes of carrying out her assigned duties, during the time period between September, 2017 and February 2018.
----------------	------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

(b) Heading in table

What Happened? On or about August 3, 2017, certain Broward College employees received a spam phishing email to their Broward College email accounts. The phishing email contained a link that asked the employee to enter their Broward College log-in credentials. On Friday, August 18, 2017, Broward College learned that certain employees had clicked on the link in the email and provided their credentials. Broward College identified and corrected the issue by contacting all potentially affected employees and ensuring that all passwords were changed. Broward College also immediately initiated an investigation, with the assistance of a third-party forensic investigator, to determine what personal information, if any, was subject to unauthorized access or acquisition.

(c) Heading in paragraph’s first line

We are writing to inform you that we recently discovered that on or around March 31, 2018 an employee file was infected with ransomware. As a result, your personal information, including your name, address, and Social Security Number may have been accessed. At this time, there is no evidence that this information has been removed from our premises. Further, rather than paying the ransom to the attackers, we recovered files from our backup copies.

(d) No heading (plain text)

Figure 2: Examples of data breach notifications using structural headings when introducing “what happened” in the breach.

line; 34 (32%) included the heading in a paragraph’s first line, reducing its salience (see Figure 2). Only 2 used the table format recommended by the California law. Interestingly, even though we analyzed data breach notifications from Maryland, among the 106 notifications with headings, 100 (94%) followed California’s wording and order requirements: “What Happened,” “What Information Was Involved,” “What We Are Doing,” “What You Can Do,” and “For More Information.” This suggests that unlike privacy policies, data breach notifications generally follow a consistent structure, thus facilitating the learning of that structure over time [58]. However, the disparities regarding heading formatting, paired with poor readability, indicates that inconsistency still remains on the content level.

Risk Communication

Risk communication for data breaches requires companies to explain the situation clearly and openly acknowledge negative consequences [92]. Maryland law [29] requires the types of compromised information to be included in the breach notification, but does not mandate other elements or how the incident should be described.

Varied specificity regarding cause and compromised data. What information was affected by a breach, as required by Maryland law, was described in 152 (94%) notifications. Similarly, the cause of the breach was specified in 150 (93%) notifications. However, notifications varied in their specificity.

Most only listed categories of exposed information generically (e.g., “the email consisted of your name, address, date of birth, account number and Social Security number” (SunTrust Bank)). In very few cases, the notification referred to the recipient’s own breached information, such as the last four digit of credit card number (e.g., in Sprint Corp’s notification). Such an individually tailored message provides clear evidence that the recipient was personally affected, which might be a good strategy to alarm consumers and motivate them to take actions [52]. A potential downside is that it may pose identity theft risks if the notification falls into the wrong hands [96].

The cause of the breach was reported in 150 (93%) notifications. Causes varied from unauthorized access (38), phishing (33), malware (19), to inadvertent human error (16), and a few others. 11 notifications used “unauthorized access” to broadly describe how the breach occurred, without indicating what was accessed or by whom. Such vagueness may confuse consumers about what really occurred in the breach, potentially causing them to underestimate the chance that their data was compromised.

Ambiguity regarding uninformed exposure time. The dates of when the breach occurred and was discovered, as two elements not mandated by Maryland law, were mentioned by fewer notifications in our analysis. Only 103 notifications (64%) included a specific date or time range for when the breach occurred. While companies may not always be able to determine the breach date, without it consumers cannot know their “uninformed exposure time” [10], i.e., how long their data has been exposed before they become aware, which is an important metric to decide how urgently actions are needed. Only 105 notifications (65%) indicated when the company discovered the breach, which, together with the date when the notification was sent, would show the company’s diligence in informing consumers about security risks.

Hedge terms downplaying risks. A data breach notification should make it clear that the recipient’s information has been breached and describe associated risks [92]. Nonetheless, companies used various strategies to downplay a breach’s magnitude and consequences. Hedge terms such as “maybe” and “likely” were used in 112 (70%) notifications, obscuring whether the recipient was personally affected by the breach. These hedge terms appeared in various places. One place is in the general incident description, e.g., “I am writing to inform you of a data security incident that **may** have affected your payment card information.” (Temecula Motorsports). Another

is in the description of breached information types, e.g., “The information **potentially** involved in this incident **may** have included your name, credit or debit card number, and card expiration date.” (Bigfoot Gun Belts). As shown in previous work [66, 67, 98], the use of hedge terms is detrimental to consumers’ ability to accurately assess risks, and usually leads to confusion and misconception.

A positive example is provided by Parkway Corporation: “A file, including information from your IRS Tax Form W-2, was sent in response to the fraudulent email.” Here, it is explicitly stated that the incident happened and which data was compromised. Unfortunately, only 22 (14%) notifications used such clear statements. Furthermore, 23 (14%) notifications stated the company had no evidence of data being compromised, which downplay respective risks. MidCap Financial Services, for instance, said: “Although we do not have confirmation that any of these forms were accessed by the attacker, we are notifying you out of an abundance of caution.”

Obfuscating risks of misuse. Many companies obfuscated the risk of breached information being misused. Here the “no evidence” argument was even more prevalent, appearing in 64 (40%) notifications. For instance, Capital Digestive Care stated: “We do not believe that the limited information could be used adversely, and we have received no reports of the misuse of anyone’s data as a result of this incident.” Hedge terms were also inserted into the ‘no evidence’ claim in 31 (19%) notifications, for example, “At this time, we **have no evidence** that your personal information **has been or is likely to be misused**” (Pershing, LLC). In other cases, the no evidence claim was used for access and misuse together: “We **have no indication** that any emails obtained in this incident were actually viewed by any unauthorized third party or that any information from those emails has been misused” (Thermo Fisher Scientific). While it might be factually accurate that companies do not have evidence of data access and misuse, lack of evidence is not evidence of absence of harm; neither does it preclude future misuse of exposed data. Thus, without further warnings about the persistent risk of misuse, such statements downplay a data breach’s significance, potentially causing consumers to underestimate risks and discouraging them from taking immediate actions.

On the contrary, 9 notifications used an effective risk communication strategy — connecting the types of breach information with potential misuse scenarios. HomeBrewIt.com, for instance, highlighted that breached information in combination may be used for identity theft: “Your credit card number ending in XXXX may have been compromised. This number in conjunction with your billing address can potentially be used to make unauthorized purchases on your credit card.” Describing possible implications can also reinforce the need for specific protective actions, e.g., “We want you to be

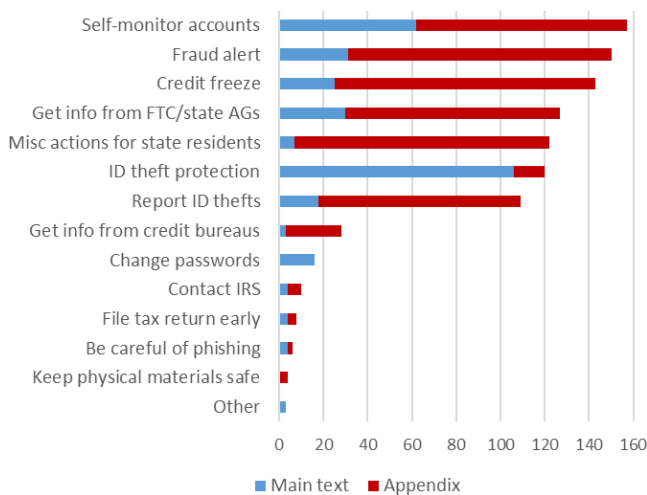


Figure 3: Frequency of recommended protective measures.

aware that because of the Incident, there is a possibility of: (1) identity theft, and (2) fraudulent filing of your tax information.” (Kinetic Systems). The second type of harm helps the recipient prioritize filing their tax return early.

Presentation of Recommended Actions

Consumers are urged to take protective actions when a data breach occurs. These actions vary in terms of effectiveness [94], including enrolling in the provided protection service, placing credit freezes, changing compromised passwords, carefully monitoring one’s financial accounts, and a few others. Several states and federal agencies provide templates for describing available measures, which companies may attach to their notifications. These templates usually include definitions of terms like fraud alert and credit freeze, contact information of major credit bureaus and the state AG, and enrollment instructions of free identity theft protection services if offered.

Choice overload with no priorities. Figure 3 shows the frequency distribution of commonly recommended actions. All notifications recommended at least one action, with a median of 8 actions (Mean=7.19, SD=2.24). 35 notifications (22%) recommended more than 9 protective measures, and the most comprehensive one included 16. The presence of so many options, compounded by lengthy explanations and poor readability, suggests the possibility of “choice overload” [77], meaning that the reader might delay the process to make a decision, pick a random option under pressure, or even avoid all options [15, 76].

We would expect key information to be presented in the main text, whereas appendices are generally used to provide supplemental materials. For each notification, we counted if

an appendix was used. For each action, we coded whether it first appeared in the main text or an appendix. The use of appendices was prevalent: 97 (60%) notifications used 1 appendix, 38 (24%) used 2 appendices, and 1 included 3 appendices. 25 (16%) notifications had main text only, ending with the sender’s contact information. Often, highly effective actions were hidden in long texts in appendices (see Figure 3), with no indication of their high priority. Credit freeze, for example, is considered an important protective measure to limit identity theft [87]. However, 118 (73%) notifications listed it as one of many options in an appendix. Even though they described options in detail, few companies compared different options directly or explicitly stated that a credit freeze provides stronger protection than a fraud alert, and thus should be prioritized.

Formatting focuses on sub-level text. Formatting, such as using lists and capitalizing important information, can effectively highlight key details and reduce the reader’s cognitive burden in processing text. Prior research suggests that certain formatting techniques, when used in a data breach notification, could enhance consumers’ perception of the affected company’s reputation [38].

We coded the presence of list and text formatting (e.g., bold, italicized, underlined, capitalized or colored text) when presenting actions. Overall, lists were scarcely used in top-level text (e.g., different actions), but became more prevalent in sub-level text (e.g., details and enrollment instructions of a specific action). The difference between top-level and sub-level formatting was sharper for bullet lists (8 vs. 83) compared to numbered lists (20 vs. 51). This indicates that while consumers are walked through details of each specific action, they may struggle to form a holistic view of what the major action options are due to the lack of list formatting on the top level. Heritage Land Bank provided a positive special example – using graphics – to encourage their consumers remain vigilant for fraud and identity theft in a vivid and salient way (see Figure 4).

Conversely, text formatting was common in both top-level (149, 93%) and sub-level text (90, 56%). However, text formatting was rarely used to highlight important details of offered identity protection services, specifically, the enrollment deadline (after which the service is no longer offered for free) and the duration of benefits (after which the protection is no longer effective). Among 124 notifications that offered such service, text formatting was used in only 46 (37%) to highlight the enrollment deadline, and even fewer (16, 13%) for the duration of benefits. In practice, the time frame to enroll in a provided service is usually short (less than a few months), and the service typically lasts for one to two years, during which consumers may easily lose track of the remaining time. When these crucial timings blend in with other

Protecting Your Financial Information



Remain vigilant for incidents of fraud or identity theft by reviewing your account statements and free credit reports for any unauthorized activity.

If you detect any unauthorized activity on financial accounts, immediately contact your financial institution, make them aware of this matter and take their advice on steps to protect your financial account.

Obtain a copy of your credit report, free-of-charge, once every 12 months from each of the three nationwide credit reporting companies by visiting AnnualCreditReport.com or call toll-free at

Figure 4: The use of graphics when recommending actions.

plain text without any visual highlights, consumers may miss out on free protection or have the illusion of being protected when they are not anymore.

6 DISCUSSION

Our analysis contributes novel insights on the readability and usability of recent data breach notifications. Building on Bisogni’s study [10], which showed that U.S. states’ breach notification laws vary substantially in terms of mandated elements, we found that mandatory elements were more frequently mentioned than elements not mandated by Maryland law: over 90% of our sample included types of compromised information, whereas only 65% reported the occurrence or discovery date. Consistent with Veltsos’s analysis [92], we observed frequent ‘no evidence of data misuse’ claims, and revealed the use of hedge terms as an additional strategy to downplay risks. Contrary to Jenkins et al.’s finding [38], formatting techniques were commonly used in our sample, but substantial difference emerged between top-level and sub-level text, and crucial information was not effectively highlighted. Next, we discuss potential limitations of our work and provide suggestions for improving the design of data breach notifications and respective legal and regulatory requirements.

Limitations

Our study has certain limitations. We only analyzed breach notifications from Maryland Attorney General’s public database, which may be different from those sent to consumers in other states. However, from cursory comparison with the databases of other state AGs, we are confident that our findings are not specific to notifications from the Maryland database. The fact that almost 70% of our sample used the structural headings mandated in California law suggests that

breach notifications are not necessarily tailored to specific states.

Furthermore, our content analysis does not provide direct evidence of how text and format of data breach notifications impact consumer behaviors. Nonetheless, given existing research on privacy policies [71, 91], where poor readability and ambiguity lead to users’ ignorance and misconception of privacy risks, we hypothesize that the issues we identified in data breach notifications would contribute to consumer inaction in a similar fashion. Our findings provide the basis for future user studies and experiments on the effects of the identified issues on consumers’ risk perception and intentions to take protective actions. Furthermore, additional research is required to better understand the overall role breach notifications play in consumers’ behavior after data breaches.

Finally, the HCI contributions of a content analysis of notification letters may not be immediately obvious. Both the cause (data breach) and consequence (the harm and protective actions users should take) of a data breach notification are rooted in technology, with the notification letter being a physical component in this overall user experience. Currently, most data breach notification laws require a mailed letter. The implications of our findings, however, are relevant for data breach notifications in general, regardless of the delivery medium. Our following recommendations can inform the design of more effective and actionable breach notifications both in letter form and in online contexts, as well as possible replacements for the paper-based process.

Design Implications

Based on our findings, we provide several design recommendations for enhancing data breach notifications’ readability by highlighting important information and communicating risks clearly.

Use clear and concise language. Data breach notifications should be readable and comprehensible by the general public, as data breaches can affect anyone. Our readability analysis, however, reveals that these notifications were written in lengthy paragraphs and complex language. While it is essential to ensure that mandated information is included, notifications can only be effective at informing consumers if the information is presented clearly and concisely. The California data breach law and the GDPR include a “plain language” requirement, but the majority of our sample are far from meeting this requirement. We advocate that government agencies and service providers who create data breach notification templates should devote more attention to craft information into short and explanatory sentences with little jargon. Furthermore, companies, when sending these notifications, should adjust the template to the corresponding

breach situation by removing unnecessary or inapplicable actions. For instance, 122 (76%) of our analyzed notifications appended a long list of contact information for different state AG offices. These notifications could be shortened by presenting only contact information of the recipient's state AG.

Support consumers in prioritizing multiple actions. Although prior research suggests providing explanation and guidance of protective measures is sufficient to empower consumers to take actions [92], our findings draw this into question. Recommended actions were often buried in long paragraphs, with little to no guidance on prioritization. Direct comparisons between actions in terms of their effectiveness or urgency were rarely made, leaving the reader overloaded with many choices, especially given that some of the recommendations are highly domain-specific and could easily be confused with other concepts (e.g., credit freeze could be mistaken as freezing one's credit cards [98]). This casts doubt on whether the provided recommendations are indeed perceivable and actionable. Thus, companies should clearly identify the recommended actions most necessary and applicable to the specific situation and list them in a certain order.

Already with minor adjustments consumers could be nudged towards preferable actions. Actions of high priority (e.g. due to high effectiveness, urgency, or easiness to initiate) should be listed before other options. For instance, credit freeze should be mentioned in the main text, and above other options such as fraud alert and credit lock, to indicate its effectiveness in preventing access to credit reports and thereby proactively reducing identity theft risks. Furthermore, notifications should explicitly recommend specific actions and explain the reasons why, instead of the current practice of presenting all explanations, definitions, and enrollment instructions together without a deliberate order. Using credit freeze again as an example, a message could be crafted in this way, "We recommend that you first place a credit freeze on your credit report, as it prevents credit, loans and services from being approved in your name without your consent. Next, you can also consider placing a fraud alert on your credit report. While less restrictive, a fraud alert tells creditors to be cautious before they open any new accounts or change your existing accounts."

Discourage hedge terms & the 'no evidence' defense. Prior research suggests that companies need to clearly articulate a breach's risks and potential threats in order to mitigate consumers' optimism bias and rational ignorance [70, 92]. We identified two strategies that could downplay the risks of a data breach: using hedge terms such as 'probably' and 'might' when describing the recipient's likelihood of being affected, and claiming that there is no evidence of breached data being misused. We do not insinuate that all companies that use

these strategies are deliberately misleading consumers. It is possible that a company might not be able to assess an individual customer's likelihood of being affected, or indeed have no evidence of data misuse. However, claims of no evidence of misuse could be misinterpreted by consumers as evidence of absence of risk, which is rarely the case. Breached data may be misused without the company's knowledge or could be misused in the future. As a result, consumers may underestimate actual risks.

Therefore, due to their potential of downplaying risks, we argue that companies should avoid "no evidence" claims and hedge terms, or at least combine them with clear warnings of potential misuse risks in the future. Consumers, when reading the claim, could be led to ignore existing risks and suffer severe consequences such as the substantial financial and emotional cost of identity theft. Recall bias based on similar events, anchoring, and other cognitive heuristics may further cause an underestimation of risks [54, 80]. Therefore, when companies face uncertainties regarding the scope or risks in the context of a data breach, overstating risks is more desirable than understating risks to balance such tendencies and trigger more immediate actions (e.g., initiating credit freezes, changing passwords, and enrolling in the provided free protection) – most of which are free and safe security practices that should be adopted regardless. Yet, overstating risks could also lead consumers to develop habituation to future breach notifications when they do not observe any actual harm in their lives. Such habituation effects could possibly be mitigated by carefully stating that measures should be taken out of precaution, and emphasizing that misuse risks can persist despite current absence of harm.

Highlight key information visually. Finally, information design literature suggests that formatting can make information visually accessible and enhance the overall user experience [40, 56]. Nevertheless, we found that list formats were common in sub-level but not top-level text regarding recommended protective actions. Although text formatting was prevalent, it was not effectively used to highlight the enrollment deadline and duration of benefits of provided compensation, increasing the chance for consumers to miss a free enrollment or forget to renew the service when it expires. We recommend the consistent use of list formats to lay out major actions; crucial information should further be highlighted with visual emphasis. While list formatting can be helpful, we do not want a notification overloaded with lists either. The choice of one particular type of list should be carefully made pertaining to the content: a bullet list indicates a parallel relationship, whilst a numbered list creates a step-by-step procedure with some prioritization and structure [9]. Each numbered or bullet point should be followed by short and

succinct sentences [37], so that the reader will not forget what it is meant to summarize after finish reading it.

Public Policy Implications

Echoing and complementing Bisogni’s study [10], our findings demonstrate the need to establish a U.S. federal data breach notification law, with stringent and unified requirements on formatting, content specificity, and possibility for more flexible delivery methods.

Clear readability expectations beyond ‘plain language’. 97% of our analyzed notifications were either fairly difficult or difficult to read. Given the suggestion that a middle-school FRES level (between 7-9) should be reached if the content is targeted to the general public, it is likely that many affected consumers cannot fully understand these notifications. Perhaps lessons can be learned from the insurance industry, where the FRES test is required as a readability assessment of insurance policies sold within some U.S. states [37]. Some data breach notification laws, notably the GDPR and the California law, have a ‘plain language’ requirement but do not define what ‘plain’ means, leaving room for interpretation. We suggest outlining specific guidance and examples on how the ‘plain language’ requirement is to be interpreted and how it could be achieved in order to make it more actionable by companies. Recommended practices should include using short sentences, commonly seen everyday words, and active voice. Standards for readability metrics can also be established, although the decision to pick one particular metric should be carefully made to ensure the formula is not biased and is applicable to the context.

Consistent standards for structure and format. Although 67% of our sample used headings, how they were formatted varied considerably, ranging from clearly separated headings to headings placed in a paragraph’s first line. Formatting techniques were inconsistent between top-level and sub-level text, and key details were not sufficiently highlighted by visual emphasis, indicating the need for clearer legal and regulatory guidance. The California law provides a promising example of the positive effects of incorporating specific formatting requirements. 62% of analyzed notifications used the headings required by the California law and 2 notifications used its recommended table layout. Similar requirements would be useful to ensure companies are appropriately prioritizing recommended actions and highlight important information, such as enrollment deadlines. Moreover, such requirements should be based on rigorous user testing to ensure that legal requirements or regulatory guidance actually improve usability. A positive example is the model notice for the annual privacy notices by financial institutions required under the Gramm-Leach Bliley Act (GLBA) [89]. The development of this model notice, which presents privacy

information and opt-out choices in a concise tabular format, was informed by usability testing [27].

Encourage notifications delivered in multiple channels. Currently, most state laws require companies to notify affected consumers in written letters or by telephone. Emails, website announcements, notices to statewide media, or other electronic methods are usually substitutes when individuals’ physical addresses are unavailable or the cost of delivering mails is too expensive. Our sample aligns with the legal requirements, with almost 95% notifications delivered by mail. However, in our digital age, attentions are increasingly shifting from papers to digital media. The slow speed of mailed letters also increases the ‘uninformed exposure time’ for consumers, which may help explain previous findings that many consumers learn about the breach even before receiving direct notifications from companies [1, 20]. A possible solution is that instead of choosing one channel, companies should be mandated to notify consumers through multiple channels, with preference to fast and reliable channels, to ensure consumers are informed of a data breach quickly and in a medium that facilitates taking action. The nature of electronic methods such as emails and in-app push notifications (small screen on the mobile device, fewer texts allowed) creates extra momentum for companies to make the message more succinct and use more visual elements to increase readability and aesthetics.

7 CONCLUSION

With the number of data breaches and their exposed records increasing over time, most affected consumers reportedly do not take effective protective actions and react indifferently to notifications sent by the breached company [61, 98]. These notifications, usually written as letters with long text, raise the question whether they are actually crafted in ways that provide essential information concisely, communicate consequences of the breach clearly, and help consumers make an informed decision of what actions to take.

We conducted a quantitative and qualitative content analysis on 161 sampled data breach notifications. Our findings reveal severe issues related to readability and usability: 97% of the analyzed notifications were classified as difficult or fairly difficult to read. Substantial disparities surfaced concerning the length and format of structural headings. Various techniques were used to downplay associated risks of the breach, failing to address consumers’ optimism bias and rational ignorance. Moreover, when presenting recommended actions, many companies overwhelmed recipients with information borrowed from existing templates, despite that these templates are overloaded with long text and too many

options, and provide little guidance for navigating and prioritizing different actions based on their effectiveness or urgency.

We provide design and public policy recommendations based on our findings. We advocate that writers and designers of data breach notifications should devote more attention to readability and visual attractiveness by using short sentences, common words, headings, lists and text formatting. Meanwhile, more efforts should be invested into achieving more effective risk communication by avoiding hedge terms and ‘no evidence’ claims, and providing actionable choices by nudging through placement and language. We further make recommend that data breach notification laws and respective regulatory guidance should (1) outline clear readability expectations and best practices for ‘plain language’ notifications; (2) make specific requirements of structure and format, unified across states; and (3) encourage the delivery of notifications in multiple channels. Essentially, our findings demonstrate the need for paying closer attention to the contents of data breach notifications and its effect on consumer behavior. Data breach notification laws miss their mark if resulting breach notifications are unreadable and do not effectively motivate consumers to take protective actions against the severe harms of data breaches.

ACKNOWLEDGMENTS

This research has been partially funded by the University of Michigan School of Information. The authors would like to thank Abraham H. Mhaidli and Justin Petelka for feedback on earlier versions of this paper.

REFERENCES

- [1] Lillian Ablon, Paul Heaton, Diana Catherine Lavery, and Sasha Romanosky. 2016. *Consumer Attitudes Toward Data Breach Notifications and Loss of Personal Information*. Technical Report. Rand Corporation.
- [2] Alessandro Acquisti, Idris Adjerid, Rebecca Balebako, Laura Brandimarte, Lorrie Faith Cranor, Saranga Komanduri, Pedro Giovanni Leon, Norman Sadeh, Florian Schaub, Manya Sleeper, et al. 2017. Nudges for Privacy and Security: Understanding and Assisting Users’ Choices Online. *ACM Computing Surveys (CSUR)* 50, 3 (2017), 44.
- [3] Alessandro Acquisti, Laura Brandimarte, and George Loewenstein. 2015. Privacy and human behavior in the age of information. *Science* 347, 6221 (2015), 509–514.
- [4] Idris Adjerid, Alessandro Acquisti, Laura Brandimarte, and George Loewenstein. 2013. Sleights of privacy: Framing, disclosures, and the limits of transparency. In *Proceedings of the Ninth Symposium on Usable Privacy and Security*. ACM, 9.
- [5] Devdatta Akhawe and Adrienne Porter Felt. 2013. Alice in Warningland: A Large-Scale Field Study of Browser Security Warning Effectiveness.. In *USENIX Security Symposium*, Vol. 13.
- [6] Hazim Almuhammedi, Adrienne Porter Felt, Robert W Reeder, and Sunny Consolvo. 2014. Your reputation precedes you: History, reputation, and the chrome malware warning. In *Symposium on Usable Privacy and Security (SOUPS)*, Vol. 4. 2.
- [7] Hazim Almuhammedi, Florian Schaub, Norman Sadeh, Idris Adjerid, Alessandro Acquisti, Joshua Gluck, Lorrie Faith Cranor, and Yuvraj Agarwal. 2015. Your location has been shared 5,398 times!: A field study on mobile app privacy nudging. In *Proceedings of the 33rd Annual ACM conference on Human Factors in Computing Systems*. ACM, 787–796.
- [8] American Bankers Association. 2018. Data Security & Customer Notification Requirements for Banks. <https://www.aba.com/Tools/Function/Technology/Pages/datasecuritynotification.aspx>. Last accessed on: 09.13.2018.
- [9] BBC. 2011. Using bullet points and numbers in lists. <http://www.bbc.co.uk/skills/wise/factsheet/en13styl-11-f-bulleted-and-numbered-points>. Last accessed on: 01.06.2019.
- [10] Fabio Bisogni. 2016. Proving Limits of State Data Breach Notification Laws: Is a Federal Law the Most Adequate Solution? *Journal of Information Policy* 6, 1 (2016), 154–205.
- [11] Cristian Bravo-Lillo, Lorrie Faith Cranor, Julie Downs, and Saranga Komanduri. 2011. Bridging the gap in computer security warnings: A mental model approach. *IEEE Security & Privacy* 9, 2 (2011), 18–26.
- [12] Cristian Bravo-Lillo, Saranga Komanduri, Lorrie Faith Cranor, Robert W Reeder, Manya Sleeper, Julie Downs, and Stuart Schechter. 2013. Your attention please: designing security-decision UIs to make genuine risks harder to ignore. In *Proceedings of the Ninth Symposium on Usable Privacy and Security*. ACM, 6.
- [13] Ronald P Carver. 1983. Is reading rate constant or flexible? *Reading Research Quarterly* (1983), 190–215.
- [14] F. H. Cate. 2010. The Limits of Notice and Choice. *IEEE Security Privacy* 8, 2 (2010), 59–62.
- [15] Alexander Chernev, Ulf Böckenholt, and Joseph Goodman. 2015. Choice overload: A conceptual review and meta-analysis. *Journal of Consumer Psychology* 25, 2 (2015), 333–358.
- [16] Lauren Lyons Cole. 2017. After the Equifax breach, consumers were advised to freeze their credit - but almost no one did it. <http://www.businessinsider.com/equifax-credit-freeze-2017-9>. Last accessed on: 01.22.2018.
- [17] Council of European Union. 2017. General Data Protection Regulation (GDPR). <http://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:32016R0679>. Last accessed on: 04.28.2018.
- [18] Lorrie Faith Cranor. 2012. Necessary but not sufficient: Standardized mechanisms for privacy notice and choice. *J. on Telecomm. & High Tech. L.* 10 (2012), 273.
- [19] Lorrie Faith Cranor, Pedro Giovanni Leon, and Blase Ur. 2016. A large-scale evaluation of US financial institutions’ standardized privacy notices. *ACM Transactions on the Web (TWEB)* 10, 3 (2016), 17.
- [20] Sauvik Das, Joanne Lo, Laura Dabbish, and Jason I Hong. 2018. Breaking! A Typology of Security and Privacy News and How It’s Shared. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM, 1.
- [21] Rachna Dhamija, J Doug Tygar, and Marti Hearst. 2006. Why phishing works. In *Proceedings of the SIGCHI conference on Human Factors in computing systems*. ACM, 581–590.
- [22] Anthony Downs. 1957. An economic theory of political action in a democracy. *Journal of Political Economy* 65, 2 (1957), 135–150.
- [23] Adrienne Porter Felt, Alex Ainslie, Robert W Reeder, Sunny Consolvo, Somas Thyagaraja, Alan Bettes, Helen Harris, and Jeff Grimes. 2015. Improving SSL warnings: Comprehension and adherence. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. ACM, 2893–2902.
- [24] Rudolf Franz Flesch et al. 1949. *Art of readable writing*. Harper.
- [25] Alain Forget, Sonia Chiasson, Paul C van Oorschot, and Robert Biddle. 2008. Improving text passwords through persuasion. In *Proceedings of the 4th Symposium on Usable Privacy and Security*. ACM, 1–12.
- [26] Brian Fung. 2018. Equifax’s massive 2017 data breach keeps getting worse. <https://www.washingtonpost.com/news/the-switch/wp/2018/03/01/equifax-keeps-finding-millions-more-people-who-were->

- affected-by-its-massive-data-breach/?noredirect=on&utm_term=.52f7af5c120a. Last accessed on: 09.09.2018.
- [27] Loretta Garrison, Manoj Hastak, Jeanne M Hogarth, Susan Kleimann, and Alan S Levy. 2012. Designing Evidence-based Disclosures: A Case Study of Financial Privacy Notices. *Journal of Consumer Affairs* 46, 2 (2012), 204–234.
- [28] Gemalto. 2017. *Data Breaches and Customer Loyalty 2017*. Technical Report. Gemalto.
- [29] General Assembly of Maryland. 2018. Md. Code Ann. Comm. Law 14-3504: Maryland’s Personal Information Protection Act. <http://mgaleg.maryland.gov/webmga/firmStatutesText.aspx?article=gcl§ion=14-3501&ext=html&session=2017RS&tab=subject5>. Last accessed on: 06.05.2018.
- [30] Joshua Gluck, Florian Schaub, Amy Friedman, Hana Habib, Norman Sadeh, Lorrie Faith Cranor, and Yuvraj Agarwal. 2016. How short is too short? Implications of length and framing on the effectiveness of privacy notices. In *12th Symposium on Usable Privacy and Security (SOUPS)*. 321–340.
- [31] Maximilian Golla, Miranda Wei, Juliette Hainline, Lydia Filipe, Markus Dürmuth, Elissa Redmiles, and Blase Ur. 2018. What was that site doing with my Facebook password?: Designing Password-Reuse Notifications. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*. ACM, 1549–1566.
- [32] Kelli Grant. 2017. Identity theft, fraud cost consumers more than \$16 billion. <https://www.cnbc.com/2017/02/01/consumers-lost-more-than-16b-to-fraud-and-identity-theft-last-year.html> Last accessed on: 06.19.2018.
- [33] Claire Greene and Joanna Stavins. 2017. Did the Target data breach change consumer assessments of payment card security? *Journal of Payments Strategy & Systems* 11, 2 (2017), 121–133.
- [34] Robert Gunning. 1969. The fog index after twenty years. *Journal of Business Communication* 6, 2 (1969), 3–13.
- [35] Erika Harrell and Lynn Langton. 2015. *Victims of identity theft, 2014*. Technical Report.
- [36] HIPAA Journal. 2017. What are the HIPAA Breach Notification Requirements? <https://www.hipaajournal.com/hipaa-breach-notification-requirements/>. Last accessed on: 09.13.2018.
- [37] Mark Hochhauser. 2001. Lost in the Fine Print: Readability of Financial Privacy Notices. <https://www.privacyrights.org/blog/lost-fine-print-readability-financial-privacy-notices-hochhauser>. Last accessed on: 09.13.2018.
- [38] Alexander Jenkins, Murugan Anandarajan, and Rob D’Ovidio. 2014. ‘All that Glitters is not Gold’: The Role of Impression Management in Data Breach Notification. *Western Journal of Communication* 78, 3 (2014), 337–357.
- [39] Carlos Jensen and Colin Potts. 2004. Privacy policies as decision-making tools: an evaluation of online privacy notices. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 471–478.
- [40] Elizabeth Keyes. 1993. Typography, color, and information structure. *Technical communication* (1993), 638–654.
- [41] Bart P Knijnenburg and Alfred Kobsa. 2013. Making decisions about privacy: information disclosure in context-aware recommender systems. *ACM Transactions on Interactive Intelligent Systems (TiiS)* 3, 3 (2013), 20.
- [42] Jeffery Kosseff. 2016. My company has had a breach: Whom do I have to notify? <https://iapp.org/news/a/my-company-has-had-a-breach-who-do-i-have-to-notify/>. Last accessed on: 09.18.2018.
- [43] Thomas Kude, Hartmut Hoehle, and Tracy Ann Sykes. 2017. Big data breaches and customer compensation strategies: Personality traits and social influence as antecedents of perceived compensation. *International Journal of Operations & Production Management* 37, 1 (2017), 56–74.
- [44] Jonathan Lazar, Jinjuan Heidi Feng, and Harry Hochheiser. 2017. *Research methods in human-computer interaction*. Morgan Kaufmann.
- [45] Joseph Lazzarotti, Jason Gavejian, and Maya Atrakchi. 2018. Security Breach Notification Laws. <http://www.ncsl.org/research/telecommunications-and-information-technology/security-breach-notification-laws.aspx>. Last accessed on: 06.05.2018.
- [46] Johnny Lieu. 2017. Terms and Conditions are too long, just ask a guy who read Amazon’s for 9 hours. <https://mashable.com/2017/03/15/reading-amazons-terms-conditions/#IQDa1u7BsOq0>. Last accessed on: 09.13.2018.
- [47] Ewa Luger, Stuart Moran, and Tom Rodden. 2013. Consent for all: revealing the hidden complexity of terms and conditions. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 2687–2696.
- [48] Bernard Mar. 2018. GDPR: The Biggest Data Breaches And The Shocking Fines (That Would Have Been). <https://www.forbes.com/sites/bernardmarr/2018/06/11/gdpr-the-biggest-data-breaches-and-the-shocking-fines-that-would-have-been/#199b5b4b6c10>. Last accessed on: 09.18.2018.
- [49] Maryland Coordination and Analysis Center. 2018. Maryland Data Breach Notification Law Updated. <http://www.mcac.maryland.gov/newsroom/Critical%20Infrastructure%20News/maryland-data-breach-notification-law-updated>. Last accessed on: 06.05.2018.
- [50] Maryland’s State Attorney General. 2018. Guidelines for businesses to comply with the Maryland Personal Information Protection Act. <http://www.marylandattorneygeneral.gov/Pages/IdentityTheft/businessGL.aspx>. Last accessed on: 06.05.2018.
- [51] Aleecia M McDonald and Lorrie Faith Cranor. 2008. The cost of reading privacy policies. *ISJLP* 4 (2008), 543.
- [52] Vyacheslav Mikhed and Michael Vogan. 2015. Out of sight, out of mind: consumer reaction to news on data breaches and identity theft. (2015). Working Paper.
- [53] Drew Mitnick. 2018. No more waiting: it’s time for a federal data breach law in the U.S. <https://www.accessnow.org/no-more-waiting-its-time-for-a-federal-data-breach-law-in-the-u-s/>. Last accessed on: 09.18.2018.
- [54] M Granger Morgan, Baruch Fischhoff, Ann Bostrom, and Cynthia J Atman. 2002. *Risk communication: A mental models approach*. Cambridge University Press.
- [55] National Conference of State Legislators. 2018. 2018 Security Breach Legislation. <http://www.ncsl.org/research/telecommunications-and-information-technology/2018-security-breach-legislation.aspx>. Last accessed on: 09.13.2018.
- [56] Jakob Nielsen. 1997. How Users Read on the Web.
- [57] Patricia A Norberg, Daniel R Horne, and David A Horne. 2007. The privacy paradox: Personal information disclosure intentions versus behaviors. *Journal of Consumer Affairs* 41, 1 (2007), 100–126.
- [58] Don Norman. 2013. *The design of everyday things: Revised and expanded edition*. Constellation.
- [59] Eyal Peer and Alessandro Acquisti. 2016. The impact of reversibility on the decision to disclose personal information. *Journal of Consumer Marketing* 33, 6 (2016), 428–436.
- [60] Justin Petelka, Yixin Zou, and Florian Schaub. 2019. Put Your Warning Where Your Link Is: Improving and Evaluating Email Phishing Warnings. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. ACM.
- [61] Ponemon Institute. 2014. *The Aftermath of a Data Breach: Consumer Sentiment*. Technical Report. Ponemon Institute LLC.
- [62] Privacy Rights Clearinghouse. 2016. What to Do When You Receive A Data Breach Notice. <https://www.privacyrights.org/consumer-guides/what-do-when-you-receive-data-breach-notice>. Last accessed

- on: 09.13.2018.
- [63] Privacy Rights Clearinghouse. 2018. Data Breaches. <https://www.privacyrights.org/data-breaches>. Last accessed on: 09.13.2018.
- [64] Robert W Proctor, M Athar Ali, and Kim-Phuong L Vu. 2008. Examining usability of web privacy policies. *Intl. Journal of Human-Computer Interaction* 24, 3 (2008), 307–328.
- [65] Ashwini Rao, Florian Schaub, Norman Sadeh, Alessandro Acquisti, and Ruogu Kang. 2016. Expecting the unexpected: Understanding mismatched privacy expectations online. In *Symposium on Usable Privacy and Security (SOUPS)*, Vol. 4. 2.
- [66] Joel R Reidenberg, Jaspreet Bhatia, Travis D Breaux, and Thomas B Norton. 2016. Ambiguity in privacy policies and the impact of regulation. *The Journal of Legal Studies* 45, S2 (2016), S163–S190.
- [67] Joel R Reidenberg, Travis Breaux, Lorrie Faith Cranor, Brian French, Amanda Grannis, James T Graves, Fei Liu, Aleecia McDonald, Thomas B Norton, and Rohan Ramanath. 2015. Disagreeable privacy policies: Mismatches between meaning and users’ understanding. *Berkeley Tech. Lj* 30 (2015), 39.
- [68] Joel R Reidenberg, N Cameron Russell, Alexander J Callen, Sophia Qasir, and Thomas B Norton. 2015. Privacy harms and the effectiveness of the notice and choice framework. *ISJLP* 11 (2015), 485.
- [69] Alex Reynolds. 2017. GDPR matchup: US state data breach laws. <https://iapp.org/news/a/gdpr-match-up-u-s-state-data-breach-laws/>. Last accessed on: 09.18.2018.
- [70] Sasha Romanosky, Rahul Telang, and Alessandro Acquisti. 2011. Do data breach disclosure laws reduce identity theft? *Journal of Policy Analysis and Management* 30, 2 (2011), 256–286.
- [71] Manuel Rudolph, Denis Feth, and Svenja Polst. 2018. Why Users Ignore Privacy Policies—A Survey and Intention Model for Explaining User Privacy Behavior. In *International Conference on Human-Computer Interaction*. Springer, 587–598.
- [72] Sonam Samat and Alessandro Acquisti. 2017. Format vs. Content: The Impact of Risk and Presentation on Disclosure Decisions. In *Thirteenth Symposium on Usable Privacy and Security (SOUPS 2017)*. 377–384.
- [73] Sonam Samat, Alessandro Acquisti, and Linda Babcock. 2017. Raise the Curtains: The Effect of Awareness About Targeting on Consumer Attitudes and Purchase Intentions. In *Thirteenth Symposium on Usable Privacy and Security (SOUPS) 2017*. USENIX Association, 299–319.
- [74] F. Schaub, R. Balebako, and L. F. Cranor. 2018. Designing Effective Privacy Notices and Controls. *IEEE Internet Computing* (2018), 1–1.
- [75] Florian Schaub, Rebecca Balebako, Adam L Durity, and Lorrie Faith Cranor. 2015. A design space for effective privacy notices. In *Eleventh Symposium On Usable Privacy and Security (SOUPS 2015)*. 1–17.
- [76] Benjamin Scheibehenne, Rainer Greifeneder, and Peter M Todd. 2010. Can there ever be too many options? A meta-analytic review of choice overload. *Journal of Consumer Research* 37, 3 (2010), 409–425.
- [77] Barry Schwartz. 2004. The paradox of choice: Why more is less. Ecco New York.
- [78] Tali Sharot. 2011. The optimism bias. *Current biology* 21, 23 (2011), R941–R945.
- [79] Robert H Sloan and Richard Warner. 2014. Beyond notice and choice: Privacy, norms, and consent. *J. High Tech. L.* 14 (2014), 370.
- [80] Paul Slovic, Baruch Fischhoff, and Sarah Lichtenstein. 1979. Rating the risks. *Environment: Science and Policy for Sustainable Development* 21, 3 (1979), 14–39.
- [81] Peter Swire and Kenesa Ahmad. 2012. *Foundations of Information Privacy and Data Protection*. International Association of Privacy Professionals.
- [82] Richard H Thaler and Cass R Sunstein. 2008. *Nudge: Improving decisions about health, wealth, and happiness*. HeinOnline.
- [83] The California State Government. 2003. California Civ. Code s. 1798.82(a). https://leginfo.ca.gov/faces/codes_displaySection.xhtml?lawCode=CIV§ionNum=1798.82. Last accessed on: 06.05.2018.
- [84] The Federal Trade Commission. 2018. Gramm-Leach-Bliley Act. <https://www.ftc.gov/tips-advice/business-center/privacy-and-security/gramm-leach-bliley-act>. Last accessed on: 09.13.2018.
- [85] The Privacy Rights Clearinghouse. 2018. Data Breaches. <https://www.privacyrights.org/data-breaches>. Last accessed on: 12.19.2018.
- [86] The U.S. Government Printing Office. 1996. Health Insurance Portability and Accountability Act of 1996, Public Law 104-191. <https://www.gpo.gov/fdsys/pkg/PLAW-104publ191/html/PLAW-104publ191.htm>. Last accessed on: 09.18.2018.
- [87] Susan Tompor. 2018. Credit freeze: A misunderstood freebie that you actually want. <https://www.freep.com/story/money/personal-finance/susan-tompor/2018/09/06/equifax-freeze-credit-breach/1156255002/>. Last accessed on: 09.13.2018.
- [88] Janice Y Tsai, Serge Egelman, Lorrie Cranor, and Alessandro Acquisti. 2011. The effect of online privacy information on purchasing behavior: An experimental study. *Information Systems Research* 22, 2 (2011), 254–268.
- [89] United States Congress. 1999. S.900 - Gramm-Leach-Bliley Act. <https://www.congress.gov/bill/106th-congress/senate-bill/00900>. Last accessed on: 09.13.2018.
- [90] Blase Ur, Patrick Gage Kelley, Saranga Komanduri, Joel Lee, Michael Maass, Michelle L Mazurek, Timothy Passaro, Richard Shay, Timothy Vidas, Lujo Bauer, et al. 2012. How does your password measure up? The effect of strength meters on password creation.. In *USENIX Security Symposium*. 65–80.
- [91] Matthew W Vail, Julia B Earp, and Annie I Antón. 2008. An empirical study of consumer perceptions and comprehension of web site privacy policies. *IEEE Transactions on Engineering Management* 55, 3 (2008), 442–454.
- [92] Jennifer R Veltsos. 2012. An analysis of data breach notifications as negative news. *Business Communication Quarterly* 75, 2 (2012), 192–207.
- [93] Melanie Volkamer, Karen Renaud, Benjamin Reinheimer, and Alexandra Kunz. 2017. User experiences of TORPEDO: tooltip-powered phishing email detection. *Computers & Security* 71 (2017), 100–113.
- [94] Paul Wagenseil. 2017. What to Do After a Data Breach. <https://www.tomsguide.com/us/data-breach-to-dos,news-18007.html>. Last accessed on: 09.13.2018.
- [95] Yang Wang, Pedro Giovanni Leon, Alessandro Acquisti, Lorrie Faith Cranor, Alain Forget, and Norman Sadeh. 2014. A field trial of privacy nudges for facebook. In *Proceedings of the SIGCHI conference on human factors in computing systems*. ACM, 2367–2376.
- [96] Kelce Wilson. 2018. Data breach notifications may facilitate identity theft. <https://iapp.org/news/a/data-breach-notifications-may-facilitate-identity-theft/>. Last accessed on: 09.13.2018.
- [97] Shomir Wilson, Justin Cranshaw, Norman Sadeh, Alessandro Acquisti, Lorrie Faith Cranor, Jay Springfield, Sae Young Jeong, and Arun Balasubramanian. 2013. Privacy manipulation and acclimation in a location sharing application. In *Proceedings of the 2013 ACM International Joint Conference on Pervasive and Ubiquitous Computing*. ACM, 549–558.
- [98] Yixin Zou, Abraham H. Mhaidli, Austin McCall, and Florian Schaub. 2018. “I’ve Got Nothing to Lose”: Consumers’ Risk Perceptions and Protective Actions after the Equifax Data Breach. In *Proceedings of the Fourteenth Symposium on Usable Privacy and Security (SOUPS)*.