

# Understanding and Improving Consumers' Adoption of Online Privacy-Protective Behaviors

by

Yixin Zou

A dissertation submitted in partial fulfillment  
of the requirements for the degree of  
Doctor of Philosophy  
(Information)  
in The University of Michigan  
2022

Doctoral Committee:

Associate Professor Florian Schaub, Chair  
Professor Alessandro Acquisti  
Professor Denise L. Anthony  
Associate Professor Adam Aviv  
Associate Professor Tawanna Ruth Dillahunt

Yixin Zou  
yixinz@umich.edu  
ORCID iD: 0000-0002-9088-705X

©Yixin Zou 2022

## ACKNOWLEDGEMENTS

The work presented in this dissertation would not be possible without the guidance and support from many people. I would like to start by thanking my advisor Florian Schaub, who has shown me what it means to be a great advisor. I knew very little about usable privacy and security when I applied to Ph.D. programs, and I am immensely grateful to Florian for seeing my potential and introducing me to the field. I have learned so much from him about not just how to conduct high-quality and high-impact research, but also how to lead by example and generously support students both within and outside work.

I would like to thank my dissertation committee members—Alessandro Acquisti, Denise Anthony, Adam Aviv, and Tawanna Dillahunt—for their continued guidance and constructive feedback. My big thanks go to Alessandro, who has always been so kind and patient and helped me become better at experimental methods day by day. Similarly, I have learned so much from Denise and the valuable perspectives she has brought from sociology and public health literature. I am lucky to have collaborated with Adam on research for multiple chapters of my dissertation; I cannot thank him enough for his insights and humor that made it such a joyful experience. In addition, I am lucky to have Tawanna as a role model, who has taught me so much about incorporating care into conducting research and building meaningful relationships with community partners. I feel extremely privileged to have had such a stellar committee accompanying me on this journey.

I have been incredibly fortunate to have many amazing collaborators over the

years, including but not limited to: Tanisha Afnan, Janet Chen, Lorrie Cranor, Shawn Danino, Nicola Dell, Hana Habib, Khue Le, Byron Lowens, Peter Mayer, Austin McCall, Allison McDonald, Julia Narakornpichit, Justin Petelka, Joel Reidenberg, Tom Ristenpart, Johann Roturier, Kevin Roundy, Norman Sadeh, Saurabh Shintre, Kaiwen Sun, Acar Tamersoy, Emily Tseng, and Yaxing Yao. Some of them are great mentors who have guided me through big decisions. Some of them are wonderful mentees who have generously tolerated my flaws and helped me improve my mentoring approaches. Some of them have become my best friends, and I look forward to continue collaborating with them and sharing our next chapters of life with each other.

I have grown so much being part of a vibrant research community at the University of Michigan School of Information. My sincere thanks go to the faculty mentors who have taught me to become a better researcher and teacher: Mark Ackerman, Robin Brewer, Nicole Ellison, Margaret Hedstrom, Tanya Rosenblat, Chris Teplovs, Kentaro Toyama, and Tiffany Veinot. I would like to acknowledge the UMSI staff who have provided crucial logistical support for my research. I am particularly grateful to Amy Eaton and Erin Krupka for their tremendous help as I navigated the doctoral program.

Within and beyond UMSI, I have been lucky to form a support network with many friends and mentors that brought me great advice, encouragement, and joy, including but not limited to: Ruba Abu-Salma, Yasemin Acar, Aparna Ananthasubramaniam, Sam Ankenbauer, Yan Chen, Camille Cobb, Sylvia Darling, Jennifer Davidson, An Vinh Doan, Yuanyuan Feng, Maximilian Golla, Cami Goray, Conor Gilsean, Gabriel Grill, Tamy Guberek, Carolyn Guthoff, Wenbo Guo, Rakibul Hasan, Licia He, Gina Herakovic, Joey Hsiao, Jane Im, Aarti Israni, Haojian Jin, Ju Yeon Jung, Vaishnav Kameswaran, Doowon Kim, Sangmi Kim, Elsie Lee-Robbins, Jingjie Li, Linfeng Li, Warren Li, Chen Liang, Alexander Liu, Alex Lu, Danaja Maldeniya, Robert Markum,

Liz Marquis, Collins Munyendo, Moses Namara, Sony Prosper, Chris Quarles, Ellen Quarles, Ashwin Rajadesingan, Ramakrishnan Sundara Raman, Reethika Ramesh, Elissa Redmiles, Olivia Richards, Kat Roemmich, Woo Suk Seo, Carlo Sugatan, April Wang, Ihudiya Finda Williams, Ritchie Wilson, Zhuofeng Wu, Shiyang Yan, Teng Ye, Yulin Yu, Ben Zhang, Lei Zhang, and Xinyan Zhao.

I am grateful to have the opportunity to work with some great community partners who have spared time out of their busy schedules to enable my research. This includes but is not limited to: Ann Arbor Senior Center, Chelsea Senior Center, Saline Area Senior Center, SafeHouse Center, and the Human Trafficking Clinic at the University of Michigan Law School. Along this line, I would like to thank all my study participants for contributing valuable insights that constitute the foundation of my dissertation.

Many people have asked me how I pivoted from studying advertising during my undergraduate years at the University of Illinois at Urbana-Champaign, and I know this pivot would not be possible without the great mentors I had there, including but not limited to: Aseel Addawood, Masooda Bashir, Ishva Minefee, Michelle Nelson, Maryalice Wu, and Mike Yao. I thank every one of them for helping me form a curiosity in research and build research experiences. I would not think about pursuing a Ph.D. without their guidance and encouragement.

My family and friends back in China have made me who I am. I want to give special acknowledgment to my parents—I am usually shy to express my love to them directly, but deep inside I know they always have my back. I am forever grateful to them for supporting my education and rooting for me as I pursue my dreams overseas. Writing about family is also a bittersweet moment as I think of my maternal grandmother, who passed away one year before I finished my dissertation. I was not able to attend her funeral due to the COVID-19 pandemic, but I hope she would be happy to learn about my achievements and I could make her proud.

Finally, I would like to thank my partner, my favorite collaborator, and my best friend Abraham Mhaidli, who has spoiled me with unconditional love and support that has powered me through the most difficult times over the past four years. I could not have done my dissertation without him.

# TABLE OF CONTENTS

<b>ACKNOWLEDGEMENTS</b> . . . . .	ii
<b>LIST OF FIGURES</b> . . . . .	xi
<b>LIST OF TABLES</b> . . . . .	xiii
<b>LIST OF APPENDICES</b> . . . . .	xvi
<b>ABSTRACT</b> . . . . .	xvii
<b>CHAPTER</b>	
<b>I. Introduction</b> . . . . .	1
1.1 Understanding Adoption . . . . .	4
1.1.1 Consumer Reactions to the 2017 Equifax Data Breach: An Interview Study . . . . .	5
1.1.2 Consumer Reactions to Data Breaches that Affected Them: A Survey Study . . . . .	6
1.1.3 An Empirical Analysis of Data Breach Notifications	7
1.1.4 Adoption and Abandonment of Security, Privacy, and Identity Theft Protection Practices . . . . .	7
1.2 Improving Adoption . . . . .	8
1.2.1 Icons and Link Texts for Conveying Privacy Choices	9
1.2.2 Computer Security Customer Support for Survivors of Intimate Partner Violence . . . . .	10
1.2.3 Nudges for Encouraging Password Changes After Data Breaches . . . . .	11
<b>II. Related Work</b> . . . . .	12
2.1 A Taxonomy of Online Privacy-Protective Behaviors . . . . .	12
2.1.1 Coping by Avoidance . . . . .	13
2.1.2 Coping by Approach . . . . .	15

2.2	Factors Behind Adoption and Non-Adoption . . . . .	19
2.3	Approaches to Improving Privacy-Protective Behaviors . . . . .	24
2.3.1	Privacy Notices and Choices . . . . .	25
2.3.2	Privacy Nudges . . . . .	28
2.3.3	Privacy Regulations . . . . .	31
<b>III. Consumer Reactions to the 2017 Equifax Data Breach: An Interview Study . . . . .</b>		<b>33</b>
3.1	Background . . . . .	34
3.2	Method . . . . .	36
3.3	Findings . . . . .	40
3.3.1	Mental Models of Credit Bureaus . . . . .	41
3.3.2	Risk Perception of the Equifax Data Breach . . . . .	45
3.3.3	Negligence of Protective Actions . . . . .	48
3.4	Discussion . . . . .	54
3.5	Recommendations . . . . .	56
<b>IV. Consumer Reactions to Data Breaches: A Survey Study . . . . .</b>		<b>60</b>
4.1	Method . . . . .	62
4.2	Findings . . . . .	67
4.2.1	Sample . . . . .	68
4.2.2	RQ1: Likelihood of Breaches . . . . .	71
4.2.3	RQ2: Perceived Causes and Impacts of Breaches . . . . .	73
4.2.4	RQ3: Awareness of Breaches . . . . .	76
4.2.5	RQ4: Emotional Responses Towards Breaches . . . . .	77
4.2.6	RQ5: Behavioral Reactions to Breaches . . . . .	81
4.3	Discussion . . . . .	83
<b>V. An Empirical Analysis of Data Breach Notifications . . . . .</b>		<b>89</b>
5.1	Background . . . . .	90
5.2	Method . . . . .	92
5.3	Results . . . . .	94
5.3.1	Readability and Structure . . . . .	95
5.3.2	Risk Communication . . . . .	97
5.3.3	Presentation of Recommended Actions . . . . .	101
5.4	Discussion . . . . .	104
<b>VI. Examining the Adoption and Abandonment of Security, Privacy, and Identity Theft Protection Practices . . . . .</b>		<b>109</b>
6.1	Method . . . . .	111
6.1.1	Taxonomy of Expert-Recommended Practices . . . . .	111



6.1.2	Survey Protocol and Recruitment . . . . .	114
6.1.3	Data Analysis . . . . .	116
6.1.4	Limitations . . . . .	118
6.2	Results . . . . .	119
6.2.1	RQ1: Commonly Adopted and Abandoned Practices	119
6.2.2	RQ2: Factors Affecting Levels of Adoption . . . . .	122
6.2.3	RQ3: Reasons for Partial Adoption and Abandonment	126
6.3	Discussion . . . . .	132

**VII. Icons and Link Texts for Conveying Privacy Choices . . . . . 136**

7.1	Background . . . . .	138
7.2	Study Overview . . . . .	140
7.3	Icon Pre-Study . . . . .	142
7.3.1	Icon Development . . . . .	142
7.3.2	Preliminary Icon Testing . . . . .	145
7.3.3	Refined Icon Testing . . . . .	150
7.4	Link Text Pre-Study . . . . .	153
7.4.1	Link Text Development . . . . .	154
7.4.2	Preliminary Link Text Testing . . . . .	155
7.4.3	Refined Link Text Testing . . . . .	156
7.5	Icon-Text Combinations Evaluation . . . . .	160
7.5.1	Method . . . . .	161
7.5.2	Findings . . . . .	163
7.6	OAG Icon Evaluation . . . . .	167
7.6.1	Method . . . . .	168
7.6.2	Findings . . . . .	169
7.7	Discussion . . . . .	172

**VIII. Computer Security Customer Support for Survivors of Intimate Partner Violence . . . . . 178**

8.1	Background . . . . .	181
8.2	Preliminary Analysis of Support Cases . . . . .	184
8.2.1	Method . . . . .	184
8.2.2	Results . . . . .	186
8.3	Focus Groups with IPV Professionals . . . . .	187
8.3.1	Method . . . . .	188
8.3.2	Suggestions for Interacting with Survivors . . . . .	190
8.3.3	Responsibilities of Customer Support . . . . .	193
8.3.4	Suggestions for External Referrals . . . . .	196
8.4	Focus Groups with Support Practitioners . . . . .	199
8.4.1	Method . . . . .	200
8.4.2	Well-Received Suggestions . . . . .	201
8.4.3	Implementation Challenges . . . . .	205

8.5	Discussion . . . . .	209
<b>IX.</b>	<b>Nudges to Encourage Password Changes After Data Breaches Using the Protection Motivation Theory . . . . .</b>	<b>216</b>
9.1	Related Work . . . . .	218
9.1.1	Protection Motivation Theory . . . . .	219
9.1.2	Password Policies, Behaviors, and Nudges . . . . .	223
9.2	Method . . . . .	226
9.2.1	Study Design . . . . .	226
9.2.2	Protocol . . . . .	228
9.2.3	Data Analysis . . . . .	235
9.3	Findings . . . . .	239
9.3.1	Sample . . . . .	239
9.3.2	Results of Confirmatory Analyses . . . . .	242
9.3.3	Results of Exploratory Analyses . . . . .	243
9.3.4	Qualitative Insights . . . . .	255
9.4	Discussion . . . . .	264
9.4.1	Limitations . . . . .	264
9.4.2	Theoretical Contributions . . . . .	266
9.4.3	Practical Implications . . . . .	270
<b>X.</b>	<b>Conclusion and Future Work . . . . .</b>	<b>277</b>
10.1	Summary of Key Contributions . . . . .	277
10.2	Reflection on Lessons Learned and Future Work . . . . .	280
<b>APPENDICES</b>	<b>. . . . .</b>	<b>289</b>
<b>A.</b>	<b>Supplemental Materials: “Consumer Reactions to the 2017 Equifax Data Breach” . . . . .</b>	<b>290</b>
A.1	Interview Protocol . . . . .	290
A.2	Screening Survey for Recruitment . . . . .	294
A.3	Qualitative Analysis Codebook . . . . .	295
<b>B.</b>	<b>Supplemental Materials: “Consumer Reactions to Data Breaches that Affected Them” . . . . .</b>	<b>303</b>
B.1	Survey Materials . . . . .	303
B.2	Qualitative Analysis Codebook . . . . .	310
<b>C.</b>	<b>Supplemental Materials: “An Empirical Analysis of Data Breach Notifications” . . . . .</b>	<b>314</b>

C.1 Analyzed Breaches Notifications . . . . .	314
C.2 Qualitative Analysis Codebook . . . . .	320
<b>D. Supplemental Materials: “Adoption and Abandonment of Protection Practices” . . . . .</b>	<b>322</b>
D.1 Survey Instruments . . . . .	322
D.2 Qualitative Analysis Codebook . . . . .	331
<b>E. Supplemental Materials: “Icons for Conveying Privacy Controls” . . . . .</b>	<b>338</b>
E.1 Survey Instruments . . . . .	338
E.2 Detailed Statistical Results . . . . .	347
<b>F. Supplemental Materials: “Computer Security Customer Support for Survivors of Intimate Partner Violence” . . . . .</b>	<b>355</b>
F.1 Focus Group Protocol: IPV Professionals . . . . .	355
F.2 Focus Group Protocol: Support Agents . . . . .	358
<b>G. Supplemental Materials: Nudges for Encouraging Password Changes After Data Breaches . . . . .</b>	<b>364</b>
G.1 Survey Material . . . . .	364
G.1.1 Screening Survey . . . . .	364
G.1.2 Main Survey . . . . .	371
G.1.3 Follow-up Survey . . . . .	375
G.2 Cognitive Walkthrough Protocol . . . . .	378
<b>BIBLIOGRAPHY . . . . .</b>	<b>381</b>

## LIST OF FIGURES

### Figure

3.1	Terminology used to present relative frequency of themes. . . . .	40
3.2	Status of FTC’s suggested actions [501] taken or not taken by participants. . . . .	49
4.1	Sample breach information shown to participants. . . . .	63
4.2	Frequency of the leaked data types for 189 breaches, excluding email address (appears in all breaches). 44 other types occurring twice or fewer. . . . .	70
4.3	Overall concern about the breach and levels of concern for the 13 most commonly leaked information types in our sample breaches. . . . .	79
4.4	Code frequencies for feelings after first learning about a breach ( $n = 792$ ); red bars indicate negative feelings, gray neutral, blue positive, according to Emolex ratings [350]. . . . .	80
4.5	Intention to take actions within the next 30 days. . . . .	83
5.1	Distribution of FRES, mapped onto Flesch’s 7-level ranking system. . . . .	96
5.2	Examples of data breach notifications using structural headings when introducing “what happened” in the breach. . . . .	98
5.3	Frequency of recommended protective measures. . . . .	102
6.1	Distribution of response options for each practice. The practices are sorted by full adoption rates in descending order. . . . .	120
7.1	Common themes that emerged in one of the brainstorming sessions for an icon that conveyed <i>opting-out</i> . . . . .	144
7.2	The percentage of participants selecting an icon that best conveys there’s an option to (1) “tell websites ‘do not sell my personal information’” (blue); and (2) “make choices about the use of my personal information” (red) in the preliminary testing. . . . .	149
7.3	Promising icons from preliminary testing in their refined versions. . . . .	150
7.4	The percentage of participants selecting for an icon that best conveys that there’s an option to “tell websites ‘do not sell my personal information’” (blue); and “make choices about the use of my personal information” (red) in the refined testing. . . . .	152
7.5	Distribution of expectations in response to “What do you think would happen if you clicked on this [link]?” in our link text pre-study. . . . .	158

7.6	Icon and link text presented on a fictitious online shoe retailer webpage used in the icon-text combination evaluation. The icon and link text were highlighted with an orange rectangle to attract participants' attention. Shown is the condition combining <i>Stylized-Toggle</i> (icon) and "Privacy Options" (link text). . . . .	162
7.7	Distribution of Likert responses across conditions in icon-text combinations evaluation. . . . .	165
7.8	Our stylized toggle, OAG's proposed opt-out button, its variant, and the iOS switch button. . . . .	167
7.9	Common expectations of what would happen after clicking based on open-ended responses in conditions with <i>Stylized-Toggle</i> ( $n=137$ ), <i>CalAG-Toggle</i> ( $n=134$ ) and <i>CalAGX-Toggle</i> ( $n=132$ ). . . . .	170
8.1	Portions of three representative customer support chats from our dataset ("C" is customer, "S" support agent). We removed specific details and slightly altered wording to protect anonymity. . . . .	188
9.1	An overview of our experiment's procedure. . . . .	228
9.2	The threat appeal (1) features negative things that could happen to affected users when a data breach involves login credentials [180, 514] (matching threat severity) and (2) addresses optimism bias and hyperbolic discounting that people may experience after data breaches [260, 602] (matching threat vulnerability). . . . .	229
9.3	The coping appeal (1) builds the connection between changing the password and reducing the chance of account compromise (matching response efficacy), (2) gives a list of concrete steps to take, including URLs when applicable (matching self-efficacy), and (3) clarifies that the effort "only takes a few minutes" and is "easy" (matching response costs). . . . .	229
9.4	Example breach information shown to participants. . . . .	233

## LIST OF TABLES

### Table

1.1	Overview of work included in this dissertation. . . . .	5
2.1	A taxonomy of online privacy-protective behaviors. . . . .	14
2.2	Heuristics and biases in consumers' privacy decision-making. . . . .	22
3.1	Demographics of participants, and scores of NFEC financial decision [362] and CFPB financial well-being scales [92]. . . . .	38
4.1	Participant demographics and breach status ( $n=413$ ). . . . .	69
4.2	Logistic regression for breach status of an email address (leaked vs. not leaked). . . . .	72
4.3	Quasi-poisson regression regarding the number of breaches per email address. . . . .	72
4.4	Logistic regression regarding prior breach awareness. . . . .	77
4.5	Logistic regression on taking actions. . . . .	84
6.1	Security, privacy, and identity theft protection practices included in our study. . . . .	112
6.2	Response options relating to adoption for our survey questions. . . . .	115
6.3	Gender, education, age and income demographics of survey participants. Census statistics from [527, 528, 529, 530] as of 2018. . . . .	116
6.4	Results of the main regression model, excluding interaction terms between the practice domain and other variables. . . . .	123
6.5	Top coded reasons for partial adoption and abandonment. . . . .	126
6.6	Participants' most frequent reasons for incomplete adoption and abandonment of security practices. . . . .	127
6.7	Participants' most frequent reasons for incomplete adoption and abandonment of privacy practices. . . . .	127
6.8	Participants' most frequent reasons for incomplete adoption and abandonment of identity protection practices. . . . .	128
7.1	Icon themes that emerged in ideation sessions for each choice-related concept, and the corresponding icons included in our preliminary testing. . . . .	144
7.2	Link texts tested in the link text pre-study. . . . .	154

9.1	Gender, age, education, income, race/ethnicity compositions among participants of the main survey ( $n=1,386$ ). Census statistics from [531, 532, 533, 534] as of 2022. Some percentages do not add up to 100% due to non-reporting. . . . .	240
9.2	Top breaches in our sample. Freq. indicates how many times the breach was featured in the main survey. . . . .	241
9.3	Pairwise comparisons between control vs. treatment for the percent of participants who reported intending to change the password in the main survey. . . . .	241
9.4	Pairwise comparisons between control vs. treatment for the percent of participants who reported having changed the password in the follow-up survey. . . . .	242
9.5	Mean, median, and standard deviation of participants' rating of the PMT constructs, divided by conditions. We adapted scales from prior work [49, 484] in measuring these variables: threat severity (Q28), threat vulnerability (Q29), response costs (Q32) were measured using 5-point Likert scales, taking the median. Response efficacy (Q30) and self-efficacy (Q31) were measured using a single 5-point Likert-type item. . . . .	245
9.6	Logistic regression for predicting password change intention in the main survey. $n=725$ after excluding incomplete or not applicable responses. Cox and Snell $R^2=0.14$ . Model $\chi^2(21)=111.05$ , $p<.001$ .	249
9.7	Logistic regression for predicting password change action in the follow-up survey. $n=615$ after excluding incomplete or not applicable responses. Cox and Snell $R^2=0.21$ . Model $\chi^2(22)=141.59$ , $p<.001$ . .	251
9.8	Top reasons for intending to change or not change the breached password in the main survey. . . . .	255
9.9	Top reasons for changing or not changing the breached password in the follow-up survey. . . . .	259
9.10	Additional actions participants took in response to learning about the breach. . . . .	262
E.1	Age, gender, education, state residence demographics of participants as well as their familiarity with CCPA in each study. . . . .	348
E.2	Participants' coded open-ended responses to "What does this symbol communicate to you?" from conditions in which the icon was shown without a link text in the icon preliminary testing, along with a code's number of occurrences. Interpretations that align with the icon's intended meaning are bolded. . . . .	349
E.3	Participants' coded open-ended responses to "What does this symbol communicate to you?" from conditions in which we showed the icon without a link text in the refined icons study, along with a code's number of occurrences. Interpretations that align with the icon's intended meaning are bolded. . . . .	350

E.4 Regression results for the four binary dependent variables (conveys *choice*, *privacy*, *misconceptions*, or *do-not-sell choices*) coded from participants’ open-ended expectations. . . . . 351

E.5 Regression results for the scenarios: “It will take me to a page with choices about the sale of my personal information” (Do-Not-Sell Choices), “It will give the website permission to sell my personal information” (Give Sell Permission), and “It will take me to a page with choices about how my personal information is used and shared by the website” (Privacy Choices). . . . . 352

E.6 Regression results for the binary dependent variables: conveys a *misconception* and perceived as a *toggle control*, coded from participants’ open-ended expectations. . . . . 353

E.7 Regression results for the scenarios: “It will immediately change the setting on this website from ‘Sell My Personal Information’ to ‘Do Not Sell My Personal Information’ ” (Do-Not-Sell Switch), and “It will take me to a page with choices about the sale of my personal information” (Do-Not-Sell Choices). . . . . 354



## LIST OF APPENDICES

### Appendix

A.	Supplemental Materials: “Consumer Reactions to the 2017 Equifax Data Breach” . . . . .	290
B.	Supplemental Materials: “Consumer Reactions to Data Breaches that Affected Them” . . . . .	303
C.	Supplemental Materials: “An Empirical Analysis of Data Breach Notifications” . . . . .	314
D.	Supplemental Materials: “Adoption and Abandonment of Protection Practices” . . . . .	322
E.	Supplemental Materials: “Icons for Conveying Privacy Controls” . . . .	338
F.	Supplemental Materials: “Computer Security Customer Support for Survivors of Intimate Partner Violence” . . . . .	355
G.	Supplemental Materials: Nudges for Encouraging Password Changes After Data Breaches . . . . .	364

## ABSTRACT

As much as consumers express desires to safeguard their online privacy, they often fail to do so effectively in reality. In my dissertation, I combine qualitative, quantitative, and design methods to uncover the challenges consumers face in adopting online privacy behaviors, then develop and evaluate different context-specific approaches to encouraging adoption.

By examining consumer reactions to data breaches, I find how consumers' assessment of risks and decisions to take action could be subject to bounded rationality and potential biases. My analysis of data breach notifications provides another lens for interpreting inaction: unclear risk communications and overwhelming presentations of recommended actions in these notifications introduce more barriers to action. I then turn to investigate a broader set of privacy, security, and identity theft protection practices; the findings further illuminate individual differences in adoption and how impractical advice could lead to practice abandonment.

Leveraging these insights, I investigate how to help consumers adopt online privacy-protective behaviors in three studies: (1) a user-centered design process that identified icons to help consumers better find and exercise privacy controls, (2) a qualitative study with multiple stakeholders to reimagine computer security customer support for serving survivors of intimate partner violence, and (3) a longitudinal experiment to evaluate nudges that encourage consumers to change passwords after data breaches, taking inspiration from the Protection Motivation Theory. These three studies demonstrate how developing support solutions for consumers requires varying approaches to account for the specific context and population studied. My

dissertation further suggests the importance of critically reflecting on when and how to encourage adoption. While inaction could be misguided sometimes, it could also result from rational cost-benefit deliberations or resignation in the face of practical constraints.

**Thesis statement.** By better understanding the challenges consumers face in adopting online privacy-protective behaviors, we can develop approaches that help consumers better protect their privacy while addressing consumers' diverse needs and practical constraints.

## CHAPTER I

### Introduction

Privacy has become one of the critical issues in the age of information. Through emails, texts, social media, and more, people reveal information to one another, to the government, or to commercial entities, making the line between private and public increasingly vague [5]. The massive flows of information, while providing numerous benefits, also fuel institutional discrimination [8], behavioral manipulation [605], harassment and abuse [167, 513], and many other negative consequences.

Prior work has documented ample evidence for elevated privacy concerns among consumers: consumers not only care about their privacy but also take action to protect their privacy [7]. For example, a 2019 Pew Research Center survey revealed that most US respondents felt “concerned, confused, and a lack of control over their personal information” harvested by companies and the government [33]. However, taking action does not mean taking action *effectively*. For instance, despite fears of account compromises, most consumers still reuse passwords [105, 158, 388], which means one password exposed in a data breach could be the key to many other online accounts they have that use the same or similar passwords. Similarly, despite concerns about companies’ data practices, most consumers agree to privacy policies without reading them in detail [33] and struggle to use privacy choices (e.g., opt-outs or data deletion options) provided by websites [197].

It is crucial to resist blaming consumers for inaction, as privacy self-protection is inherently difficult. Inaction could be rational from an economic perspective due to the marginal benefits and far greater indirect costs in the form of effort [218]. With surveillance capitalism as a common business model [606], companies often use dark patterns that trick consumers into disclosing more information than intended [186], making privacy self-protection a rigged game. Adopting privacy-protective behaviors becomes more challenging for already vulnerable and marginalized populations when they face unique challenges and competing needs [335, 562].

Still, there are many reasons to motivate consumers to take action when they need to. Harms as a result of privacy violations take many shapes and forms [83], such as physical harms (e.g., physical injuries enabled by online stalking), economic harms (e.g., financial loss due to identity theft), reputational harms (e.g., in the case of revenge porn), and psychological harms (e.g., emotional distress after a data breach). Minor harms happened to individuals can aggregate into substantial harm to the society, such as in the case of chilling effects [258]. Furthermore, certain populations such as undocumented immigrants [192] and sex workers [335] often face heightened threats and consequences of privacy violations.

I argue that by better understanding the challenges consumers face in adopting online privacy-protective behaviors, we can develop approaches that help consumers better protect their privacy while addressing consumers' diverse needs and practical constraints. To this end, my dissertation seeks to address the following two questions:

**RQ1: What prevents consumers from adopting online privacy-protective behaviors?**

**RQ2: How can we develop context-specific approaches that encourage consumers to adopt online privacy-protective behaviors?**

In response to RQ1, I conduct research that contributes new knowledge of con-

sumers’ reactions to data breaches and reasons for inaction, readability and usability issues of breach notifications sent to consumers, and the broader struggles consumers face in adopting expert advice for security, privacy, and identity theft protection. The findings showcase the bounded rationality consumers have in privacy decision-making, as well as systemic issues in technology design that can limit consumers’ privacy-protective behaviors.

In response to RQ2, I explore how to develop approaches that encourage consumers to adopt privacy-protective behaviors in three studies with complementary methods. First, we can base the approach on consumers’ preferences directly, as my collaborators and I designed and evaluated icons to help consumers better find privacy controls based on their preferences via a series of MTurk experiments. However, seeking input from direct users is not always possible and could even introduce additional risks or retraumatization. As this is the case for survivors of technology-enabled intimate partner violence, I identified computer security customer support—an existing avenue that survivors use for seeking help—and engaged with IPV and customer support professionals to obtain their insights on how customer support could better address survivors’ needs. While the previous two studies both involved seeking input from consumers or relevant stakeholders, in the last study, I demonstrated how to take inspiration from established behavioral change theories and known issues in building the intervention, as I developed and evaluated nudges based on the Protection Motivation Theory that encourage consumers to change passwords after data breaches. My final chapter synthesizes the takeaways from the different studies, such as the need for a more nuanced view of inaction and the individual differences in consumers’ adoption of privacy-protective behaviors.

A few term-related clarifications for this dissertation: I use the term “consumers” to highlight the target audience (i.e., an average person that consumes rather than develops technology, in contrast to professionals such as developers and system ad-

ministrators). I choose “consumers” rather than “users” to include non-users of a system or interface that could simultaneously be affected. The terms “privacy” and “security” are highly related but have nuanced differences: privacy emphasizes one’s control over their personal information and how it is used, whereas security emphasizes how one’s personal data is protected [40]. As I will discuss in Chapter II, while some of the actions I cover in my dissertation are security-related, I argue that all of these actions have privacy implications—e.g., data breaches are security incidents, but many potential consequences of being affected by data breaches corresponding to existing taxonomies of privacy harms [83, 475]. Lastly, I limit the scope of privacy-protective behaviors to the online context since online consumers face a greater extent of information asymmetry and a lack of cues that indicate privacy violations [5, 7]. In contrast, many offline privacy-protective behaviors (e.g., lowering our voice during intimate conversations or leaving a group of people to take a personal call) happen more ubiquitously with little conscious awareness [22].

For the rest of this chapter, I will briefly summarize individual projects included in my dissertation. Table 1.1 shows the research question each chapter corresponds to, methods used, and venue (if published). I write in the first person when presenting my own arguments or reflections (e.g., in the discussion chapter), and I use “we” when writing about specific projects to acknowledge the collaborative nature of the work presented.

## 1.1 Understanding Adoption

Motivating consumers to adopt online privacy-protective behaviors requires a fundamental understanding of what online privacy-protective behaviors consumers adopt or do not adopt and why. I conducted four studies to address this question, including three studies focusing on data breaches and one study that examines the adoption and abandonment of protective behaviors in privacy and adjacent contexts.

Chap.	Topic	Methodology	Venue
III	Consumer reactions to the 2017 Equifax data breach (RQ1)	interview	SOUPS 2018
IV	Consumer reactions to data breaches that affected them (RQ1)	survey	USENIX Security 2021
V	An empirical analysis of data breach notifications (RQ1)	content analysis	CHI 2019 and IEEE S&P magazine
VI	Adoption and abandonment of security, privacy, and identity theft protection practices (RQ1)	survey	CHI 2020
VII	Icons for conveying privacy controls (RQ2)	design workshop, experiment	CHI 2021
VIII	Computer security customer support for survivors of intimate partner violence (RQ2)	content analysis, focus group	USENIX Security 2021
IX	Nudges to encourage password changes after data breaches (RQ2)	experiment	Manuscript in prep.

Table 1.1: Overview of work included in this dissertation.

### 1.1.1 Consumer Reactions to the 2017 Equifax Data Breach: An Interview Study

Data breaches, defined as the unauthorized exposure, disclosure, or loss of personal information, are security incidents that also bear substantial privacy implications [476]. The types of breached data, such as financial and health information, can be highly sensitive and touch on an individual’s intimate details. In Chapter III, I present a qualitative study in which we conducted 24 semi-structured interviews with US-based consumers focusing on the 2017 Equifax data breach. The breach put 146 million US consumers—almost half of the nation’s population—at risk of identity theft, as the breached data types included social security numbers, driver’s license numbers, and other personally identifiable information.

We found that most participants were aware of this breach and the associated risks of identity theft and privacy invasion. Nevertheless, less than half of our participants had checked whether they were affected on Equifax’s website, and even fewer took concrete protective measures such as freezing their credit reports. Participants’ inac-



tion occurred due to perceived costs associated with protective measures, optimism bias in estimating one’s likelihood of victimization, and a general tendency to delay action until harm has occurred. Based on the findings, we draw public policy recommendations that give consumers free and more frequent access to credit reports. We also outline the need for fixing usability issues related to existing protective measures and the importance of educational efforts.

### **1.1.2 Consumer Reactions to Data Breaches that Affected Them: A Survey Study**

Going beyond the 2017 Equifax data breach, how do consumers react to data breaches in general? Will things be different if consumers know that the breach has compromised their personal information? In Chapter IV, we addressed these questions via an online survey ( $n=413$ ) by presenting each participant with up to three data breaches known to have exposed their email addresses and asked about their reactions. This method ensured high ecological validity in collected responses and helped mitigate recall bias as participants were confronted with specific real-life breaches that affected them.

We found that 74% of participants were unaware of the breaches displayed to them even though most of our participants were affected by at least one breach. While participants expressed various emotions upon learning about the breach (e.g., upset, surprised, and fatigued), the overall concern level remained low, and most participants believed the breach would not impact them. Our regression analyses further showed that prior awareness and greater concern were significantly correlated with a higher likelihood of taking action, suggesting that raising awareness and concern is key to motivating action. Our findings highlight that we need better mechanisms to raise consumers’ awareness of breaches and stricter regulatory requirements that hold breached organizations accountable for providing more effective protective measures.

### **1.1.3 An Empirical Analysis of Data Breach Notifications**

Many consumers first learn of a breach from notifications sent by affected businesses [1, 329]. Breach notifications are also legally required in many jurisdictions across the world [48]. Nonetheless, consumers' inaction, fatigue, and low awareness of data breaches imply that current ways of informing consumers about breaches might be problematic. In Chapter V, I present a content analysis of 161 consumer-facing breach notifications focusing on their readability, structure, risk communication techniques, and presentation of recommended actions.

We found that most notifications were lengthy and required advanced reading skills compared to materials addressed to the general public. Vague terms such as “may” and “likely” were used to describe the breach’s compromised data and potential consequences, which may downplay the breach’s risks to consumers. Moreover, available protective actions were often described in lengthy paragraphs with little information regarding an action’s effectiveness or urgency. From these findings, we provide concrete ideas of nudges in data breach notifications that could help consumers better understand potential risks and become motivated to take action. Meanwhile, future data breach notification laws could incorporate these design recommendations—or even make them mandatory—to create stronger incentives for businesses to deliver more usable and useful breach notifications to consumers.

### **1.1.4 Adoption and Abandonment of Security, Privacy, and Identity Theft Protection Practices**

As reacting to data breaches is a small component of one’s online privacy posture, in Chapter VI, I document a survey ( $n=902$ ) in which we took a broader perspective by examining a comprehensive set of expert-recommended practices in privacy and adjacent fields. We selected 30 practices from past literature, including 12 related to security protection, 12 related to privacy protection, and six related to identity

protection. We also looked into when/why a practice had been abandoned and demographic differences in practice adoption and abandonment.

We found that security practices in our sample received wider adoption than privacy or identity theft protection practices. Participants preferred manual practices that fully rely on users' cognitive efforts (e.g., checking the URL when visiting a website) and automated practices that require no effort after the initial adoption of a technology (e.g., using antivirus software) over practices requiring recurring user interactions with a system or interface (e.g., using a password manager). Furthermore, participants' gender, education, technical background, and prior negative experience were significantly correlated with the level of adoption (no adoption, partial adoption, or full adoption). Some practices were inconsistently adopted or abandoned when participants perceived them as irrelevant, impractical or overrode them with subjective judgment. Our findings echo prior work by illuminating the socioeconomic differences in adopting privacy-protective behaviors and usability issues being a substantial barrier to adoption regardless of context. More fundamentally, focusing on usability issues alone is insufficient, as privacy self-protection becomes prohibitively difficult when the recommended action is unrealistic, inconvenient, or communicated with obfuscated language.

## **1.2 Improving Adoption**

Having gathered insights on the struggles consumers face in adopting privacy-protective behaviors, the second part of my dissertation tackles the question: How can we help consumers adopt privacy-protective behaviors? I answer this question in three distinct yet interconnected contexts. First, the approach to encouraging adoption can be directly built on consumers' preferences, as I demonstrated how to do this by designing and evaluating icons to convey privacy controls (Chapter VII). The approach could also be informed by input from relevant stakeholders, especially when

conversations with the direct users could introduce additional risks or retraumatization. I demonstrated how to engage with stakeholders in making computer security customer support better serve survivors of technology-enabled intimate partner violence who face unique threats and challenges (Chapter VIII). Lastly, we could draw from established behavioral change theories and known impediments to action in finding the approach, as I demonstrate how to develop nudges that encourage password changes after data breaches and evaluate them with consumers in a longitudinal experiment with high ecological validity (Chapter IX).

### **1.2.1 Icons and Link Texts for Conveying Privacy Choices**

As shown in Chapter VI and in other pieces of my prior work not included in this dissertation [197, 198], consumers struggle to find and exercise privacy controls effectively. One way to increase the adoption of privacy controls is by making privacy interfaces more usable and useful [438]. In Chapter VII, I describe a series of design workshops and online experiments in which we designed and evaluated novel icons and accompanying link texts that convey privacy controls. We focused on generic privacy controls and opt-outs of the sale of personal information—a privacy control required by the California Consumer Privacy Act (CCPA).

We identified that a blue stylized toggle icon paired with “Privacy Options” best conveyed general privacy choices; the CCPA-mandated link texts (“Do Not Sell My Personal Information” and “Do Not Sell My Info”) effectively conveyed do-not-sell opt-outs in combination with most icons. Our findings suggest that icons for conveying privacy controls should be rooted in simple and familiar concepts, as evidenced by the toggle icon’s performance over other candidates. Our findings also highlight link text’s importance in aiding comprehension and reducing misconception, especially for new icons. This research has generated concrete impacts on policymaking, as the California Office of Attorney General (OAG) cited our work in CCPA’s rulemak-

ing process and adopted the blue toggle icon we recommended in the official CCPA regulations [378].

### **1.2.2 Computer Security Customer Support for Survivors of Intimate Partner Violence**

Chapter VII demonstrates how to develop privacy icons and link texts based on consumers' preferences. In other contexts, helping consumers adopt privacy-protective behaviors requires adapting existing support infrastructures to ensure that those who need privacy-related support and guidance can get it when needed. In Chapter VIII, I present a case study illustrating how to incorporate stakeholders' insights in improving computer security customer support to address the unique needs and challenges faced by survivors of intimate partner violence (IPV). Rather than engaging with IPV survivors directly on this topic, which could cause unnecessary re-traumatization, we conducted focus groups to elicit insights from two sets of stakeholders: IPV professionals (e.g., social workers and lawyers) and customer support professionals.

Our focus groups with IPV professionals surfaced suggestions to make computer security customer support more attuned to survivors' threat models and safety risks, such as using trauma-informed language, avoiding promises to solve problems, and making referrals to external resources. In focus groups with customer support practitioners, we confirmed the practicality of these suggestions while uncovering real-world constraints of deploying customer support, such as challenges in identifying potential survivors and frontline agents' limited capacity. As a practical outcome of this research, we have developed IPV-focused training materials for customer support agents and shared them with participating security companies.

### 1.2.3 Nudges for Encouraging Password Changes After Data Breaches

Chapters VII and VIII both involve seeking input from direct users or relevant stakeholders in developing the approach to encouraging adoption. In Chapter IX, I demonstrate how known psychological constructs and well-established behavioral change theories could be another source of inspiration for finding the approach that helps consumers better understand risks and guides consumers toward taking action. Specifically, we developed nudges based on the Protection Motivation Theory (PMT) to encourage consumers to change passwords after data breaches, including a threat appeal that highlights the risks of compromised passwords and a coping appeal that provides links and instructions for changing the breached password. We then evaluated the nudges' impact on participants' password change intention and behavior in a longitudinal experiment.

We found that a threat appeal alone was most effective at motivating password change intention, and threat and coping appeals combined were most effective at motivating password change behavior. Both generated a statistically significant yet small difference compared to the control condition. Participants further described challenges they experienced in attempting to change their passwords and alternative strategies when they did not see the benefit of changing the password for the particular breach site. Our findings add to prior PMT literature by highlighting the necessity of combining threat and coping appeals in motivating password changes. Moreover, our results contribute a new view in understanding and theorizing inaction—we should not view inaction as a negative outcome by default, but rather connect the inaction to consumers' practical constraints and competing needs they have in life.

## CHAPTER II

### Related Work

In this section, I review related work that serves as background knowledge for the research questions I tackle in my dissertation: a taxonomy of online privacy-protective behaviors (RQ1), factors behind adoption and non-adoption (RQ1), and mainstream approaches to encouraging adoption (RQ2).

#### 2.1 A Taxonomy of Online Privacy-Protective Behaviors

Prior research has mostly examined online privacy-protective behaviors in individual contexts, such as social media [79, 249, 596], targeted advertising [469, 597], and the Internet of Things [132, 581], indicating the need for a comprehensive taxonomy. Example taxonomies in prior work include Son & Kim [477], Caine [66], Coopamootoo [93], and Redmiles et al. [415]. Specifically, Son and Kim’s work is based on a literature review and has six categories: refusal, misrepresentation, removal, negative word-of-mouth, complaining directly to companies, and complaining indirectly to third parties [477]. However, the taxonomy developed in 2008 does not include many more recent privacy-enhancing technologies (PETs). The other three taxonomies are all based on empirical user studies. Caine’s taxonomy is based on focus group data, uses a simpler categorization (avoidance, modification, and alleviatory), and includes examples of offline privacy-protective behaviors as well [66].

Coopamootoo’s taxonomy is based on survey data, focusing on PETs and the types of protection afforded (e.g., anonymity, tracking prevention, or data leaking prevention) [93]. Redmiles et al.’s taxonomy is based on scraped web data about expert advice as well as input from both experts and end-users, but the taxonomy focuses more on computer security and does not include any high-level categorization.

Below I present my taxonomy of online privacy-protective behaviors (Table 2.1) as an organizing framework for this dissertation. In line with prior work [23, 151, 400, 429, 593], I divide behaviors based on the coping mechanism’s nature (approach/active vs. avoidance/passive). Example approach/active protective behaviors include problem-solving and seeking social support [593], whereas avoidance/passive behaviors usually means distancing oneself from the source of threats [220]. Inspired by Son & Kim’s taxonomy [477], I use the level of social interactions required (private vs. public) as a second dimension for categorization. Private behaviors are those one can implement autonomously, usually because they post personal information themselves or have some control over the information. Public behaviors are those that require the cooperation of and potential confrontation with others. I then fill out the taxonomy with specific examples extracted from prior taxonomies [66, 93, 415, 477] as well as studies of online privacy-protective behaviors in individual contexts.

Notably, one can choose *not* to cope with privacy risks by continuing their regular tech use and accepting or ignoring known risks. The inaction, limited actions, or inconsistent actions that individuals take in response to privacy concerns could stem from a sense of resignation that surveillance is inescapable, and prior work has contributed ample evidence of the “digital resignation” phenomenon [129, 207, 317, 523].

### **2.1.1 Coping by Avoidance**

In avoidance-driven behaviors, consumers disengage from activities that could lead to the collection, use, and dissemination of their data [66]. This behavior could



Mechanism	Interaction	Category	Behavior	Examples
avoidance	private	n/a	avoid technology use	avoid going online in general; avoid certain types of technology
			limit technology use	limit use of social media; avoid using public or shared devices
approach	private	checking	read privacy notices	read privacy policies and terms of service; read cookie banners
			review privacy settings	review profile visibility on social media; review mobile app permissions
			track digital footprint	check exposure to data breaches; check spyware; egosurfing; review archives of personal data; use dark web monitoring services
approach	private	modification	adjust privacy settings	disable third-party cookies; enable “do not track”; opt out of targeted ads; restrict post visibility; switch off location tracking
			erase digital footprint	clear cookies and browsing history; delete unused accounts; delete old emails or social media posts; opt out of mailing lists
			selective disclosure	limit the amount of information shared; self-censor in social media posting
			misrepresentation	provide inaccurate contact info; use pseudonyms; use different personas online
			use PETs	PETs for anonymity: Tor, proxy, virtual machine, encryption; PETs for tracking prevention: DuckDuckGo, anti-tracking browser extensions, privacy browsing mode; PETs for filtering communications: ad-blocker, firewall; PETs for data leakage prevention: anti-spyware and anti-malware tools
approach	public	social support	advice exchange	teens seeking advice on online behaviors from parents; spread negative experiences with companies that disrespect user privacy
			collective management	family members sharing IoT devices; romantic partners sharing online banking accounts
			observation	see what information others disclose; see what security features friends adopt on Facebook
approach	public	confrontation	activism	public outcry after big news events, such as the Snowden leaks
			complaint	file a complaint with government agencies after a data breach; ask friends on Facebook to take down offensive photos
			flaming	flame entities that send unsolicited emails

Table 2.1: A taxonomy of online privacy-protective behaviors.

stem from pessimistic views of threats posed by information technology [129], overwhelming requests for personal data [311, 477], or adverse consequences caused by privacy violations [171]. Consumers might avoid engagement with online activities in general [93, 421] or avoid certain technologies that have privacy-invasive features and can trigger privacy concerns, such as social media [93, 306] and smart home devices [283].

A more calculated form of avoidance-driven behavior is to use technology but with reduced frequency, e.g., avoiding public computers in places like hotels and libraries due to malware risks, shoulder surfing, and WiFi spoofing [171]. However, limiting technology use or avoiding it altogether can cause considerable inconvenience to the non-user [171] who faces the loss of social and economic benefits provided by an Internet-connected world [207, 421]. Furthermore, not everyone can find alternatives that provide the same utilities as the abandoned technology. Even though using public devices is not advised, it might be one of the few (if not only) options for certain populations such as homeless people [145] or people in developing regions [203, 547].

### 2.1.2 Coping by Approach

In approach-driven behaviors, consumers actively seek to modify or eliminate the conditions that introduce privacy risks [66]. I further divide approach-focused behaviors into four categories: checking the privacy status quo, modifying privacy management behaviors, confronting privacy violators, and seeking social support.

**Checking.** Through checking, consumers learn about potential privacy risks and determine the required course of action. An example source for checking is the privacy policies and terms of service agreements many websites and mobile apps have in response to the notice and choice regulatory framework [418]. While privacy notices are mostly known as privacy policies [347], they can take other forms such as cookie

banners [539] and LED lights that indicate active recording from a device [437]. Consumers can also review their privacy settings, ranging from profile visibility on social media [594] to permissions given to mobile apps [21].

Many options exist for tracking one’s digital footprint beyond making use of privacy notices and choices. For instance, consumers can require companies to send them an archive of personal data processed by the platform under the European Union General Data Protection Regulation’s “Right of Access” [144]. For cross-platform tracking of digital footprints, one can search for their name on a search engine and review results (i.e., egosurfing) [309, 415] or use dark web monitoring services [575]. Tracking digital footprint is also relevant when there are privacy violations, e.g., services like Have I Been Pwned [232] and Firefox Monitor [354] track the occurrence of data breaches and notify sign-up users when they are affected. Lastly, there are tools for spyware detection and cleanup in response to the evolving landscape of spyware.

**Modification.** While checking is mostly about learning about privacy risks, modification refers to behaviors that help alleviate existing risks. By adjusting privacy settings, consumers can, for example, restrict content visibility to only certain contacts [486] or opt out of targeted advertising [295]. Consumers can also more fundamentally change the amount of information collected by service providers via anti-tracking settings, such as by enabling the “Do Not Track” option in browser settings [415], blocking third-party cookies, and switching off location tracking [93]. Regarding the modification of digital footprints, consumers can make use of account deletion options [357, 444], delete old emails or social media posts [357, 445, 560], clear cookies and browsing history [93, 311], or opt out of mailing lists [93, 118].

Modification also applies to online disclosure, as consumers can reduce the amount of information given to companies (i.e., selective disclosure) [415] or provide inaccurate information (i.e., misrepresentation) [477]. Example behaviors for selective disclosure

include reducing posting frequency, being wary of friend requests, or deliberately choosing not to share controversial contents [79, 465]. Example behaviors for misrepresentation include using pseudonyms [93, 311], fabricating contact information in creating online accounts [381, 477], or using separate accounts for different activities to mitigate cross-platform inferences [415]. Notably, both selective disclosure and misrepresentation require constant vigilance from users and may be hard to follow consistently [603]. Both behaviors could result in inconvenience: e.g., missing important notifications due to incorrect contact details [322] as well as fewer opportunities for self-representation and building social capital [137].

Lastly, a large family of online privacy-protective behaviors centers around using privacy-enhancing technologies (PETs)—algorithms, tools, and applications built on cryptography and computer security research designed to support privacy and data protection [93, 496]. Coopamootoo divides PETs into four types depending on the goal of protection: preserve anonymity (e.g., Tor, proxy, virtual machine, and encryption), prevent tracking (e.g., the DuckDuckGo search engine, anti-tracking browser extensions, and private browsing mode), filter communications (e.g., adblocker and firewall), and prevent data leakage (e.g., anti-spyware and anti-malware tools) [93]. PETs can be built-in features in browsers and applications, such as in the case of encryption, while other PETs like Tor and VPN are standalone technologies. In addition, many PETs have alternatives that do not require adopting the technology. For instance, tools like “Sign in with Apple” [30] and “Firefox Relay” [355] facilitate the generation of unique and random email addresses during account registration, but this can be done manually too. While most PETs are free, they have not acquired a large user base because they are still new concepts to many consumers, require advanced technical skills, or have usability issues [93, 603].

**Social support.** Past research has shown the role of social influence on consumers' privacy attitudes and behaviors [9, 107, 138, 340, 568]. Below I discuss three ways of social support: advice exchange [411], observation [107], and collective management of digital resources [568].

Advice regarding privacy self-protection strategies can come from various sources, including media, workplaces, friends, family members, service providers, and reputable organizations [248, 293, 340, 411]. Different sources can influence the advice recipient's behavior differently [340, 411], and the influence is further moderated by the recipient's privacy concerns and self-efficacy [45, 340, 387]. In addition to seeking advice, people may advise others in their social circle, e.g., informing acquaintances of companies with invasive privacy practices [477].

In contrast to informational social influence, which occurs when consumers seek information and guidance, normative social influence happens when individuals try to fit in with a group or look to others for cues on how to act in uncertain circumstances [108, 117]. Normative messages with social proof have been shown to motivate privacy-protective behaviors, e.g., the observation that many friends and acquaintances use a privacy feature on Facebook affects one's likelihood to adopt the feature [109]. Peer influence from observations can also mislead, rather than encourage, privacy-protective behaviors. For example, consumers may tend to divulge sensitive information when told that others have made sensitive disclosures [10].

Individuals sometimes form social groups and jointly navigate privacy decisions in sharing accounts and devices [568]. Such sharing behaviors could occur between family members, romantic partners, friends, neighbors, colleagues, and more [57, 324, 385, 488]. Sharing digital resources creates learning and intervening opportunities during teachable moments, such as exchanging knowledge about data breaches or other breaking news events [107, 568]. Nonetheless, sharing behaviors may lead to tensions when group members have conflicting interests or preferences or when the re-

relationship collapses [385]. As an example, older adults with cognitive impairment tend to rely on caregivers to manage their online accounts, and sharing personal information with caregivers creates issues of agency, autonomy, or embarrassment [341, 342]. Even though shared decision-making might be a common goal, most existing privacy controls are not fine-grained or context-adaptive enough to meet users' evolving needs [341, 385, 488].

**Confrontation.** Through confrontation, consumers negotiate with sources that generate privacy risks or violations. For example, consumers impacted by a data breach may express their frustration by contacting the breached organization, government agencies, or non-profit consumer protection organizations [425]. A social media user who does not want to be tagged in a photo may ask the photo owner to take it down or escalate the request to a group administrator [228, 311]. Furthermore, individual efforts can form collective activism [590], as in public outcry against the US government surveillance after the Snowden leaks [311] and against corporate surveillance after the Cambridge Analytica Data Breach [217, 427]. However, in provoking systemic changes, activism efforts often are confronted with organized lobbying of big tech companies that go against them [7, 116, 129, 523]. Sending highly negative messages, i.e., flaming [452], is a more extreme form of protest against privacy violations than complaining and activism. An example is flaming senders of unsolicited emails or spam callers.

## 2.2 Factors Behind Adoption and Non-Adoption

Prior privacy research has often drawn inspiration from behavioral theories in social science literature to explain the social and psychological processes behind adopting or not adopting online privacy-protective behaviors. The Protection Motivation Theory (PMT) [422], the Theory of Planned Behavior (TPB) [14], and the Technol-

ogy Acceptance Model (TAM) [112] are just a few examples. Individuals' decision-making process behind adoption is more nuanced than rational risk-benefit analysis, often subjected to information asymmetries, heuristics and biases, market power, and dark patterns [4, 7]. Below I summarize several key factors that inform my work.

**Attitude.** According to the Theory of Planned Behavior [14], attitude—how an individual thinks of a particular behavior, usually in the form of positive, negative, or neutral [98]—is one of the major determinants of behavioral intention. For example, privacy concern, the desire to keep personal information out of the hands of others [59], is a negative privacy-related attitude and has been extensively studied [87]. Another privacy-related attitude is trust: an individual's inherent tendency to rely on institutions (institutional trust) or other people (interpersonal trust) under risky conditions.

Attitudes can guide privacy decision-making, especially when situational cues for risk-benefit analysis are unavailable [25]. For instance, trust in an organization's integrity may raise consumers' intention to disclose information [123]. Nonetheless, past literature on the privacy paradox has demonstrated the complex relationship between privacy attitudes and behaviors: consumers may state high privacy concerns while being reluctant to adopt privacy protective behaviors [41, 177, 274, 474]. A popular explanation is that privacy concerns can be generic or context-dependent: one may care deeply about privacy in general but may not seek privacy protection depending on the benefits and costs in a specific situation [41, 474]. Furthermore, there are interplays between general attitudes and situational risk-benefit analysis in privacy decision-making [7]. For example, the pre-existing trust may interact with situational cues, such as the purpose of data use and who requests the data, in determining one's willingness to disclose health information [25].

**Emotions.** Unlike attitudes, emotions are not stable mental states learned through experience but temporary physiological conditions, often accompanied by changes in facial expressions, gestures, and posture [65]. Major emotion categories include anger, disgust, fear, guilt, joy, shame, and sadness, among others [441]. The “risk as feelings” theory highlights the central role of emotions in individuals’ cognitive processing of risks: individuals experience affective responses such as fear, dread, and anxiety in the face of risks, in addition to considering objective factors such as probabilities and outcomes; when the two processing systems produce different outputs, the affective-based system usually prevails [304].

Research on the interplay between emotions and privacy decision-making has only emerged in recent years, typically by measuring pre-existing emotional states and correlating them with other constructs, such as risk beliefs, trust, and intentions to disclose information [25, 298, 553]. Prior work has found that consumers underestimate information disclosure risks with an interface that elicits positive emotions [261, 394]. Major types of emotions triggered by privacy concerns include anxiety, anger, and disappointment; each emotion further leads to different coping behaviors [254]. In the case of virtual private networks (VPN), emotional considerations drive consumers’ adoption, such as fear of government and corporate surveillance [360]. Altogether, these studies suggest the importance of considering emotions in examining consumers’ privacy decision-making and behaviors, but more work in this space is needed [25, 254].

**Heuristics and biases.** Heuristics are mental strategies or rules people apply to simplify the evaluation of risks [525]. Although heuristics are valid in some circumstances, they lead to persistent biases, misjudged risks, or unwarranted confidence at other times [467]. Heuristics and biases were originally studied in behavioral economics experiments involving probability-based settings [525], but they also apply to



Challenge	Coping Mechanism	Heuristics & Biases	Examples
Information overload	Notice things primed in memory	availability bias [524]	judge online tracking practices based on a company's name rather than reading its privacy policy [537]
		representativeness bias [256]	perceive privacy invasions as low-probability events as they often are not observable [4]
	Focus on unusual or surprising things	affect bias [152]	underestimate the privacy risks from liked things [261, 394]
	Focus on changes	anchoring effect [525]	exposure to provocative selfies changes the perception of what information is appropriate to share [75]
	Analyze benefits and costs	framing effect [526]	change the perception of a privacy notice based on a reference point [12]
		rational ignorance [128]	skip privacy policies [370, 550]
Not enough meaning	Look for stories and patterns	anthropomorphism [213]	how much a technology or interface is anthropomorphized impacts privacy concerns [194, 591]
	Rely on stereotypes	authority bias [190]	trust the security and privacy of smart home devices provided by well-known companies [283]
	Project current mindset to the future	pessimism bias [221]	people feel resigned about protecting their personal information [129]
Need to act fast	Believe in one's own judgment	overconfidence [351]	claim to know about a privacy-enhancing technology but cannot answer simple questions about it [247]
		optimism bias [573]	users see themselves less vulnerable to privacy infringement than other individuals [36, 81, 602]
		illusion of control [402]	greater perceived control leads to higher willingness to disclose sensitive information [52]
	Focus on immediate things	hyperbolic discounting [82]	share information on social media for short-term bonding, but the shared content leads to future hiring discrimination [4, 8]
	Focus on invested efforts	loss aversion [257]	the willingness-to-accept price for selling personal information is consistently higher than the willingness-to-pay price for protecting such information [11]
		post-completion error [104]	forget to erase browsing history after using a public computer to perform sensitive searches [4]
Limited memory	Avoid irreversible decisions	status quo bias [255]	stick with default privacy settings [7]
	Store memories based on the experience	modality effect [130]	the presentation of privacy policies influences recall [263]
	Reduce information to key elements	recency bias	believe that recent content (e.g., on social media [349] and in cloud storage [267]) is more important and relevant

Table 2.2: Heuristics and biases in consumers' privacy decision-making.

privacy decision-making where the objective risk is hard to quantify and often unavailable to end-users [4, 5]. Table 2.2 summarizes major heuristics and biases with supporting evidence in privacy research. I organize them based on Benson’s taxonomy of four dilemmas in people’s decision-making [43]: information overload, lack of meaning, the need to act fast, and limited memory capacity.

**Norms.** According to Nissenbaum’s theory of privacy as contextual integrity, privacy violations occur when there is a breach of contextual norms governing information flows; the five parameters that help assess the privacy impact of information flows include data subject, sender, recipient, information type, and transmission principle [364]. In essence, norms affect people’s beliefs regarding what is private and what is public.

Consumers’ perception of social normative pressures or beliefs affects their ways of coping with privacy risks [14, 38]. For instance, consumers may turn to tech-savvier social connections for advice after a data breach [602]. Teenagers adhere to parental rules for disclosure on social media [543]. College students set their social media profiles as private after observing such behaviors from their friends [297]. In other cases, compliance with social norms can decrease one’s privacy protection. For instance, watching other individuals reveal sensitive personal information increases the likelihood that an individual will reveal it themselves [10]. Consumers may further be bounded by the rule of reciprocity, e.g., feeling compelled to disclose their relationship status on social media if their significant other does this [177].

**Perceived usefulness and ease of use.** TAM explains an individual’s usage of technological systems [112], thus providing a useful theoretical background for understanding the adoption of PETs [93]. Specifically, TAM posits that technology use is influenced by attitude toward the usage, which in turn is influenced by the perceived usefulness of the system (how the system increases the individual’s efficacy in dealing

with a task) and perceived ease of use (whether using the system is free of effort). Privacy research has used TAM to study users' acceptance of specific PETs such as encrypted messaging service [2], proxy client [204], and VPN [360]. Extrapolating from these examples, TAM broadly highlights usability as an important factor in adopting PETs and online privacy-protective behaviors. Adoption is more likely to occur when the involved technology or system is usable. Conversely, the absence of practical needs, access to alternatives, and burden in usage contribute to the abandonment of PETs [360, 603].

**Self-efficacy.** Self-efficacy refers to an individual's judgment of their self-capability to execute courses of action to attain a goal. Several mainstream behavioral science theories, including the Social Cognitive Theory [38], the Theory of Planned Behavior [14], and the Protection Motivation Theory [423] all highlight the role of self-efficacy in determining one's behavior. Applying self-efficacy to privacy research, consumers can carry out privacy-protective behaviors only when they have sufficient confidence or control. For example, empowering users of mobile apps with more granular controls over the information accessed helps users make better privacy decisions [558]. By contrast, a lack of perceived control explains the digital resignation phenomenon [129] when inaction happens not because consumers do not care but because they are resigned and convinced that surveillance is inescapable. Importantly, self-efficacy may inculcate confidence when it should not and produce paradoxical effects. More control may not guarantee enhanced privacy protection if people behave recklessly under the illusion of control [52].

### 2.3 Approaches to Improving Privacy-Protective Behaviors

Calo identified notice, nudge and code as three primary types of behavioral interventions: notice as the mere provision of information, nudge as soft paternalistic

mechanisms to guide individuals toward more beneficial outcomes, and code as rules or regulations to make negative behaviors difficult [68]. In line with Calo’s taxonomy [68], I discuss existing approaches that seek to facilitate the adoption of online privacy-protective behaviors: privacy notices and choices, privacy nudges, and privacy regulations.

### 2.3.1 Privacy Notices and Choices

Privacy notices and choices are products of the “notice and choice” regulatory framework [100]. The framework originates from Westin’s theory of privacy in Western democracies as the “claim of individuals, groups, or institutions to determine for themselves when, how, and to what extent information about them is communicated to others” [576]. As privacy refers to the ability to limit others’ access to one’s personal information, there should be mechanisms for individuals to understand where and under what conditions their information may flow and to control that flow [100]. Many data protection laws since then have embraced this concept by requiring businesses to provide consumers with notices about data practices and allow consumers to exercise choices over these practices [72]. Nevertheless, it is a consensus among privacy scholars, advocates, practitioners, and regulators that the notice and choice framework and resulting privacy notices and choices are insufficient and undesirable [7, 72, 100, 316, 418, 466, 556]. Walker identified five problems of privacy notices: overkill, irrelevance, opacity, non-comparability, and inflexibility [556]. Other work has surfaced dark patterns as an additional problem powered by businesses’ interest in capitalizing on consumers’ data [186, 321].

**Overkill.** McDonald and Cranor’s study conducted in 2008 estimated that if an individual were to read the privacy policy at every website visited once a year, they would spend 244 hours on average reading privacy policies per year [332]. Since

the ubiquity of smartphones, cloud services, and smart home technologies, the required time to read all privacy policies encountered is only likely to increase. To comply with regulations and avoid liability, companies tend to enumerate all potential uses and risks of information flow without highlighting specific practices that may unpleasantly surprise consumers [407]. Consequently, most privacy policies are long and complicated legal documents that require at least college reading levels to understand [29, 184].

**Irrelevance.** As a result of the overkill, most privacy notices are rendered meaningless. Prior work has consistently shown that most consumers accept terms of service agreements without reading them [67, 370]. While reading through privacy notices and choices might better inform consumers of product choices and shield them from privacy violations, the action also applies far greater indirect costs in the form of time and effort [218]. Most PETs that aim to give consumers more control over their personal information similarly have low awareness and adoption [93, 457].

**Opacity.** Ambiguity and vagueness are common in privacy notices and choices [417]. Privacy policies frequently use vague terms that dance around conditionality (“appropriate” “as needed”), generalization (“mostly” “normally”), modality (“might” “likely”), and numeric quantifiers (“certain” “some”) in describing data practices [47]. These vague terms are companies’ precautionary attempts to inoculate themselves against lawsuits, but introduce privacy risks to consumers as the vague statements may conceal privacy-threatening practices [47].

**Non-comparability.** Privacy notices and choices differ significantly across industries and organizations, presenting a daunting learning curve for consumers [29]. The requirement of subheadings and standard formats stipulated in the US Gramm-Leach-Bliley Act has substantially improved the readability of financial institutions’ privacy

policies [272]. Since then, there have been numerous efforts to pursue standardized formats of privacy policies [100, 263, 417]. Empirical studies have shown that consumers comprehend standardized formats faster but at the cost of accuracy as important details are hidden in layered notices [333].

**Inflexibility.** Privacy notices and choices struggle to keep up with evolving business models and consumer demands. The highly-publicized Facebook–Cambridge Analytica scandal reflects the increasing difficulty for companies to track information flows while ensuring constant alignment between stated and actual data practices [443]. Some companies use an “umbrella privacy policy” to cover several subsidiaries when the policy itself is not universally applicable [198]. Moreover, it becomes a tricky endeavor to choose the appropriate channel for privacy notices when it comes to emerging technologies such as smart home devices [437]. For instance, most fitness tracking device manufacturers still put privacy policies on their websites, which consumers rarely read, but the actual device does not offer any form of privacy notice [395].

**Dark patterns.** Dark patterns try to manipulate consumers into making decisions against their best interests through malicious interaction flows [186, 356]. Academic researchers [13, 50, 186, 321, 356, 368], consumer protection organizations [162], white papers [143], and popular press [462] have documented the prevalence of dark patterns and provided examples: having privacy-intrusive default settings, requiring lengthy efforts to select privacy-friendly choices, and illusory or take-it-or-leave-it choices. While regulators have listed common dark patterns as examples of non-compliance [144, 377], these efforts do not seem effective. For example, after the GDPR went into effect, consent banners are still fraught with dark patterns [323, 368, 539], as illustrated by highlighted “I agree” buttons and the absence of granular choices.

Prior research has contributed frameworks for best practices of privacy interface

design. Schaub et al. discussed that an ideal privacy notice should appear at different times tailored to the user’s needs in that context, be delivered across multiple channels, target the most effective and accessible modality, and come with usable control options [437]. Feng et al. developed a similar framework for privacy controls among IoT devices, introducing type (what kinds of choices are offered) and functionality (capabilities to support the choice process) as two additional dimensions to consider [150]. The authors also noted that privacy choices should be integrated with privacy notices and ideally be presented through data-driven privacy assistants that further reduce consumers’ burden and honor their privacy preferences [150].

### 2.3.2 Privacy Nudges

Thaler and Sunstein defined nudges as soft paternalistic interventions that “predictably alter people’s behaviors without forbidding any option or significantly change their economic incentives” [490]. Below I discuss examples of privacy nudges along the six dimensions summarized by Acquisti et al. [4] that target different aspects of decision-making: information, presentation, defaults, incentives, reversibility, and timing.

**Nudging with information.** Nudges can provide information to create awareness of privacy risks, especially considering that consumers are bounded by information asymmetries and do not have a complete picture of companies’ data practices [7]. Information can be provided through education (informing users of benefits and risks associated with a system or application before the usage) or feedback (information provided alongside or after the use) [4]. Ideally, information-based nudges should contain concise and relevant information about privacy risks and come with other nudges that mitigate the burden of reading and comprehending such information.

**Nudging with presentation.** How information and choices are presented can address many hurdles and biases consumers experience in privacy decision-making. Nudging toward privacy-friendly options can occur through framing risks as a clear loss (to address loss aversion), framing risks in exaggerating manners (to address overconfidence), or differentiating the representativeness of a risky event from its probability of occurrence [4]. In addition to framing, which mainly concerns text changes, privacy nudges can adjust the order of presenting different options and target post-completion errors by placing the most important action on top [600]. Changes in framing and ordering can result in varying saliency among different privacy options, and saliency can be adjusted through visual cues.

**Nudging with defaults.** Privacy nudges can target the status quo bias, i.e., people’s tendency to stick with default choices. An idea from the Privacy by Design framework [73] is that privacy settings should protect an average user’s information by default while allowing expert users to customize settings toward their specific needs. There are mass defaults (in which everyone gets the same default) vs. personalized defaults (catering to the needs of individual users, which also come at the cost of collecting additional information). However, defaults can be too sticky or slippery [582].

**Nudging with incentives.** Incentives can take the form of rewards or punishments. Providing desired incentives can help mitigate hyperbolic discounting, i.e., struggles in recognizing long-term negative consequences [4]. In privacy-related contexts, rewards can be monetary (e.g., selling personal information for a given price) or non-monetary (e.g., an honorary badge for users who regularly check in on location-sharing apps [301]). Punishment can mean the time and effort required to configure privacy settings, and in more extreme cases, termination of service when rejecting privacy policies [162]. Based on the loss aversion theory, incentive-based nudges can



frame a punishment as a loss to make it sound scarier than framing it as a simple penalty [4].

**Nudging with error reversibility.** An ideal system should ease the pain of correcting errors or do as much as possible to prevent errors from happening. Examples include an “undo” button that enables the removal of regrettable emails and a short period of delay after clicking the “post” button on social media. While some consumers may find such a delay nudge beneficial as it helps them stop and think, others may find it annoying or unnecessary [560].

**Nudging with timing.** Timing options for delivering privacy notices range from periodic, just-in-time to on-demand [437]. Prior research has documented the importance of timing choices, e.g., as online shoppers become willing to pay significantly more for increased privacy when seeing privacy indicators before visiting websites than after arriving at the site [136]. The just-in-time approach stands out among all options as it provides relevant and contextual information while minimizing disruptions to the user’s workflow. As evidence, prior work has found that privacy notices shown while using a mobile app receive more attention than those appearing in the app store [37].

**Limitations of nudging.** Nudges are not the silver bullet to solve every privacy problem. By focusing exclusively on the decision-making aspect of individuals, nudges may fail to address the causes of or provide the solution to many other problems [339]. Techniques that nudge individuals to opt out of the sale of personal information may not be as effective as regulations that fundamentally ban such selling between businesses. Moreover, consumers’ responses to nudges may differ across individual, social, economic, and cultural contexts [55, 392]. Efforts to create a one-size-fits-all nudge may fail to match the targeted group, context, or behavior, thereby triggering

unintended consequences worse than the original behavior [419]. Researchers should also be attuned to the ethical aspects of deploying nudges: nudges should result in equal benefits to all recipients and be supported by solid scientific reasoning for predictable impacts [419]. Nudges should not be deployed when they misalign with the privacy needs of recipients, when unintended consequences may override potential benefits, or when they present unresolvable conflicts of interests between multiple stakeholders [4].

### **2.3.3 Privacy Regulations**

As of 2021, 145 countries and independent territories, including nearly every country in Europe and many in Latin America and the Caribbean, Asia, and Africa, have enacted data privacy laws [189]. In addition, at least 23 other countries have official Bills for new laws in various stages of progress [189]. The General Data Protection Regulation (GDPR) [144] enacted by the Council of the European Union is one of the most stringent data protection regimes with extending impacts across the globe. GDPR requires that data controllers and processors implement procedures that make privacy by design and respect the rights of data subjects, ranging from the right of access to the right to be forgotten [144]. GDPR's influence continues as other non-EU countries start to enact new laws that are comparable or stronger [189].

In the United States, there has not been an all-encompassing federal privacy law. Partial regulations exist concerning certain states, sectors, or populations, such as the Health Insurance Portability and Accountability Act (HIPAA) which applies to health information, and the Children's Online Privacy Protection Act (COPPA) which applies to children under 14 years. The California Consumer Privacy Act (CCPA), effective in 2020, has been one of the strongest US privacy laws with an array of new consumer rights. Under CCPA, consumers can opt out of the sale of personal information and access/transfer personal information to third parties in a "readily

usable format” [373]. While CCPA ostensibly applies to Californian consumers, California’s population and economic importance mean that CCPA will likely generate impacts beyond California’s borders and serve as a testing ground for a long-awaited US federal privacy law [205].

While privacy legislative activities are growing worldwide, more stringent data protection laws do not guarantee better privacy outcomes. Most privacy regulations nowadays still take the “privacy self-management” model that burdens individual consumers to make their own decisions [473]. The privacy self-management model fails to result in meaningful controls [100] and does not address structural problems like market forces behind the expanding encroachment of surveillance into consumers’ lives [7]. Moreover, privacy regulations may generate unintended consequences by deepening the gap between large and small businesses, being cost-prohibitive for compliance, and giving too much power to the government [285]. As privacy advocates call for more substantive and comprehensive policies, there should be more research on the impacts of privacy regulations beyond economic metrics [6]. As Ari Waldman wrote in “Industry Unbound,” we need privacy laws that enable unequivocal bans of harmful technologies rather than vaguely written clauses that allow businesses to stay inside the law without fundamentally improving their data practices [555].

## CHAPTER III

# Consumer Reactions to the 2017 Equifax Data Breach: An Interview Study<sup>1</sup>

Data breaches are security incidents with grave privacy implications for affected consumers. This chapter presents a qualitative study featuring consumer reactions to the 2017 Equifax Data Breach. Equifax is one of the “Big Three” credit bureaus (also called credit reporting agencies) in the US; the other two are Experian and TransUnion. These organizations create aggregated reports of individual consumers’ credit information and offer this information to businesses that need to assess their customers’ creditworthiness, such as lenders, landlords, and employers. In 2017, Equifax suffered a data breach in which hackers stole the sensitive data of over 146 million consumers [251], which were almost half of the US population. The compromised data included names, addresses, birth dates, social security numbers (SSN), and driver’s license numbers. The combination of these data types—many of which are personally identifiable and hard to change—poses significant risks of identity theft. Nonetheless, a survey conducted after the breach found that only 8% of respondents reported having frozen their credit reports [371].

---

<sup>1</sup>This chapter is based on: Yixin Zou, Abraham Mhaidli, Austin McCall, & Florian Schaub. 2018. “I’ve Got Nothing to Lose”: Consumers’ Risk Perceptions and Protective Actions after the Equifax Data Breach. In *Fourteenth Symposium on Usable Privacy and Security (SOUPS 2018)* (pp. 197–216). [602]

Our study shed light on the reasons behind consumers’ lack of action after the Equifax data breach. While most participants were aware of the breach and resulting risks, few knew whether they were affected, and even fewer took proactive protective measures. We find that this inaction was not primarily influenced by the accuracy of mental models or risk awareness but rather by costs associated with protective measures, optimism bias in estimating one’s likelihood of victimization, and a general tendency towards delaying action until harm has occurred. In addition, sources of advice influenced actions taken; many participants who took action acted on advice from people they trusted. Yet, taken actions also created a false sense of security for some participants, leading them to overlook other measures. Our findings suggest technical, legal, and educational approaches to protect consumers’ credit data better and empower consumers with usable protection measures.

### 3.1 Background

**Consequences of data breaches.** Data breaches have multifaceted consequences. Breached organizations can bear substantial costs to repair the aftermath, including patching system vulnerabilities, compensating affected individuals, and resolving potential lawsuits [424]. There are also invisible and hard-to-measure costs in rebuilding the breached organization’s reputation [268, 579] and affected individuals’ trust [1]. For affected individuals, exposed data puts them at risk of account compromise [449, 514], phishing [393], and identity theft [426, 475]. Though it may take years before leaked data is misused, the harm can be profound when it happens. For instance, victims of identity theft may have ruined credit reports or have to file for bankruptcy due to abuse of credit [26]. Identity theft is also traumatizing: in a 2017 survey by the Identity Theft Resource Center [235], 77% of respondents reported increased stress levels, and 55% reported increased fatigue or decreased energy. Thus, some researchers have argued that data breaches cause compensable harm due to

the substantial risk of future financial injury and the emotional distress imposed on victims [113, 475].

**Available protective measures.** Breached organizations are often legally required to notify affected victims [144, 397] and offer compensations such as discounts [482] and free credit or identity monitoring services [446]. Services like HIBP [232] and Firefox Monitor [354] routinely examine third-party breach reports and notify signed-up users. Some companies automatically reset passwords for users whose credentials appeared in password dumps [183, 587]. Additional measures for victims include two-factor authentication (2FA) that increases the difficulty of misusing leaked credentials and warnings that flag social engineering and phishing attacks [281, 396].

In the US, most consumers have credit reports that curate information about their credit activity and current credit situation [491]. As such, breach victims are further urged to check their credit reports regularly and place a credit freeze or fraud alert on credit reports if personally identifiable information like SSN is exposed [506]. Credit freezes block inquiries for one's credit report, thus preventing new accounts or loans that require credit checks to be opened under the consumer's name [90]. A credit report with fraud alerts signals that the consumer is at risk of credit fraud. Creditors are expected to perform identity verification, but sometimes they may ignore such alerts and take no action, making fraud alerts less restrictive than credit freezes [90].

Nevertheless, no solution is perfect. Attackers can bypass 2FA without obtaining the secondary token [126], and phishing warnings have low adherence rates [16, 133]. A credit freeze only limits access to a credit report and thus does not protect against other types of identity theft that do not require credit checks, such as tax fraud. Placing a credit freeze is a burdensome process, with a long waiting time and frequent error messages [303].

**Consumer reactions.** Past research and news media have shown that consumers rarely take recommended protective measures in response to data breaches [1, 187, 403, 471] even with increased concern about identity theft [28, 403] and diminished trust in the breached organization [359]. In a 2014 survey by Ponemon Institute [403], the most frequent response to a breach notification was to “ignore it and do nothing” (32%), followed by the acceptance of free identity theft protection services (29%) and contacting the breached organization for more information (21%). RAND Corporation’s US national survey in 2016 [1] reports a higher acceptance of free credit monitoring (62%), but only 11% of respondents stopped conducting business with the breached organization. Motivated by the apparent contradiction between the prevalence of data breaches and consumers’ lack of action, our research aims to gain insights into why the contradiction occurs by examining potential factors, including awareness, risk perception, and hurdles in following recommended protective measures.

## 3.2 Method

Our study seeks to answer the following research questions:

**RQ1: What are consumers’ mental models of how credit bureaus operate?**

**RQ2: What do consumers perceive as risks of the Equifax data breach?**

**RQ3: What protective actions did consumers take or not take in reaction to the perceived risks, and what drove their decisions?**

We conducted semi-structured interviews with 24 participants in January and February 2018. All interviews were audio-recorded and lasted 40 minutes on average,

ranging from 20 minutes to 61 minutes. Each participant was compensated with \$10. The study was determined to be exempt by the University of Michigan’s Institutional Review Boards (IRB).

**Interview procedure.** We developed and refined our script for the semi-structured interviews through multiple pilot sessions. The final interview script is included in Appendix A.1. We started by asking participants how they manage their personal finances, leading to a discussion about their experiences with and understanding of credit bureaus. Next, we asked about their awareness of Equifax and the 2017 data breach before providing a basic description for those who had not heard of it. We probed participants’ risk perception by asking what they saw as consequences of the breach, reactions when hearing about the breach, and feelings about their data collected by Equifax. Then we asked whether participants had taken any protective action and, if yes, their experiences and interpretations of an action’s outcomes. Finally, we asked participants to recall previous experiences with data breaches and identity theft before giving participants opportunities to ask clarification questions.

At the end of the interview, participants were asked to complete two questionnaires measuring their financial decision-making ability [362] and self-determined financial well-being [92]. We collected financial-related information after the interview to minimize potential priming. For instance, participants might otherwise think the study is about one’s financial management and overstate how often they check credit reports. Conversely, the interview questions should have little impact on participants’ responses to the exit survey, as they did not touch on the same topics.

**Recruitment.** We recruited participants via online platforms (e.g., Reddit, Craigslist, and Facebook) and emails to a university research pool and campus mailing lists. We recruited for “a study on personal finance and credit bureaus,” purposefully not mentioning Equifax to avoid priming participants and limit self-selection bias. Prospective



ID	Gender	Age	Education	Income	NFEC (0-8)	CFPB (0-100)
P1	W	60-69	Bachelor	\$125-150k	8	88
P2	M	30-39	Master	\$25-50k	6	61
P3	M	60-69	Bachelor	<\$25k	5	35
P4	M	18-29	Some college	\$125-150k	7	73
P5	W	50-59	Master	<\$25k	3	41
P6	M	50-59	Bachelor	\$50-75k	6	45
P7	W	18-29	Bachelor	\$25-50k	4	50
P8	W	50-59	Some college	<\$25k	6	47
P9	M	60-69	Bachelor	<\$25k	8	48
P10	W	18-29	Some college	\$150k+	7	81
P11	W	18-29	Bachelor	NA	8	54
P12	M	40-49	Master	\$50-75k	7	65
P13	W	30-39	Professional degree	\$50-75k	5	58
P14	W	18-29	Some college	<25k	5	56
P15	M	40-49	Bachelor	<25k	8	49
P16	M	50-59	Master	\$75-100k	7	57
P17	W	30-39	Master	\$150k+	6	75
P18	M	30-39	Bachelor	\$25-50k	6	57
P19	W	50-59	Master	\$100-125k	7	56
P20	W	18-29	Master	\$50-75k	7	64
P21	M	50-59	Some college	\$125-150k	8	82
P22	M	18-29	Bachelor	\$25-50k	6	52
P23	W	40-49	Master	\$75-100k	8	60
P24	W	40-49	Associate	\$50-75k	7	56

Table 3.1: Demographics of participants, and scores of NFEC financial decision [362] and CFPB financial well-being scales [92].

participants provided basic demographic information in an online screening survey (see Appendix A.2). We only recruited US citizens and permanent residents who had lived in the country for more than five years, as recent immigrants might not be familiar with the US credit reporting system or may not be included in credit bureaus’ databases yet. Because prior literature suggests that demographic factors can influence people’s financial experiences and responsibilities [307], we selected a diverse sample of 24 participants in terms of age, gender, education, occupation, and income.

Table 3.1 summarizes the demographics of our participants. 11 participants identified as men, and 13 as women. Their ages ranged from 21 to 68, with a median age of 44. Five participants had no college experience, ten had a Bachelor’s or Associate’s degree, and nine had a graduate degree (e.g., Master’s or Professional degree). Eight

participants worked in a university setting as students or staff, and the rest worked in various medical, business, and social work industries. P16 was the only participant with a cybersecurity background. Our participants' annual household income ranged from less than \$25k to more than \$15k, with the median income in the range of \$50-75k. For the NFEC financial decision test [362], 19 of our 24 participants got a score of 6 or higher, indicating they are financially literate enough to make entry-level financial decisions [362]. The CFPB financial well-being score [92] ranged from 35 to 88 with a median score of 56.5 out of 100, suggesting average financial well-being in our sample [92].

**Qualitative data analysis.** With participants' permission, we audio-recorded and then transcribed all interviews. We then conducted thematic analysis [346], a standard analysis approach in qualitative studies on human-computer interaction [286] and usable privacy and security [163, 592]. Two researchers co-developed the initial version of the codebook by coding a subset of interviews independently and grouping them into themes. They achieved good inter-coder reliability through multiple rounds of collaborative refinement (Cohen's  $\kappa=0.79$ ) [156]. The final codebook included 14 overarching themes (e.g., "understanding of credit bureaus," "attitudes toward the breach," and "actions suggested by participants") and 53 unique codes (see Appendix A.3). One researcher then coded the remaining interviews and re-coded previous ones using the final codebook.

As our study was primarily qualitative, we do not report exact numbers when presenting most of our study findings. However, following recent qualitative work in usable security and privacy [139, 197], we adopted the terminology in Figure 3.1 to provide a relative sense of the frequency of major themes.

**Limitations** Our study has certain limitations. First, as is typical for qualitative research [286], our sample size cannot support quantitative conclusions about the




Figure 3.1: Terminology used to present relative frequency of themes.

general US population. That being said, we believe our study provides rich qualitative insights into people’s risk perception and hurdles in taking action after the Equifax data breach. We recruited a demographically diverse sample to gather a wide range of issues, perceptions, and perspectives. Findings like optimism bias and a reactive approach to dealing with risk are unlikely to be limited to specific demographics. Second, we conducted our study four months after the Equifax breach was made public. The lapse between the breach’s occurrence and the interview might lead to a decrease in breach awareness. We chose the timing deliberately to ensure participants had sufficient time to take any action if they wanted. Although most participants could not remember details of the breach, they seemed to have a clear memory of actions they took and did not take and could articulate the reasons behind it. Third, interviews are based on self-reported data, and participants may over-claim security and privacy concerns due to social desirability bias. To mitigate this bias, we designed our interview script to avoid priming questions and give participants space to re-collect their memories before prompting them about specific measures.

### 3.3 Findings

We discuss our findings in three areas: mental models of credit bureaus, risk perceptions of the Equifax breach, and protective actions.

### 3.3.1 Mental Models of Credit Bureaus

Almost all participants were aware of the big three credit bureaus and most correctly interpreted the bureau's function as assigning credit scores to individual consumers. Yet, no participant could fully describe the types of information collected by credit bureaus and corresponding information exchange entities, making their mental models incomplete or inaccurate.

**General awareness of the big three bureaus.** While almost all participants knew that there are three big credit bureaus in the US, only some could list the specific names of all three. A few participants mentioned that other smaller-scale credit bureaus might also exist. As P11 said, *"I wouldn't be surprised if there are other smaller companies that track and monitor credit scores."* A few participants had difficulty mapping the names they had heard of with the concept of credit bureaus. P15 shared: *"I don't know if the credit bureau is separate, or if Equifax, Experian, et al. are considered credit bureaus."* P3 considered Credit Karma, a company offering credit monitoring services, as a credit bureau based on how he checked credit scores via Credit Karma: *"It is on the same level as those three major ones . . . You can go to them any day . . . and they're not charging, but they have that same information."*

**Purpose of credit bureaus.** Most participants described credit bureaus as companies that assign credit scores to individual consumers. They understood that these scores represent one's creditworthiness and help lenders, insurance companies, and others make decisions. By contrast, some participants gave inaccurate descriptions of credit bureaus. P11 viewed credit bureaus as government agencies, whereas these are private, for-profit organizations in reality. Some participants confused credit bureaus with other organizations such as credit unions, debt collectors, and loan companies. P23, for instance, confused credit bureaus with credit rating agencies that rate the

creditworthiness of companies and governments rather than individual consumers: *“I guess they need to support the rating . . . and maybe the credibility of that organization. Maybe any complaint from the customer [about] how they use their funding, and if it’s a bank, how they use the customer’s money.”* P4 referred to credit bureaus as loan companies: *“They loan out money to [your] credit cards that they expect you to pay back. Then if you don’t pay back, they just charge you more interest.”*

**Incomplete understanding of information flows.** Regarding the information collected by credit bureaus, personally identifiable information such as names, addresses, and SSN and financial-related information (e.g., number of credit cards and loans, credit limits, late payments) were noted most frequently. About half of the participants mentioned the collection of employment history, public records (e.g., tax lien and bankruptcy), and inquiries made by creditors in recent years. Some stated that the information collected by credit bureaus is “a lot,” “a variety of different things,” or “almost everything,” yet no participant covered all types of collected data. Participants’ knowledge was closely tied with their personal experience with credit bureaus. Those who checked their credit reports more frequently could recall more details but still showed uncertainty. As P24 described, *“I think they use past accounts and maybe employment history. I know they use the length of credit. But like I said, I don’t know, random guessing.”*

Some participants thought credit bureaus collected certain data that credit bureaus do not collect. For instance, P9 thought credit bureaus would monitor the consumer’s social media accounts, indicating that he might confuse credit bureaus with data brokers: *“[It] would just show things like their hobbies and travel like to go to Europe or Las Vegas . . . It would give you an idea of their lifestyle, and if they’re throwing money around.”*

Almost all participants noted that financial institutions are the primary informa-

tion providers for credit bureaus. As P13 described, “*people who provide information are like banks, loan companies, loan providers, debt collectors . . . people who you’ve rented with before and haven’t paid back . . . stores or credit card companies.*” Some participants mentioned auto dealers, governments, and utility companies as information providers, although these were brought up much less frequently. As for customers of credit bureaus, participants mentioned creditors and lenders more regularly than other businesses, such as car dealerships and landlords. Some participants noted that information providers of credit bureaus are simultaneously their customers, and collaborations exist between these organizations. According to P16: “*I imagine that they also send some of that information back to banks . . . It’s a two-way street, I assume, and there’s probably a data sharing agreement between them.*”

**Interactions with credit bureaus.** Most participants knew their right to obtain a free credit report annually. A few participants mentioned other products offered by credit bureaus that might cost money, such as credit monitoring services. Most participants noted that although they knew they could check credit scores directly at credit bureaus, they preferred to check their scores through third-party financial management tools such as Credit Karma due to their low cost, convenience, and frequent updates. Notably, low-income participants generally knew about the offerings of credit bureaus but chose not to take advantage of them or refused to interact with credit bureaus altogether. P5 and P15 said they had no interest in checking their credit reports. According to P15: “*I can find out my credit score . . . But I have been reluctant to do that because (a) I know my credit’s terrible, and (b) I don’t want to give them any information.*”

**Negative sentiments.** Some participants expressed a moderately or strongly negative sentiment towards credit bureaus or the whole credit reporting system. Sometimes, the negative perception stemmed from doubts about the credit reporting sys-

tem's fairness to consumers. P19 explained: *"I don't like the idea that [for] things like auto insurance and getting an apartment ... people come up with cash upfront and they still get denied because of a credit report ... It does make sense that there is something like this, but not the way it's running right now."* P14 described how credit bureaus exacerbate inequality by worsening the financial well-being of already less affluent people: *"It's really like a bad cycle. If you don't have enough money and then you need a loan ... you can't get a loan or your interest rate is really high, and you can't afford to pay it."* P24 said that credit bureaus work to serve lenders' interest, with little concerns about individual consumers: *"For the interest of who? Those in power to make these laws ... I'm assuming they probably all have lobbyists and things that could potentially benefit collaborators of credit bureaus, like lenders, businesses and car companies."*

Other negative perceptions originated from unpleasant personal experiences with credit bureaus. P1 said that her husband was once denied a credit card because credit bureaus provided the credit file of another person with the same name to the credit card company. P5 described going through a burdensome process of disputing erroneous credit card charges with little support from credit bureaus, leading to a loss of faith in the system: *"[The dispute process] is probably all automated and they only take what people give them ... I feel powerless to try to get that stuff off. I just give up. I don't care. That's why I say I don't want to look [up my credit report]. Because how much stress and time would that take?"*

Moreover, a few participants expressed confusion or concern over the data exchange between credit bureaus and their information providers or customers. P3 expressed his frustration when finding out information about transactions between him and other businesses would inevitably be obtained by credit bureaus, which constituted a breach of confidentiality: *"When it comes to credit bureaus, I don't think there is any such thing as confidentiality ... Whatever I'm talking to [banks], whatever*

*they do, that should be strictly between them and me. Okay? But somehow, in my mind, the credit bureau ends up with this information.”*

### **3.3.2 Risk Perception of the Equifax Data Breach**

About half of our participants had heard of the Equifax data breach before the interview. They viewed identity theft as the primary risk and described how it could happen. A few participants also noted privacy invasion as a secondary risk. Nevertheless, participants seldom associated these risks with themselves, implying the existence of optimism bias.

**Vague memory of the event.** Participants showed a high awareness of the breach’s occurrence. Almost all participants knew that a data breach happened to one of the big three bureaus. Most knew that Equifax suffered the breach; the rest either did not remember the name or attributed the breach to Experian. Participants generally had a vague idea that the company was “hacked,” leading to the disclosure of “a lot of” information, but many stated that they could not remember the details. For example, P2 said: *“I don’t know the specifics, [like] if it was a hacker attack or something like that, but I know that a lot of information got out and millions of people were affected.”* As for types of information that were exposed, name, address, date of birth, and SSN were mentioned most frequently, followed by bank account numbers and credit card numbers.

**Identity theft as a primary risk.** Almost all participants mentioned the risk of identity theft as a direct consequence of the Equifax breach. Some participants explained how identity theft could happen in their opinions. As P2 shared, *“The consequences? ...It could make it very easy if somebody wants to steal somebody’s identity. [Identity thieves] could get hold of ...the name, SSN, and birth date, and could just open up a bunch of accounts under their name.”* Most of the examples



participants gave focused on opening new accounts and fraudulent charges on existing accounts. Only a few participants considered potential misuse of stolen personal information that did not require credit checks, such as tax fraud. P12 mentioned that this breach prompted him to consider filing his tax return earlier this year: *“It could lead to some fraud around tax time. I heard the other day where . . . criminals take other people’s tax returns. I’m going to file my tax returns as soon as I can.”*

Participants’ knowledge of exposed data influenced their risk perception about identity theft. The loss of SSNs triggered more identity theft concerns than other types of personally identifiable information. P13 differentiated the sensitivity of exposed information based on how publicly accessible it was: *“you can find someone’s date of birth and name online, but the Social Security Number should be harder to find.”* P19 was concerned as Social Security numbers are hard to replace: *“You can’t get a new Social Security Number . . . the government is not very accommodating about that . . . I would prefer not to think about it because you’ll just be screwed.”* Both P13 and P19 mentioned that the combination of different types of data scared them the most. As P19 said, *“It’s not like someone will take a credit card out in your name or somehow try and use your bank, and you have some recourse . . . If they’ve got everything, I have no idea what you would do.”*

**Privacy invasion as a secondary risk.** A few participants stated that the exposure of such sensitive data is an invasion of privacy. Although P5 did not use the word ‘privacy’ explicitly, she described the panic when thinking about how much the hackers could know about her: *“[The hackers] would find out my personal information, which really scares me. I don’t want people to know where I live. I don’t want people to know whatever information they have.”* P16 noted the possibility of knowing one’s personal life in detail based on the exposed data: *“As they aggregate that data, they can get more and more information about you. For example, if there’s detailed*

*credit card information, which I hope not, they would know your shopping habits, they might know where you live, what kinds of cars you drive.” P22 said he would view his financial information as private information, but he would not value it as highly as SSN due to the latter’s repercussions for identity theft: “I guess I would value my Social Security number, number one, because I don’t want my identity stolen. I also value my privacy, but I feel like I haven’t gotten to a point yet where I’ve made lots of . . . credit-based purchases, so not yet at a point where that’s my number one.”*

**Eroded trust.** Some participants noted that this breach eroded their trust in Equifax’s ability to ensure the security of consumers’ data. P14 said that consumers had no choice but to trust Equifax: *“They’re gonna get your information whether you wanted them to or not.”* P12 shared that his trust in Equifax decreased to the point that he did not accept the company’s free credit monitoring service: *“Well, you didn’t handle the other information, why should I trust you to monitor information anymore?”* Interestingly, P24 provided a counterexample that she would trust Equifax more because Equifax would now have better security practices: *“I’d probably go back to them just because they’re probably going to be a little bit more cautious than the one that didn’t get hit.”*

**Perceived low likelihood of being personally affected.** While almost all participants demonstrated an understanding of the breach’s resulting risks, most did not assume they would be personally affected. Considering that the Equifax breach affected almost half of the US population, such perception could be optimism bias [448]. We identified multiple reasons for the underestimation of personal risk.

A few participants mentioned checking the Equifax website to see if they were affected and received the message “your personal information was not impacted by this incident.” Additionally, there was a notion of ‘I have nothing to lose,’ especially among low-income participants. P5 said: *“I don’t have any credit. I have a bad*

record, so I wouldn't [check if I was affected]. Nobody can hurt me; it's already at the lowest place." Some participants mentioned the absence of signals indicating negative repercussions: Equifax did not send notifications to individual consumers, and they did not notice suspicious account activities since the breach occurred. P7 said: "[Equifax] was like ... there was a breach, and if you were directly affected, we would let you know. [Interviewer: But then you never received it?] No, so I was fine." Furthermore, some participants presumed they were not included in Equifax's database or had limited information in the database. For instance, P6 believed he could not be affected because he had never held any credit cards. P8, who held a credit card but never checked her credit reports, believed the information shared with credit bureaus was not as extensive as someone who regularly checks their credit reports or interacts with credit bureaus in other ways.

Even for participants who thought they might be affected by the breach, none claimed it assertively. Among a few participants who received the 'Your personal information might have been impacted by this incident' message from Equifax's website, most were doubtful about the word 'might.' P13 interpreted it as a public relation strategy that did not necessarily reflect the truth, causing little concern to her: "If they say no and then you get affected, you might be like 'you said I wasn't gonna be affected so I didn't worry, and I wasn't monitoring' ... But if they say yes, then, of course, you're gonna freak out and start calling them, asking them for advice or services ... If they say maybe, that's like a safe, middle ground for a company to say." P2 did not check the website but felt he could be affected since "[if] these many people were affected, it's likely that I was affected."

### 3.3.3 Negligence of Protective Actions

Figure 3.2 lists actions suggested by FTC for reacting to the Equifax data breach [501]. Despite perceived risks, most participants did not actively take any protective mea-

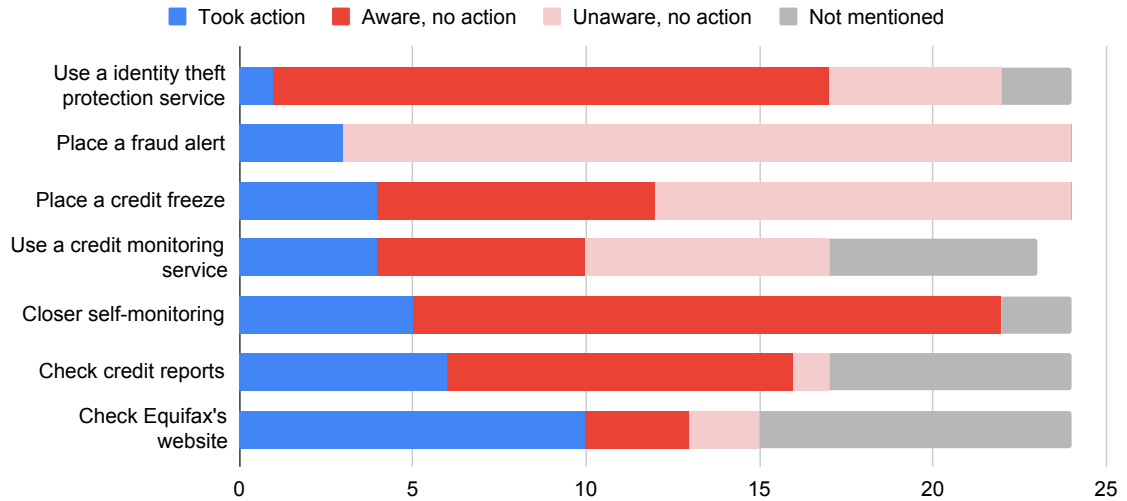


Figure 3.2: Status of FTC's suggested actions [501] taken or not taken by participants.

asures after the breach. Participants were either unaware of available measures or intentionally avoided using them for various reasons.

**Insufficient knowledge.** The high portion of participants unaware of available protective measures suggests insufficient knowledge as a primary reason for inaction. Only a few participants correctly described fraud alerts, and all learned it because of being affected by previous data breaches and receiving the service as compensation. The remaining participants either said they did not know what fraud alerts were or associated them with alerts sent from banks and credit card companies when fraudulent activities occurred. Similarly, about half of the participants interpreted credit freezes as freezing credit cards. Participants generally considered the measures banks and credit card companies offered to prevent identity theft. However, the unawareness of fraud alerts and credit freezes provided by credit bureaus prevents them from utilizing these measures to protect their credit reports.

**Financial cost inhibited action.** Financial cost appears to be a significant determining factor of whether an action was adopted. Actions with no cost were more

favored: checking Equifax’s website was taken by most participants, followed by checking credit reports through `annualcreditreport.com` or a third-party service for free and a closer self-monitoring of existing accounts. For the remaining actions in Figure 3.2, most occurred before the Equifax breach, e.g., when the participant had been affected by another breach and received free services as compensation. Some participants questioned the effectiveness of identity theft protection services compared to the associated cost. P22 said: *“It feels like you’re giving them money for nothing . . . I don’t know if I believe them because they can’t own all of my data, so how are they actually protecting me?”*

For a few participants who had placed a credit freeze, only P19 paid to have her report frozen at all big three credit bureaus after the Equifax breach. P16 froze his credit reports at all three bureaus too, but he did it for free since he was a documented victim of another breach. P12 placed a credit freeze only at Equifax, which was offered for free, and P20 placed only a free credit freeze at ChexSystems, a smaller-scale credit bureau. Seven participants expressed that freezing and unfreezing credit reports should be free at Equifax and at other credit bureaus.<sup>2</sup>

A few participants viewed identity theft protection services as a waste of money. P3 said: *“I’m poor. I’m having a hard time keeping my head above water or staying. I’m not giving [credit bureaus] money for their profits.”* P8 considered these services an unwise investment: *“It wouldn’t be worth paying for something like that, but if I had a lot of assets, then I would pay . . . because I’d be more likely to lose money.”*

**Optimism bias.** A few participants attributed their reason for inaction to the perceived low likelihood of being personally affected, assuming that whoever had access to the stolen data would target more affluent people with a better credit history. P9 described himself as ‘a small fish in the pond:’ *“Why would they come*

---

<sup>2</sup>Since the study has been conducted, the Economic Growth, Regulatory Relief, and Consumer Protection Act (S.2155) has mandated that credit freeze and unfreeze actions should be free of charge [502].

*after me? If they're going to go to all the bother of stealing my identity, why don't they go after somebody with some real wealth?"* Notably, this attitude was not limited to low-income participants. P1, who reported being affluent, did not consider herself a target either: *"These days people can be so tricky about applying for things in your name ... especially people who had good credit. I would think they would definitely target them. Someone like my son who has a high FICO score ... they might make [him] a priority."*

**Tendency to delay action.** Another reason for inaction is a general retroactive way of dealing with risks. A few participants explained that nothing terrible had happened to them since the breach occurred and saw this as reassurance that no protective actions were needed. According to P10: *"I haven't had any problems with my credit since that happened ... that I heard about, so I'm not too concerned."* A few participants further cautioned that such a reactive approach might not be the most effective, but knowing this limitation was not enough to trigger action. As P9 described his general attitudes towards risks in life, *"Right now I don't have any problems, so I'm not really going to worry about it, and that's probably a very bad attitude, but I have enough problems in life without looking for trouble."* Similarly, P23 reflected that *"I wait until something bad happens and then I will react to it ... Maybe it's not as good as a proactive approach, [but] so far I think I'm okay with all the finance and nobody's stolen my identity yet."*

**Influence from the source of advice.** Of participants who took action, some said they were motivated or reminded to take action after receiving advice from various sources. News media were brought up most often, primarily informing participants about the breach and available options rather than prompting action. For instance, P9 reflected that he heard of the breach from NBC but did not follow their recommendations: *"I'm not sure which company it was, Equifax or which one, but I remember,*

*it's been a while . . . Consumers were supposed to take action and do something, but I didn't pay much attention to it because I didn't feel threatened.*" A few participants mentioned TV advertisements of LifeLock as the information source of identity theft protection services, but none of them had signed up for the service.

By contrast, a few participants who learned about available actions from sources they trusted (e.g., family members, colleagues, and experts) followed the given advice. For example, P19 described how she decided to place credit freezes after hearing a colleague's recommendations: *"He's our tech guy. He put together an e-mail . . . about how to find out and what to do, and so I finally did something."* P16 followed a security expert's advice to place credit freezes after being affected by a previous data breach: *"There's a gentleman named Brian Krebs who is active in the security community. He gave a very informative article about what's involved with credit freezes and why he chose to take that path to protect his credit. Given the way I use credit, and given [that] I have a very good credit rating, and given the data breaches . . . it made a lot of sense for me to do that."* P16 further mentioned that he shared Krebs' article with family members after knowing his son was affected by the Equifax breach.

**False sense of security.** A few participants mentioned that taking one action created a *"false sense of security"* in P16's words, which led to the negligence of other actions. P19, for example, shared that she did not continue to monitor her credit reports closely after freezing her credit reports: *"I downloaded the reports so I had a copy of it then, but I haven't done anything since then with regards to looking at it since I assumed that the freeze is working. I guess I am trusting the freeze, and I just don't want the hassle of having to worry about it all the time."* P16 mentioned that he checked his reports once a year instead of more frequently after being affected by a previous breach. He was also aware that a credit freeze could not fully eliminate the risk of identity theft: *"I think the credit freeze can help with some of it, but again*

*it depends on the institution that they're using . . . Let's say, for example, it was a car loan. If somebody was able to misrepresent themselves as me, they might be able to get the loan . . . Maybe there's an agency or [something] else besides the big three that is used to verify someone's credit information."*

**Usability issues.** Usability issues did not necessarily deter participants from taking action but still affected their experience. A few participants mentioned that using a PIN to lift the freeze was inconvenient. P8 described how a credit freeze created a lot of hassles for her elderly parents: *"One time I was with [my father], and he wanted to buy something. . . . He had to call the company, TransUnion, but then he couldn't remember his account . . . It just seemed like it was a lot of trouble . . . since you always have to know so many different passwords."* P20 initially tried to place a credit freeze at one of the big three bureaus but found it *"costs money and delays things,"* so she eventually chose a smaller bureau.

Participants also suggested making information flow around credit bureaus more transparent. For instance, P5 was not satisfied that consumers could only know a partial picture of what data credit bureaus collect about them: *"I think that I can learn what they know about me, but I don't have the power and the access to find out . . . It should be equal. Those reporting should be the same as those who have their name reported, but I'm skeptical."* P23 expressed confusion about different credit scores provided by different credit bureaus and argued that they should all be the same. Some participants pointed out that Equifax should have more proactive communications about the breach instead of assigning the responsibility of finding it out to individual consumers. As P21 said, *"They should have reached out to . . . their customers. Be very open about what exactly happened to the extent possible on a personal basis and communicate that to me personally."* P16 offered a general suggestion to the credit reporting system: *"The best way to regulate them is to define the boundaries around*



*privacy and data . . . to come up with means and standards to protect that information. On top of that, there should be means for consumers to work with those companies to have them respond to errors and misinformation, and to meet consumer needs.”*

### 3.4 Discussion

Our findings reveal that consumers’ adoption of protective measures was primarily influenced by financial cost, sources of advice, and an optimism bias of “the rich will be targeted” or “I’ve got nothing to lose.”

**Awareness does not lead to action.** Participants’ mental models of credit bureaus and risk awareness were not the primary factors affecting their protective behaviors. In line with prior work [54, 563, 592, 599], we found connections between certain components of participants’ mental models and their identity theft risk perception. For instance, participants who mentioned the possibility of tax fraud also specified government agencies as information providers of credit bureaus. However, most participants demonstrated awareness of identity theft risks regardless of how sophisticated their mental models were. Yet, most did not seem to turn such awareness into action. Participants’ decisions to take action or not seemed to be influenced by how they interpreted an action’s outcomes and their own biases rather than by a lack of knowledge of how credit bureaus operate.

**Costs as a barrier for action.** Prior work has shown how the US credit reporting system could exacerbate income inequality, as lower-income and minority groups are more likely to have inaccurate credit scores and face more challenges in getting loans and mortgages due to limited credit histories [27, 273]. Our participants expressed similar negative sentiments toward the credit reporting system and noted fees associated with protective measures as an additional barrier to action. For instance, some

participants perceived identity theft protection services as an untrustworthy and unwise investment. For participants who had placed a credit freeze, some did not do it at all big three bureaus due to cost considerations, and thus their credit freezes were not entirely effective. Participants who took action mainly chose economical options, such as checking their annual credit reports for free and keeping close monitoring of existing accounts.

**Advice as a trigger for action.** In line with prior work showing how advice sources influence security behavior adoption [411, 414], our findings provide additional insights on the effects of different sources in reacting to data breaches. Some participants learned about the breach and available protective measures from news media, but the awareness was insufficient to trigger action. Among those who took action, many instead followed recommendations from sources with perceived expertise and trustworthiness. A possible explanation is that participants received high-level information from news media but were more likely to resonate with the detailed and personal experiences of social connections. Future research can examine how different source characteristics (e.g., accessibility, quality, and credibility) affect advice adherence.

**Optimism bias and the “I’ve got nothing to lose” fallacy.** Our findings are consistent with prior work highlighting cognitive and behavioral biases in people’s security and privacy decision-making [4]. For example, optimism bias—the general tendency of underestimating the possibility of being affected by negative events [448]—occurred regardless of participants’ income levels: they tended to think ‘the rich’ were more likely to become targets of identity theft. Similar to Camp’s work [69], we found that participants tended to ignore action when there was no sign of negative consequences resulting from previous risky behavior, reinforcing the notion that protective measures are unnecessary.

Among low-income participants, the lack of motivation to take action is similar to the ‘I’ve got nothing to hide’ argument as a common misunderstanding of privacy [472], that there is no need to be concerned about privacy as long as one does not have secrets to hide. Similarly, some participants did not exhibit strong motivations to react because they thought they had nothing to lose, given their limited income or assets. Nevertheless, this notion contradicts findings in a 2013 Pew survey [290], which suggested that median-income households were more likely to be victims of identity theft than high-income households. In addition, existing US Department of Justice data has shown that at least one-third of identity theft victims live in low-income households [210], indicating that identity theft does not exclusively happen to affluent consumers. As Greene unpacked, thieves likely target low-income individuals because they have fewer resources to pursue a complaint [188]. Furthermore, as thieves use low-income stolen identities to receive public benefits, open utility accounts, or present to authorities when arrested, all of these activities can happen even when the stolen identity has low credit [188].

### 3.5 Recommendations

Our findings provide implications for public policy, technical, and educational approaches for protecting consumers against data breaches.

**Public policy recommendations.** Our findings demonstrate the need to amend the US Fair Credit Reporting Act (FCRA) to better protect consumers’ sensitive information held by credit bureaus. We argue that consumers should be able to obtain their detailed credit reports from the big three bureaus for free at any time, not just once per year. Even though participants reported checking their credit reports through third-party financial management services, many of these services only show the credit score from one bureau and a simplified report that might omit important

details. Other services may aggregate credit scores from different credit bureaus but do not explain the calculation clearly, leading to confusion among participants. Given these issues and the severe impacts of identity theft and erroneous credit reports, free and frequent access to credit reports is essential to lower the barriers for consumers to monitor their credit information.

When we conducted the study, a freeze or unfreeze operation of one’s credit report could cost up to \$10 per credit bureau depending on the consumer’s state of residence, and this action has to be performed at each bureau separately. In the published paper [602], we suggested that credit freezes, currently the most effective way of limiting undesired access to one’s credit data, should be free under any circumstance in all US states. The Economic Growth, Regulatory Relief, and Consumer Protection Act—a federal law effective since September 2018—has made this suggestion a reality [502]. The new law also allows parents to freeze for free the credit report of their children under 16 [502].

The magnitude of the 2017 Equifax breach further indicates a need for more stringent oversight of credit bureaus and better audits of their data practices. Our findings show that some participants held a negative sentiment towards credit bureaus due to inaccurate credit files, opaque data aggregation processes, and inadequate responses regarding the Equifax breach, among other reasons. In January 2020, Equifax agreed to a settlement with the US Federal Trade Commission (FTC), including up to \$425 million to help consumers affected by this breach [505]. In addition to such remedial enforcement, we argue that preemptive oversight measures, such as detecting and preventing misconduct through auditing, should take place simultaneously to ensure the security and accuracy of consumers’ credit data.

**Technical recommendations.** Accompanying public policy reform, there should be technical solutions to ensure that regulatory efforts result in improved and usable

protective measures for consumers. Our findings reveal issues with existing protective measures: participants experienced hassles when placing credit freezes or avoided certain measures due to low trust. Educational efforts about protective measures would not make sense unless the usability issues are addressed. We argue that credit freezes, for instance, should be offered in an integrated and user-friendly system. Similar to fraud alerts, credit freeze requests should be automatically communicated between the big three bureaus so that consumers do not need to go through the process with each bureau.

More efforts should also be devoted to making information flows around credit bureaus more transparent and giving consumers more agency in deciding how their information is collected and used. Our findings show that the opaque data collection and aggregation processes led to misconceptions. For example, some participants believed they were not included in credit bureaus' databases since they had no credit cards. But in reality, credit bureaus can still collect data about them from other providers, such as car dealerships and utility companies. Moreover, even though consumers can avoid using paid services from credit bureaus, they have no control over the information exchange between credit bureaus and their data providers.

An incremental yet promising step to involve consumers in the information flows is to develop just-in-time notifications that inform consumers whenever companies request a copy of their credit report, new data is added to their credit report, or a file about them has been created at any bureau. Once a consumer signs up for this service and passes the identity verification, they could receive notifications through multiple channels (e.g., mobile app, text message, and email). The notification should be quick to read and understandable. The notification could further come with an approval option when a third party makes a credit request, and consumers could allow or deny those requests. For instance, consumers could immediately raise a red flag when they notice inaccurate data being added to their credit report. This option not

only makes the dispute process more efficient but also may result in higher-quality credit data, benefiting credit bureaus and their customers.

**Educational efforts.** The implementation of regulatory and technical measures should be accompanied by educational efforts, as our participants showed a limited understanding of available measures and often misinterpreted their outcomes. While making resources widely accessible online is necessary, our findings suggest that provisioning resources alone is insufficient to guarantee the reach of most consumers. Our participants tended to act primarily on advice from trusted people rather than news or online resources. No participant mentioned the abundant online resources on identity theft protection, such as those on FTC’s website, illuminating the significant gap between how consumers obtain knowledge and how educational resources are offered. An interesting opportunity is to enlist financially literate or tech-savvy consumers as ‘influencers’ in educating other community members. Rather than creating ‘one-size-fits-all’ materials, it might be more effective to help people who are already motivated and well versed in the subject matter better communicate knowledge and recommendations to others.

## CHAPTER IV

# Consumer Reactions to Data Breaches: A Survey Study<sup>1</sup>

Chapter III provides qualitative insights into consumers' risk perception and behavior in the context of the 2017 Equifax data breach. Overall, the accuracy of consumers' mental model and risk awareness did little to explain the inaction; instead, factors such as insufficient knowledge of protective measures, optimism bias, and a tendency to delay action played a much more prominent role. An intuitive next step is to quantitatively understand how consumers react to breaches in general beyond focusing on individual breaches in isolation. In this chapter, I describe an online study ( $n=413$ ) in which we leveraged the Have I Been Pwned (HIBP) database to present participants with, and have them reflect on, specific data breaches that had exposed their email addresses and other personal information.

Methodologically, prior work primarily asked participants to recall past experiences with generic breaches [1, 403] or describe intended reactions in hypothetical scenarios [211, 260]. We apply a novel approach to examine participants' responses to specific breaches that exposed their information; such an approach increases ecological validity and mitigates recall bias, as participants learned about the breach and

---

<sup>1</sup>This chapter is based on: Peter Mayer\*, Yixin Zou\*, Florian Schaub, & Adam J. Aviv. 2021. "Now I'm a bit angry:" Individuals' Awareness, Perception, and Responses to Data Breaches that Affected Them. In *30th USENIX Security Symposium (USENIX Security '21)* (pp. 393–410). [329]  
\*The first two authors contributed equally to this research.

immediately shared their (intended) responses. Instead of focusing on one breach as a case study [187, 345, 522, 602], We gathered 792 detailed breach-specific responses (up to three per participant), covering 189 unique breaches and 66 different exposed data types. Our findings answer the following research questions:

**RQ1 [Breach status]: What factors influence the likelihood that an email address is involved in a data breach?**

Overall, 73% of our participants experienced at least one breach and 5.36 breaches on average. An email address's likelihood of being exposed in a breach significantly correlated with the email account's age and purpose of use.

**RQ2 [Perception]: What do participants perceive as the cause of being involved in data breaches and related impacts, and to what extent do their perceptions align with reality?**

Only 14% of our participants accurately attributed the cause of being affected by a breach to external factors, such as breached organizations and hackers. Others blamed their email or security behaviors for making themselves victims or viewed breaches as inevitable. Most participants expected little impact from shown breaches despite realizing certain risks.

**RQ3 [Awareness]: What factors influence participants' awareness of data breaches that affected them?**

Participants were unaware of most data breaches presented (74%). Those who knew they were affected by a specific breach had primarily learned about it from the breached organization or third-party services. Participants were more likely to be aware of older rather than recent breaches.

**RQ4 [Emotional response]: What are participants' emotional response to data breaches that affected them?**



Most participants rated their concern regarding breaches as low (56% slightly / somewhat concerned, 19% no concern). Certain breached data types, such as physical addresses and passwords, raised more concern than others. Participants expressed emotions ranging from upset, angry, annoyed, frustrated, and surprised (or not) to violated and fatigued.

**RQ5 [Behavioral response]: What factors influence participants' likelihood to take action in response to data breaches that affected them?**

Participants reported having already or being very likely to change their passwords and review credit reports/financial statements in response to over 50% of shown breaches. In addition, participants were more likely to take action with increased concern and prior awareness, suggesting that better communication about breaches could increase individuals' tendency to take action.

Our findings demonstrate that rather than burdening consumers to take action, breached organizations should be held responsible for increasing awareness and providing appropriate mitigations. Furthermore, our findings highlight the need for better tools to help affected individuals be more resilient against future breaches.

## **4.1 Method**

We conducted an online study with data pulled from HIBP by building a survey platform that queried the HIBP web service API using email addresses provided by study participants. To protect participants' confidentiality, we only maintained email addresses in ephemeral memory to query HIBP. At no point did we store participants' email addresses. We then used the query results, i.e., the breaches in which a participant's email address was exposed, to drive the remainder of the survey. The survey

## Your email address was part of the following breach

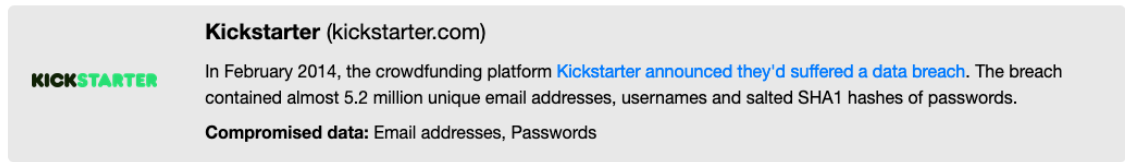


Figure 4.1: Sample breach information shown to participants.

consisted of three main parts (see Appendix B.1). The study was approved by the University of Michigan’s Institutional Review Boards (IRB).

**Survey procedure.** After consenting, we asked participants for their most commonly used email addresses. We clearly noted that the email address will only be used to query HIBP and that we will never see it. Once a participant entered an email address, we asked a few questions about it. Participants who indicated that the email address belonged to someone else or was fabricated could enter a different email address or leave the study. Next, we asked participants about their email habits as a potential influencing factor of the email’s involvement in breaches. The specific questions include how often they checked the email account, the primary use of the account (professional/personal correspondence or account creation), how long it has been used, and the number of other email accounts the participant used. We then used the provided email address to query HIBP.

We next informed participants whether their email address was exposed in any data breaches without stating the specific number or giving more details. To answer RQ2, we asked participants to speculate why their email address was or was not part of data breaches. Participants whose email address was not part of any breach could enter a different email address until a provided email address had associated breaches. If they did not provide another email, they were redirected to demographic questions toward the end.

We randomly selected up to three breaches, displayed one by one, to ask breach-related questions while limiting potential fatigue. First, we showed a breach’s description, logo, name, and types of compromised data provided by HIBP (see Figure 4.1). We explicitly stated that these were actual breaches, and no participants doubted the validity of shown breaches in their qualitative responses. For each breach, we asked about participants’ awareness (RQ3), emotional response (RQ4), and actions taken or intended to take (RQ5). For emotional response, participants provided open-ended responses, then rated their concern level on five points from “not at all concerned” (1) to “extremely concerned” (5) regarding the breach in general and for each type of exposed data. For behavioral response, participants described their reactions (open-ended) before rating their intention to take (or whether they had taken) ten provided actions based on online resources [352, 506]. The respective breach information was visible at the top of the page when participants answered all these questions.

We collected participants’ demographics, including age, gender, education, whether they had a background in IT or law, and household income. We also included two attention check questions: one asking them to identify the name of a breach shown during the study (only for participants whose email address was part of at least one breach) and a generic attention check. Finally, we showed participants a list of all breaches associated with their provided email address and links to resources on data breach recovery to help them process and act on this potentially new information.

**Recruitment.** We recruited participants via Prolific,<sup>2</sup> an online research platform similar to Amazon’s Mechanical Turk with more demographically diverse subjects [382], between August and October 2020. We balanced participants’ age and gender distributions in data collection. After the first 171 participants, we realized and corrected a storage error that caused missing data in income and ratings for taken/intended actions. Participants were compensated \$2.50 for an average completion time of 13.37

---

<sup>2</sup><https://prolific.co>

minutes (\$11.22/hour).

We collected data from 416 participants; three participants were excluded as they did not respond to any open-ended questions meaningfully, resulting in 413 participants. We based our sample size on our planned analyses: Bujang et al. [62] suggest  $n=500$  or  $n=100 + 50 \times \#IVs$  as the minimum sample size for logistic regressions. For the linear regression (RQ4), G\*Power suggests  $n=127$  for detecting medium effects ( $f^2=.15$ ), with  $\alpha=.05$ ,  $\beta=.80$ . We met or exceeded these thresholds with 413 participants (435 email-specific responses; 792 breach-specific responses).

97% of participants passed our generic attention check. Of the 302 participants who were shown at least one breach, only 55% passed the breach-specific attention check, whereas the rest chose “none of these” (42%) or a decoy option (3%). We reviewed open-ended responses from participants who failed this attention check, and all of them were detailed and insightful. We also did not find significant correlations between this attention check’s performance and participants’ breach-specific responses about awareness (chi-squared test,  $\chi(1)=.06$ ,  $p=0.8$ ), concern level (Mann Whitney test,  $W=58395$ ,  $p=0.2$ ), and whether they had taken action (chi-squared test,  $\chi(1)=.29$ ,  $p=0.6$ ). Thus, we did not exclude any of these participants as our findings suggest the question was not a reliable exclusion criterion.

**Qualitative analysis.** We analyzed participants’ open-ended responses using inductive coding [434]. For Questions 7, 10, 14, 16, and 18, a primary coder created an initial codebook based on all responses. Multiple coders then iteratively improved the codebook. A second coder analyzed 20% of responses to each question to ensure high inter-rater reliability. Cohen’s  $\kappa$  were 0.89 (Q7), 0.73 (Q10), 0.74 (Q14), 0.81 (Q16), and 0.78 (Q18). We resolved all coding discrepancies through discussions. Appendix B.2 includes the codebook, with common themes highlighted.

**Statistical analysis.** We conducted regressions to identify influential factors for breach status (RQ1), awareness (RQ3), emotional response (RQ4), and behavioral response (RQ5). We included a random intercept for individual participants to account for repeated observations between multiple breaches. For models corresponding to RQ1, the random effects were close to zero and caused a boundary singularity fit, so we conducted single-level regressions instead.

For all models, we treated participant demographics as control variables. We report a model’s output with participant demographics when it has a significantly better fit than the model without; otherwise, we opt for the simpler model in reporting the results. To avoid model fit problems caused by too few observations in a category, we binned the demographic data in Table 4.1 into fewer categories for the regression analyses: gender (binary, men or women), age (three levels: 18-34, 35-54, 55+), and educational attainment (three levels: no Bachelor’s degree, Bachelor’s degree, Graduate degree). We treated participants’ responses of concern level as a continuous variable in our regressions, which has limitations, as we discuss below.

**Limitations.** As with most surveys, parts of our findings rely on self-reported data, which is prone to biases. For instance, prior work has shown a gap between self-reported behavioral intentions and actual behaviors in security contexts [245] and beyond [453]. We do not imply that all participants would take the actions they reported. Participants’ self-reported intentions are still valuable insights to inform ideas about how to protect consumers against data breaches.

HIBP’s API does not return sensitive breaches such as those involving adult sites. Accessing these breaches requires sending a confirmation message to participant-provided email addresses for ownership verification. We decided not to do this as participants might infer that we store their email addresses even though we do not.

Our study only included data breaches involving email addresses, which may not

represent all breaches (e.g., only 4% of breaches recorded by Privacy Rights Clearinghouse [405] included email addresses). Relatedly, the email-focused nature of these breaches means it is difficult to track whether and how breached organizations in our sample notified affected individuals and how that impacts consumer reactions because existing breach notification databases mostly document letter-based notifications [600]. Future research can look into breaches that expose a broader range of data types and consider organizations' handling of breaches when feasible.

Regarding our analyses, we considered several options for treating the Likert responses of concern level: ordinal, nominal, or continuous. Treating concern as ordinal would introduce square and cubic effects into the model—these effects are difficult to interpret and inconsistent with the scale. Treating concern as nominal would lose information about the scale's ordering and prevent comparisons across all levels (e.g., with “not at all concerned” as the baseline, the regression would not describe the difference when moving up or down the scale between “slightly concerned” and “extremely concerned”). Treating concern as continuous would require a more cautious interpretation of the p-values in the analysis, and it assumes equal differences between the scale items. After discussions with our university's statistical consulting service, we followed their advice and decided to treat concern as a continuous variable. While this comes with the limitations mentioned above, it also allows a more straightforward and meaningful interpretation of results, which we prioritize to make the results more accessible.

## 4.2 Findings

We first describe the participants and breaches in our sample, then summarize findings corresponding to each research question.

### 4.2.1 Sample

Table 4.1 summarizes our 413 participants’ demographics and breach status. Our participants were evenly distributed between men and women but skewed educated and younger. 122 (30%) described having a background in information technology; 25 (6%) in law. In total, participants provided 435 email addresses. 421 (97%) accounts belonged to the participant exclusively, and ten were shared with someone else. Four were either someone else’s account or a made-up address for the study and were removed from the data. Participants whose initial email address was not exposed in any breach could scan another: 393 participants (95%) scanned only one email address, 18 scanned two addresses, and only two scanned three addresses.

For the 431 owned or shared email accounts, we further asked participants how long they had been using the email account, how frequently they checked it, and what they primarily used it for. The majority of email accounts were used for an extended period (mean: 8.75 years, median: 8). Most (81%) were checked daily; the rest were checked less frequently (14% weekly, 4% monthly, and 1% yearly). In addition, participants reported multiple uses for their email address (mean: 2.74, median: 3): 74% were used for personal correspondence, followed by signing up for medium-sensitive accounts like social media (68%), signing up for sensitive accounts like banking (51%), signing up for low-value accounts (49%), and professional correspondence (32%).

**Overview of breaches.** We observed 189 unique breaches across 431 email addresses queried against HIBP. 302 (70%) email addresses, or 73% of participants, were exposed in one or more breaches. The average number of breaches per email address was 5.12 (median: 3, sd: 6.21, max: 46), or 5.36 per participant (median: 3, sd: 6.23). The number of breaches per email address formed a long-tail distribution: 34% of email addresses appeared in 1 to 5 breaches, and only 2% were associated with 21 or more breaches.

	Total	Num. (%) W/ Breaches	Num. (%) W/o Breaches	Avg. (Med./Std.) Breaches
Men	199	139 (70%)	60 (30%)	4.49 (2/5.97)
Women	212	162 (76%)	50 (24%)	6.11 (4/6.28)
Non-Binary	2	1 (50%)	1 (50%)	11.00 (11/11.00)
18-24	77	56 (73%)	21 (27%)	3.90 (2/5.15)
25-29	51	35 (69%)	16 (31%)	4.25 (2/4.90)
30-34	42	33 (79%)	9 (21%)	6.55 (3/8.72)
35-39	49	29 (59%)	20 (41%)	4.63 (1/7.05)
40-44	45	26 (58%)	19 (42%)	4.36 (2/5.04)
45-49	32	29 (91%)	3 (9%)	6.59 (4/6.05)
50-54	39	30 (77%)	9 (23%)	6.72 (6/6.16)
54-59	34	30 (88%)	4 (12%)	6.12 (5/4.82)
60-64	27	19 (70%)	8 (30%)	6.52 (3/6.85)
65+	17	15 (88%)	2 (12%)	8.24 (8/6.06)
Some High School	1	0 (0%)	1 (100%)	0.00 (0/0.00)
High School or Equiv.	46	35 (76%)	11 (24%)	4.59 (3/4.61)
Some College	88	70 (80%)	18 (20%)	5.67 (3/6.63)
Associate (voc./occ.)	14	14 (100%)	0 (0%)	8.07 (6/6.51)
Associate (aca.)	20	19 (95%)	1 (5%)	6.10 (4/5.99)
Bachelor	140	108 (77%)	32 (23%)	6.04 (4/6.56)
Master	83	46 (55%)	37 (45%)	4.10 (2/5.68)
Professional	5	4 (80%)	1 (20%)	11.60 (13/7.71)
Doctorate	16	6 (38%)	10 (62%)	1.44 (0/2.26)
IT Background	122	67 (55%)	55 (45%)	3.82 (1/6.30)
No IT Background	278	224 (81%)	54 (19%)	5.91 (4/6.06)
Prefer not to say	13	11 (85%)	2 (15%)	8.00 (9/6.41)
Law Background	25	14 (56%)	11 (44%)	5.80 (2/9.63)
No Law Background	374	278 (74%)	96 (26%)	5.29 (3/5.93)
Prefer not to say	14	10 (71%)	4 (29%)	6.36 (5/6.25)
No Data	170	115 (68%)	55 (32%)	4.45 (2/6.21)
<\$15K	16	15 (94%)	1 (6%)	7.81 (4/8.59)
\$15K-\$25K	22	20 (91%)	2 (9%)	6.77 (4/5.79)
\$25K-\$35K	28	26 (93%)	2 (7%)	5.89 (3/5.37)
\$35K-\$50K	26	19 (73%)	7 (27%)	4.58 (2/5.35)
\$50K-\$75K	45	40 (89%)	5 (11%)	8.04 (7/6.50)
\$75K-\$100K	38	28 (74%)	10 (26%)	6.95 (4/6.61)
\$100K-\$150K	37	22 (59%)	15 (41%)	4.05 (2/4.63)
>\$150K	24	13 (54%)	11 (46%)	3.92 (2/5.34)
<b>Total</b>	<b>413</b>	<b>302 (73%)</b>	<b>111 (27%)</b>	<b>5.36 (3/6.23)</b>

Table 4.1: Participant demographics and breach status ( $n=413$ ).






Figure 4.2: Frequency of the leaked data types for 189 breaches, excluding email address (appears in all breaches). 44 other types occurring twice or fewer.

For the 189 unique breaches, we examined their date, the total amount of breached accounts, and the types of compromised data according to HIBP. The majority (69%) of breaches occurred in 2015–2019; 15 breaches occurred in 2020. The average number of breached accounts captured by HIBP was 46.52M (median: 4.79M; sd: 125M), indicating a distribution skewed by several large breaches (max: 772.90M). Our sample’s breaches covered 66 different data types. The average number of leaked data types per breach was 4.86, and the maximum was 20 (median: 4, sd: 2.58). Aside from participants’ email addresses (which were present in all breaches as HIBP uses them as references), the other commonly breached data types included passwords (162, 86%), usernames (110, 58%), IP addresses (82, 43%), names (74, 39%), and dates of birth (47, 25%). The frequency distribution of data types in our sample’s breaches falls steeply (see Figure 4.2), suggesting a broad range of leaked data types with a much smaller set of commonly leaked data.

We used Cisco’s website content taxonomy<sup>3</sup> for cross-referencing breached orga-

<sup>3</sup><https://talosintelligence.com/categories>

nizations’ industry, excluding 25 (13%) non-applicable cases.<sup>4</sup> Gaming companies were represented the most in our sample (40, 21%). Other represented industries included general business (17, 9%), computers/Internet (16, 8%), shopping (10, 5%), and online communities (10, 5%). We used Alexa’s ranking of global websites<sup>5</sup> as of October 14, 2020, as a proxy for a breached organization’s popularity.<sup>6</sup> Excluding 33 organizations with missing data, the average ranking was 650.73k (median: 24.85k, sd: 1,768k). There were 19 organizations in the top 1K list, indicating that while the majority of organizations in our sample were not mainstream, a few were relatively well-known.

#### 4.2.2 RQ1: Likelihood of Breaches

We conducted a logistic regression on whether an email address had been breached in relation to the email account’s age, checking frequency, and purpose of use. Results in Table 4.2 show that an email address was significantly more likely to be breached as the account’s age in years increased ( $OR_{age}=1.35, p<.001$ ), as it was checked daily instead of weekly ( $OR_{daily}^{weekly}=2.30, p=.03$ ), and as it was used for personal correspondence ( $OR_{yes}^{no}=2.13, p=.02$ ). Additionally, the significant intercept indicates that an email address was significantly unlikely to be associated with any breach if the email account was just created, checked weekly, and not used for any correspondence or account creation purposes ( $OR_{intercept}=0.14, p=.002$ ). Essentially, the less frequently used and newer an email address is, the less likely it is to be exposed in a data breach.

We further conducted a quasi-Poisson regression on the number of breaches per email address with the same independent variables as above. We chose quasi-Poisson

---

<sup>4</sup>These breaches were spam lists or aggregate credential stuffing lists, or the breached organizations were no longer active.

<sup>5</sup><https://alexa.com/topsites>

<sup>6</sup>We used rankings at the time of analysis rather than historic ranking (i.e., the ranking when the breach occurred) because (1) Alexa only provides ranking data for the last four years; and (2) we anticipate that current ranking would better reflect participants’ impression of the organization’s popularity at the time when they took our study.

	Est.	OR	95% CI	p-value
(Intercept)	-1.95	0.14	[0.04, 0.49]	.002
Freq. Checked daily (vs. weekly)	0.83	2.30	[1.07, 4.99]	.03
Prof. Corr. yes (vs. no)	-0.02	0.98	[0.51, 1.87]	.94
Pers. Corr. yes (vs. no)	0.76	2.13	[1.13, 4.03]	.02
Acct. Creat. yes (vs. no)	0.31	1.36	[0.60, 3.07]	.46
Email age years	0.30	1.35	[1.26, 1.46]	< .001
Age: 35-54 (vs. 18-34)	-0.51	0.60	[0.29, 1.23]	.16
Age: 55+ (vs. 18-34)	-0.60	0.55	[0.27, 1.10]	.09
Gender: men (vs. women)	-0.24	0.79	[0.43, 1.45]	0.45
Edu.: =Bach. (vs. <Bach.)	0.25	1.28	[0.65, 2.53]	0.48
Edu.: >Bach. (vs. <Bach.)	-0.62	0.54	[0.25, 1.16]	.11
Occu.: IT/law yes (vs. no)	-0.51	0.60	[0.31, 1.17]	.14

Table 4.2: Logistic regression for breach status of an email address (leaked vs. not leaked).

	Est.	Exp (Est.)	SE	p-value
(Intercept)	0.67	1.94	0.26	.01
Freq. Checked daily (vs. weekly)	0.36	1.43	0.19	.06
Prof. Corr. yes (vs. no)	-0.11	0.89	0.12	.33
Pers. Corr. yes (vs. no)	0.29	1.34	0.15	.06
Acct. Creat. yes (vs. no)	-0.18	0.83	0.15	.22
Email age years	0.08	1.08	0.01	< .001
Age: 35-54 (vs. 18-34)	-0.29	0.75	0.14	.045
Age: 55+ (vs. 18-34)	-0.35	0.71	0.14	.02
Gender: men (vs. women)	-0.18	0.84	0.12	.13
Edu.: =Bach. (vs. <Bach.)	0.17	1.18	0.12	.18
Edu.: >Bach. (vs. <Bach.)	-0.17	0.84	0.16	.29
Occu.: IT/law yes (vs. no)	-0.05	0.95	0.14	.70

Table 4.3: Quasi-poisson regression regarding the number of breaches per email address.

regression because the dependent variable is count data with a skewed distribution [598]. Results in Table 4.3 show how the number of breaches increases with an email account's age: for every one year of increase in age, the expected number of breaches increases by a factor of  $exp(0.08) = 1.08$  ( $p < .001$ ). In other words, the number of breaches increases 8% per year of use, compounding yearly. A possible explanation is that the older an email address is, the more it has been used for account registrations, which increases its presence in organizations' databases. The significant intercept in Table 4.3 confirms this finding: a new and rarely used email address is more immune to breaches. Furthermore, the number of breaches per email address differed among age groups: compared to young adults (18-34), the number of breaches decreases by a factor of  $exp(-0.29) = 0.75$  ( $p = .045$ ) for middle-aged adults (35-54) and by a factor of  $exp(-0.35) = 0.71$  ( $p = .02$ ) for older adults (55+).

### 4.2.3 RQ2: Perceived Causes and Impacts of Breaches

We asked participants to speculate why or why not their email address was part of a data breach and name any experienced or anticipated future impacts from a specific breach.

**Perceived reasons for being affected by breaches.** We analyzed 302 open-ended responses to Question 10 in which participants speculated why their email address was exposed in one or more data breaches. The most common explanation, cited in 159 (53%) cases, was that it was due to participants' email-related practices. Specifically, 70 (23%) mentioned using the email address to sign up for many different sites (e.g., *"it's on the website of every business I have an online relationship with"*). Another 31 (10%) mentioned the email's age as a relevant factor, saying it had been used for a long time. 23 (8%) expressed that breaches were inevitable, especially for an old or widely-used email address (e.g., *"there are a lot of companies or organizations that have my email [address] and chances are one of them is going to get hacked"*). Furthermore, in 31 (10%) cases, participants mentioned using the email to sign up for seemingly sketchy websites, sometimes with a clear intention to do so despite knowing that the website might be insecure.

Participants mentioned other insecure behaviors as potential reasons for being affected by a breach in 31 (10%) cases. 13 cases referred to password-related behaviors, such as using simple passwords, reusing a password across accounts, or not changing passwords frequently. Incautious clicking behavior was mentioned five times (e.g., *"because I was not careful with what emails I clicked"*). Other participants indicated their exposure to breaches was due to infrequent monitoring of the email account, easily guessed answers for security questions, or being signed into the email account for too long. While these are indeed insecure behaviors, password choices do not impact one's likelihood of being involved in a breach; they impact a breach's consequences

by increasing the possibility of account hijacking due to credential stuffing. Similarly, clicking on untrustworthy links may make the email address appear in spam lists, which will be reported by HIBP if found on the public web. However, this action does not increase one’s vulnerability to breaches.

Only 42 (14%) of participants accurately attributed the cause of being affected by a breach to external factors unrelated to their behaviors. 26 (9%) blamed it on lax security measures by the breached organization (e.g., *“these companies did not try hard enough to keep information private”*). 16 (5%) blamed it on bad actors such as hackers and scammers targeting the breached organization (e.g., *“hackers are devious devils and learn to adapt faster than organizations can protect users”*). Another 15 (5%) suspected their email address was sold by the breached organization or a third party. Nevertheless, nine participants incorrectly blamed their email provider’s security (e.g., *“I feel like Hotmail has poor security and cannot block as many spam emails compared to Gmail”*).

**Perceived reasons for not being affected by breaches.** Question 7 asked participants to speculate why their email address was *not* involved in any data breach. Among the 136 provided responses, 78 (57%) mentioned cautious email practices. Specifically, 31 (23%) reported using their email address to sign up for trusted sites only, sometimes with careful examination of the website (e.g., *“I try as much as possible to scrutinize websites before dropping any of my details”*). 18 (13%) mentioned that their email address was relatively new or did not get used much. Ten further mentioned limiting the email to specific purposes, such as correspondence with friends and family members only.

Eight participants described using multiple email accounts for different purposes, e.g., using one email address for correspondence exclusively and another for account registration on “low-value” sites. Such behavior would likely reduce the likelihood of

breaches involving high-value email addresses. However, breaches involving low-value email addresses may still have real impacts, such as account hijacking.

21 (15%) participants cited their security practices as reasons for not being affected. Nine participants mentioned their password practices, such as using strong/unique passwords and changing passwords regularly. Less frequently mentioned were two-factor authentication, anti-virus, firewall, and VPN. None of these behaviors are likely to prevent data breaches despite potentially having other positive security outcomes.

**Experienced and anticipated impacts of data breaches.** Participants with at least one breach were asked to describe a given breach’s experienced or potential impacts (Question 16). Of the 792 responses, more than half assessed the breach’s impact as none (343, 43%) or very little (85, 11%); another 77 (10%) were unsure. Only 19 (4%) breaches were perceived as having a significant impact. In 135 (17%) cases, participants described emotional feelings without naming concrete impacts, such as *“no impact just rage.”*

In 149 (19%) instances, participants described specific experienced impacts or anticipated future impacts. The most prevalent was an increase in spam emails, text messages, etc. Some participants reported scam phone calls, and others anticipated identity theft as a potential impact (e.g., *“I suppose now that someone has all that information about me they could impersonate me, open credit lines in my name, scam my family and friends”*). Participants who experienced adverse events described emotional stress and resulting behavioral changes, such as avoiding phone calls due to frequent scams or checking emails for suspicious activities after account compromises.

Notably, participants with and without experienced impacts differed in assessing the impact’s severity. Most participants who described anticipated impacts but had not experienced them did not foresee real consequences (e.g., *“the only things that [would] really happen is . . . scammers . . . occasionally attempt to access some of my*

*older accounts that hold no sensitive information*”). This underlines that participants’ perception of impacts after being affected by breaches largely depends on individual circumstances. The finding also aligns with prior work [602, 603] showing that people don’t adopt secure behaviors until experiencing actual harm.

#### 4.2.4 RQ3: Awareness of Breaches

Among the 792 breach-specific responses, 590 (74%) reported unawareness of being affected by the breach before our study. Only 143 (18%) reported prior awareness, and the other 8% were unsure. Participants who were previously aware of the breach mostly learned about it from the breached organization (45, 31%) or third-party notification services (45, 31%). Less common sources included news media (17, 12%), credit/identity monitoring services (14, 10%), bank or credit card companies (3, 2%), experiencing adverse events (3, 2%), and someone else (3, 2%). In nine instances, participants could not remember how they learned about the breach.

Using a mixed-effect logistic regression to identify factors that might impact awareness (excluding “unsure” responses), we included the same email-related factors from Table 4.2 as independent variables. Additionally, we included breach age (i.e., the time lapse between a breach’s occurrence and the participant taking our study), hypothesizing that participants are more likely to recall and report awareness of recent breaches.

Results in Table 4.4 show a significant intercept, indicating that participants were more likely to be unaware of a breach if they have a newer email address and the breach just occurred ( $OR_{intercept}=0.01$ ,  $p<.001$ ). Participants were also significantly more likely to be aware of a breach as the breach’s age in years increased ( $OR_{breach\_age}=1.22$ ,  $p<.001$ ). Older participants were less likely to be aware of breaches than young participants ( $OR_{55+}^{18-34}=0.39$ ,  $p=.049$ ), and men were more likely to be aware of a breach than women in our sample ( $OR_{men}^{women}=2.09$ ,  $p=.049$ ), though p-values in both cases

	<b>Est.</b>	<b>OR</b>	<b>95% CI</b>	<b>p-value</b>
(Intercept)	-4.24	0.01	[0.002, 0.09]	< .001
Freq. Checked daily (vs. weekly)	0.31	1.37	[0.45, 4.16]	.58
Prof. Corr. yes (vs. no)	-0.06	0.94	[0.45, 1.98]	.88
Pers. Corr. yes (vs. no)	0.22	1.25	[0.50, 3.10]	.63
Acct. Creat. yes (vs. no)	0.77	2.15	[0.70, 6.63]	.18
Email age years	0.04	1.04	[0.98, 1.11]	.17
Breach age years	0.20	1.22	[1.09, 1.35]	< .001
Age: 35-54 (vs. 18-34)	-0.41	0.66	[0.27, 1.61]	.36
Age: 55+ (vs. 18-34)	-0.94	0.39	[0.15, 1.00]	.049
Gender: men (vs. women)	0.74	2.09	[1.00, 4.37]	.049
Edu.: =Bach. (vs. <Bach.)	-0.79	0.45	[0.20, 1.00]	.051
Edu.: >Bach. (vs. <Bach.)	-0.18	0.84	[0.31, 2.22]	.72
Occu.: IT/law yes (vs. no)	0.50	1.65	[0.72, 3.77]	.23

Table 4.4: Logistic regression regarding prior breach awareness.

are close to 0.05. These findings align with prior work in which adopting protective behaviors differed by age [266] and gender [454, 603]. Other demographic variables and email-related factors are not significantly correlated with prior awareness.

#### 4.2.5 RQ4: Emotional Responses Towards Breaches

Participants indicated their concern using a 5-point Likert item for each shown breach (Question 15) and each data type leaked in a breach (Question 17). We also asked participants to describe their feelings regarding the breach (Question 14, open-ended).



**Quantitative ratings of concern level.** Among 792 breach-specific responses, the median concern level regarding the breach was “somewhat concerned.” Less than half reported either no concern (151, 19%) or being very/extremely concerned (197, 25% combined). Figure 4.3 shows concern levels for commonly leaked data types. Participants were most concerned about leaks of physical addresses (52% very/extremely), passwords (47% very/extremely), and phone numbers (42% very/extremely). Other leaked data types that participants felt less concerned about were employer information (38% not at all), social media profile (42% not at all), job title (46% not at all), and gender (65% not at all).

We sought to identify factors that might impact concern level through a mixed-effect linear regression on overall concern Likert responses. We included email address-related factors and prior awareness as independent variables, hypothesizing that participants would be more concerned about frequently used email addresses or if they had not been aware of a breach. We also included the number of breached data types and the breach status of data types for which more than 50% of responses were “somewhat concerned” or above in Figure 4.3, namely password, physical address, phone number, date of birth, IP address, and name.<sup>7</sup> We hypothesized that the concern level would increase as the amount or sensitivity of leaked data types increases. Treating concern (measured as a 5-point Likert item) as a numeric variable has its limitations, as we discussed in Section 4.1. Additionally, we included the breaches’ age since participants might be more concerned about recent breaches.

The regression results do not reveal any significant factors impacting overall concern except the intercept ( $b_{intercept}=2.52$ ,  $SE=.31$ ,  $p<.001$ ), indicating that participants likely default to between “slightly concerned” and “somewhat concerned.” The model’s  $f^2 = 0.03$  indicates a small effect size. The absence of influential factors on concern may be due to data types known to trigger more concerns, such as finan-

---

<sup>7</sup>Email address was not included because it was exposed in all breaches in our sample.




Figure 4.3: Overall concern about the breach and levels of concern for the 13 most commonly leaked information types in our sample breaches.

cial information and social security numbers, being underrepresented in our sample’s breaches (see Figure 4.2). Even relatively sensitive data types in our sample still had a fair number of “not at all/slightly concerned” responses.

**Various emotions in qualitative responses.** Figure 4.4 shows the wide range of emotions reflected in participants’ open-ended responses about their feelings after learning of a breach affecting them. In 237 (30%) cases, participants reported feeling upset (including annoyed, frustrated, mad, and angry), mainly toward the breached organization. The upset came from not having been properly informed (e.g., “*I was very disappointed . . . they hid the fact that there was a data breach from everyone for three months*”), the organization’s poor security measures (e.g., “*don’t run an entirely online business if you cant do basic security*”), or violation of consumers’ trust (e.g., “*I joined this site to read a story my granddaughter had written and thought it was completely safe*”). These emotions align with the “risk as feelings” theory, which highlights that people experience dread and outrage in comprehending risks [468], and




Figure 4.4: Code frequencies for feelings after first learning about a breach ( $n = 792$ ); red bars indicate negative feelings, gray neutral, blue positive, according to Emolex ratings [350].

such affective responses greatly influence their subsequent decision-making, sometimes overriding cognitive assessments [304].

Mirroring the Likert responses, feeling unconcerned about a breach was common (185, 23%). Many participants believed that the exposed data was not sensitive (e.g., “*I had only used the free version of that site, so I had not entered any payment information*”). Others were unconcerned because they rarely interacted with nor knew the breached organization (e.g., “*I don’t even know what this site is, so I don’t think that them having my info ... is a huge deal*”). Some were unconcerned due to confidence in their security habits, including regularly changing passwords (25), avoiding password reuse (10), and enabling 2FA (4). A few participants were unconcerned due to a lack of experienced impacts (e.g., “*I’m not especially worried because I haven’t detected any suspicious activity*”) or optimism bias (e.g., “*I feel like a drop in the bucket since there were 711 million emails affected*”). 104 (13%) responses reported feeling unsurprised, whereas 66 (8%) reported feeling surprised. Unsurprised participants explained that they never trusted the breached organization or already knew about

the breach. Conversely, surprised participants stated that they had never used the breached organization’s service or trusted the organization.

In another 75 (9%) cases, participants expressed confusion due to unfamiliarity with the breached organization or not remembering having an account. Other prominent emotions included fatigued (43, 5%), violated (40, 5%), indifferent (33, 4%), scared (29, 4%), unsafe (18, 2%), relieved (18, 2%), or curious about why the breach happened (13, 2%). Those who expressed fatigue stressed that breaches were inevitable (e.g., *“It’s the internet and things WILL be leaked somehow, either by hackers or by incompetence at the company that is holding your information anyhow”*). This attitude is akin to the “digital resignation” phenomenon [129]: many people’s inaction in the face of privacy infringements is not necessarily because they do not care but because they are resigned and convinced that surveillance is inescapable. Notably, neutral emotions, like curiosity, or positive emotions, like relief, were rare. Participants were relieved when sensitive data like financial information was not involved or that they were now aware of the breach and could take proper action.

#### **4.2.6 RQ5: Behavioral Reactions to Breaches**

For the 143 breaches participants were already aware of before our study, we further asked if they had taken any action in response (Questions 18). The most common action was changing passwords (87, 61%): 15 specified they changed the password for the breached account, and 27 mentioned changing the password across multiple accounts that might use the leaked password. Five further mentioned changing their email account’s password; this could be due to a misconception that their email account, not the account with the breached organization, was compromised. Participants also described other password-related practices triggered by the breach, such as using unique passwords, using a password manager, and making passwords more complicated.

Participants reported taking various actions related to their account with the breached organization. For example, 18 (13%) deleted or deactivated the account, and one mentioned reviewing accounts on other websites and deleting them as needed. Five mentioned enabling 2FA for the account on the breached site, other accounts, or their email account. Four reported checking the breached site’s account to see if it stored any sensitive data or if there had been any suspicious activity. In 31 (22%) cases, participants reported doing nothing in reaction; the percentage was lower than that in Ponemon’s 2014 survey (32%) [403], but still substantial.

Additionally, we asked all participants with at least one breach to indicate, for each breach, how likely they were to initiate ten provided actions within the next 30 days or whether they had taken action already. We only include 500 breach-specific responses in the following analysis due to a data storage issue, excluding incomplete responses. Figure 4.5 shows the results. Of the ten provided actions, changing the password for the breached site’s account or other accounts was the most popular, receiving more than half of likely/already done responses. “Review credit reports and/or financial statements” had the highest percentage of already done (30%). By contrast, most participants selected “not likely” for four actions—“use a credit/identity monitoring service,” “place a credit freeze on my credit reports,” “file a complaint with a consumer protection agency,” and “take legal action against the breached organization.” This finding is understandable given that most leaked data types such as email addresses and passwords are considered “non-sensitive records” according to ITRC’s report [236].

We sought to understand factors that would impact the likelihood of having taken any of the ten provided actions through a mixed-effect logistic regression. For independent variables, we discarded variables related to email habits since many of the listed actions were unrelated to one’s email account. We kept all other independent variables from the concern regression model: prior awareness, the breach’s age, the number of breached data types, and the breach status of six data types with rela-




Figure 4.5: Intention to take actions within the next 30 days.

tively high concern levels. We further included overall concern Likert responses as an independent variable. Results in Table 4.5 show a significant intercept, indicating that participants were likely to default to inaction with no leaked data and no prior awareness or concern ( $OR_{intercept}=0.04, p=.02$ ). Being aware of a breach significantly increased the likelihood of having taken any of the listed actions ( $OR_{yes}^{no}=390.48, p<.001$ ). This is unsurprising given that participants unaware of being affected had little motivation to engage in protective measures. Additionally, more concern was significantly correlated with a higher likelihood of having taken action: for a one-unit increase of concern on the 5-point scale, the odds of having taken action increase by 2.22 ( $OR_{concern}=2.22, p=.005$ ).

### 4.3 Discussion

We examined individuals' awareness, perception, and responses to specific data breaches that had exposed their email addresses and other information. Compared to RAND's 2016 survey [1], in which 44% reported already knowing about a breach before receiving a notification, participants' prior awareness was much lower in our

	<b>Est.</b>	<b>OR</b>	<b>95% CI</b>	<b>p-value</b>
(Intercept)	-3.27	0.04	[0.002, 0.61]	.02
Awareness yes (vs. no)	5.97	390.48	[45.72, 3334.79]	< 0.001
Breach age years	-0.03	0.97	[0.77, 1.21]	.77
Num. of types numeric	.12	1.13	[0.85, 1.50]	.39
Password yes (vs. no)	-0.18	0.84	[0.18, 3.79]	.82
Physical Addr. yes (vs. no)	-0.26	0.77	[0.16, 3.71]	.75
Phone Num. yes (vs. no)	-0.29	0.75	[0.19, 3.02]	.69
Date of birth yes (vs. no)	-0.24	0.79	[0.17, 3.62]	.76
IP Addr. yes (vs. no)	-0.20	0.82	[0.26, 2.64]	.74
Name yes (vs. no)	-0.19	0.83	[0.21, 3.22]	.79
Concern numeric	0.80	2.22	[1.28, 3.86]	.005

Table 4.5: Logistic regression on taking actions.

sample. This finding is concerning as our results suggest that unawareness creates a substantial barrier to taking mitigating action. Participants also reported a lower level of overall concern than in prior work [260, 403]: this might result from a methodological difference, as our participants reflected on specific breaches affecting them rather than on breaches in general [1, 403] or on hypothetical scenarios [260]. Another possible reason is that the leaked data types in the HIBP database are mostly categorized as non-sensitive records [236]. While participants named potential consequences of data breaches such as more spam and increased risks of identity theft, similar to prior work [260, 602], many considered these events would have little to no impact on their lives. Most participants also exhibited misconceptions about what led to themselves being affected by breaches, blaming their email or password practices rather than the breached organization.

**Set stricter legal requirements for notifying consumers.** Our study reflects a sad reality that many individuals are unaware that they are affected by breaches, at least for breaches exposing email addresses. This finding indicates that current breach notification requirements, mechanisms, and tools fail to reach data breach victims. Nonetheless, according to our regression results, awareness was a crucial trigger for taking action.

Stricter regulatory requirements may help establish high standards for breach notifications, which could raise more awareness. However, simply requiring companies to send notifications is not enough as the notification also needs to be effective [48, 600]. For instance, prior work highlights the role of media reports in informing and shaping attitudes toward data breaches [1, 110]. Our findings indicate that notifications from breached organizations or third-party services are more relevant. Given that individuals may not stick with one channel to learn about breaches, breached organizations could be mandated to notify consumers in multiple channels instead of the most convenient channel and seek confirmation from the recipient. Regarding when to notify, Art. 34 of GDPR specifies that consumer-facing notifications are only needed for breaches that “result in a high risk” to data subjects [144]. We argue that this should be done for *all* breaches, given that many court cases struggle to assess risks and harms caused by data breaches [475]. Alternatively, less ambiguous criteria should be set for high-risk breaches, e.g., in California, consumer-facing notifications are mandated when the breach involves unencrypted personally identifiable information [481].

**Use novel approaches in notifying consumers.** Prior research on SSL warnings [16, 149] shows that in-browser warnings effectively raise threat awareness and encourage safer practices. Similarly, data breach notifications could explore approaches beyond letters and emails, such as in-situ methods whereby visiting affected sites leads



to a notification [114], as recently pursued by some browsers and password managers that warn users if saved passwords appear in credential dumps [278, 398].

Notifications should also consider non-adherence: among participants who were already aware of a breach before our study, 22% reported doing nothing in response to that breach; emotions like fatigue and resignation were also noted. Drawing from warning design literature on mitigating fatigue in email-based notifications, one could build systems that highlight unread breach notifications in email clients, similar to Gmail’s reminders to reply to emails [56]. The contents of such emails could also be automatically parsed and reformatted to guide attention to important details.

**Address misconceptions.** Participants commonly blamed their own email habits or security practices for data breaches, and such misconceptions exacerbate a power asymmetry—rather than demanding that organizations improve security measures or that regulators hold them accountable, participants blamed themselves. Consumers should be reminded that the root cause of breaches is security issues in the breached organization, and there are actions that can hold the breached organization accountable, such as filing a complaint with a consumer protection agency (e.g., the Federal Trade Commission for breaches in the US).

Participants also differed regarding the perceived impacts of breaches. Those who had not experienced negative consequences in real life tended to take the breach less seriously. Conversely, those who had experienced an adverse event reported emotional distress and behavioral changes. Indeed, not everyone would experience the negative consequences of not reacting to data breaches, but the cost is real and immediate when the consequences manifest. Therefore, breach notifications and education materials should stress that good security practices, such as using unique passwords and 2FA, can dampen the severity of a breach’s impact, even though they do not decrease one’s likelihood of being affected by a breach. While these preventive measures

might not provide instant gratification, they could be worthy investments considering the substantial hassles and trauma in recovering from identity theft [235] or other repercussions.

**Develop tools to help consumers react to breaches.** While consumers may not be able to prevent breaches, there are actions one could take to mitigate the aftermath of a breach. Our findings show that some straightforward actions, such as changing passwords, had high adoption rates or intention to adopt. Yet, the majority of actions we provided were much less popular, indicating the need to offer more relevant and usable protective measures to affected individuals.

One of our key findings is that extensive use of an email account (e.g., use it for a long time and check it frequently) significantly increased the email address’s likelihood of being involved in a breach. Yet, simply asking users to reduce their usage or abandon their email accounts is not a viable solution, as it also diminishes the email account’s utility. Instead, drawing from some participants’ descriptions of creating dedicated email accounts for registration on low-value sites, we see the promise of more automated tools to offer unique email aliases for account registration. Such features could further be integrated into other technologies with broader adoption, such as browsers or password managers, to create a more streamlined experience (e.g., through auto-filling). Recent respective efforts include “Sign in with Apple” [30] and “Firefox Relay” [355], both of which support the generation of a unique, random email address during account registration, which is forwarded to a user’s actual inbox. However, both products are currently limited to their respective ecosystems. The effectiveness, awareness, and adoption of such tools, as well as how individuals manage multiple email aliases in general, are open questions for future research.

**Increasing responsibilities of breached organizations.** Our participants exhibited a low awareness of data breaches, which in turn serves as a precursor to the

low intention for certain protective measures. Participants' lack of awareness and self-protection indicates that breached organizations should play a more active role in protecting affected individuals. Notifying victims should not absolve breached organizations from further responsibility—they should further ensure that consumers have viable remediation solutions and assist in the recovery process, such as offering support in identity restoration. For example, rather than defaulting to conventional credit and identity monitoring services, which are known to provide little preventative protection [276], breached organizations could offer victims email alias generators, password managers, or other more promising mitigation tools by partnering with respective service providers. Regulators should also set and frequently revisit requirements for the types of services breached organizations must offer as compensation.

Importantly, breached organizations have financial incentives for transparent post-breach communications and active mitigation. Prior work shows that data breach notifications provide a venue for impression management and repairing damaged trust [244]. Moreover, breached organizations that provide affected individuals with free credit monitoring services face a lower likelihood of lawsuits [425]. Regulators should also create meaningful incentives for organizations to act accordingly. For instance, GDPR's threat of substantial fines has resulted in a heightened effort by organizations worldwide to overhaul their privacy and security programs.

## CHAPTER V

# An Empirical Analysis of Data Breach Notifications<sup>1</sup>

In the previous two chapters, I document two studies with complementary methods to examine consumer reactions to data breaches. Among other factors, optimism bias, insufficient knowledge of available protective measures, and a general tendency to delay action until harm has occurred dissuaded consumers from reacting to the 2017 Equifax data breach. For breaches curated by HIBP that exposed individuals' email addresses and other personal information, most participants were unaware that they were affected by these breaches; while some reported intending to take action, many participants believed the breach would not impact them. Consumers' lack of awareness and action begs the question: are current data breach notifications effective?

In this chapter, I present a content analysis of 161 notifications sent by companies to US consumers between January and June 2018. We analyzed their readability, structure, risk communication, and presentation of recommended actions. We found that 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

---

<sup>1</sup>This chapter is based on: 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 *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (pp 194:1–194:14). [600] A modified version of the paper was published as a magazine article: Yixin Zou & Florian Schaub. 2019. Beyond mandatory: Making data breach notifications useful for consumers. *IEEE Security & Privacy Magazine*, 17(2), 67-72. [604]

incident description’s specificity. Consequences and risks of the 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 were 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.

## 5.1 Background

**The legal landscape of data breach notifications.** Data breach notification requirements vary widely across jurisdictions. For example, in the European Union, GDPR requires data breach notifications to both the supervisory authority within 72 hours and affected European consumers ‘without undue delay’ [144]. The notification must include a clear and plain description of the breach’s nature and recommended protective measures [144]. Substantial fines for non-compliance with GDPR [314] pose incentives for companies to disclose mandated information and establish more robust security incident procedures and training [420].

In the US, no equivalent federal data breach notification law exists [348]. Instead, a patchwork of sector-specific federal laws outlines various requirements. For example, GLBA [535] regulates data breach notifications for financial institutions, prescribing several mandatory elements to be included. HIPAA [512] establishes a 60-day notification deadline for data breaches that compromise consumers’ health information via a mailed letter written in plain language. In addition to these sectoral laws, all 50 US states have enacted their own data breach notification laws. These state laws vary substantially in stringency, resulting in inconsistent notification requirements among states and different definitions of PII [287], 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 set clear expectations about the structure and formatting of breach notifications in their law with a template [481]. The template includes specific wording of the title (“Notice of Data Breach”) and headings and 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 require companies to offer identity theft protection services to help consumers deal with potential harms; similar legislation is pending in the State of New York [361].

Our content analysis sourced breach notifications from the State of Maryland due to the comprehensive records in the Maryland Attorney General’s public database. Maryland’s Personal Information Protection Act [176] defines that PII encompasses traditional types of PII (e.g., name and Social Security number), government-issued IDs, and health information, among others. The law requires data breach notifications to be sent to individual consumers ‘as soon as possible’ and within 45 days upon discovery of a data breach [318]. In addition, the notification must specify the types of breached information and offer the contact information for different entities such as major credit bureaus, the FTC, and the Maryland Attorney General [318].

**Data Breach Notification Analysis.** There has been limited research on the content, language, and structure of consumer-facing data breach notifications. Jenkins et al. [244] found that visual elements (e.g., italics, bold, and underlining of subject lines), while rarely used (in less than 30% of their analyzed notifications), contributed to the restoration of the affected company’s reputation in a follow-up experiment.

Veltsos [548] found that the claim ‘lost data might not be used at all’ appeared in two out of 13 templates to ‘soften the bad news’, which, according to the author, does not help affected consumers overcome optimism bias [448] and rational ignorance [128]. Golla et al. [180] examined real-world notifications for password breaches and found that most notifications, while successfully raising participants’ concerns, did not lead to intentions to change compromised passwords or adopt other secure practices against 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 [48]. Bisogni’s work 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 [48]. Due to the identified inconsistencies, the paper concluded that a US federal breach notification law is needed to standardize notifications’ timing and mandatory elements, both of which are crucial for informing consumers of potential risks [48]. We expand on and complement Bisogni’s study by (1) analyzing more recent data breach notifications and (2) focusing on readability and usability issues.

## 5.2 Method

**Data collection.** Many US state laws require companies to submit data breach notifications sent to consumers to the state’s Attorney General’s office when the breach affects the state’s residents. To date, California, Iowa, Maryland, New Hampshire, and Vermont have made 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 broadest range of data breaches. It also provides additional metadata, such as the cause of the breach and the types of information compromised.

From Maryland’s database, we downloaded all data breach notifications on record

from January 1<sup>st</sup> to June 30<sup>th</sup> 2018 to narrow the scope of our analysis while ensuring data recency. In total, we obtained 548 data breach notifications. Next, we cleaned the dataset by removing duplicates and entries that did not include consumer-facing notifications (i.e., some files only had 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 ( $\approx 50\%$ ) from them to analyze. Appendix C.1 provides a complete list of our analyzed notifications.

**Sample.** The 161 notifications in our sample came from 159 unique companies. The majority (154; 96%) were letters; four were delivered via email, one company also sent in-app messages and push notifications to consumers using their app. Fourteen companies included multiple versions of notifications in their uploaded files. We analyzed the version with a larger audience (e.g., adults instead of guardians of affected minors, 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%), SSN (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 database [405] 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 PRC, breaches in our sample exposed 151.93 million records across the US, which constituted 18.5% of the total exposed records coming from all data breaches listed by PRC (820.93m) for this time frame. Of the 151.93m records, 150m were exposed in one breach (Under Armour); 9 breaches (16%) exposed over 10k records, and 46 (84%) breaches exposed over 100 records.



**Quantitative analysis.** Our quantitative analysis focuses on readability, which shows to what extent a particular breach notification is understandable by the general public. Two metrics we used, the Flesch Reading Ease Score (FRES) and the Flesch Grade Level (FGL) [157], are calculated based on the sentence length and word length of the text. We also looked into the Gunning Fog index (FOG) [193], 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 quantitative 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 [286]. Three research team members 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'  $\kappa=0.75$ ). The final codebook (see Appendix C.2) has nine 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.

### 5.3 Results

Our analysis shows that current data breach notifications were not readable. Most notifications followed a similar structure, yet their style, length, and content specificity varied considerably. Furthermore, companies downplayed the severity and con-

sequences resulting from a data breach. Even though all notifications recommended protective measures, there was little guidance on how consumers should prioritize among them.

### 5.3.1 Readability and Structure

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

**Advanced reading skills required.** Severe readability issues surfaced from the analysis. FRES [157] 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 5.1 shows our sample’s FRES distribution mapped onto Flesch’s 7-level ranking system [222] 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). These findings indicate that 97% of the notifications were fairly difficult or difficult to read. Conversely, only one notification was rated “easy” and one “standard”, and no notification was rated as “fairly easy” or “very easy.”

Converting FRES to FGL [157], our sample’s median FGL was 10.0 (mean: 10.02, sd: 1.18), i.e., the 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 an FGL score higher than 9.4. By contrast, prior literacy research suggests that materials addressed to the general public should aim for a junior-high reading level (i.e., 7 to 9) [222]. Using FOG [193], the result was even poorer, with a median of 11.6 (mean: 11.55, sd: 1.33), indicating that the existence of long words and jargon aggravated the readability of

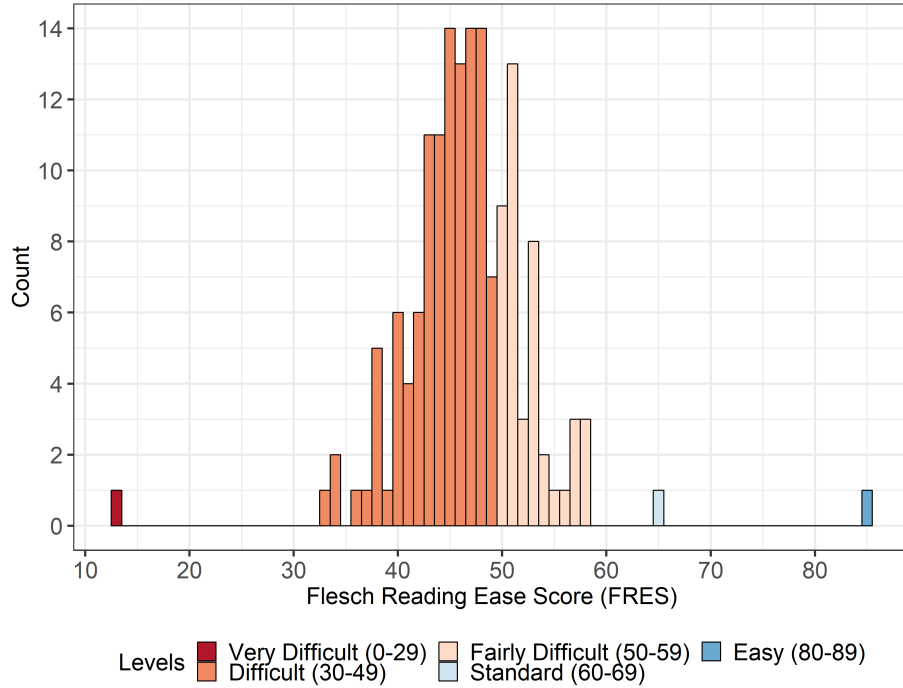


Figure 5.1: Distribution of FRES, mapped onto Flesch’s 7-level ranking system.

these notifications. Essentially, most analyzed notifications would likely not meet “plain language” requirements in California’s data breach law or GDPR—average readers will struggle to understand these breach notifications.

**Long 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 (mean: 1,539, sd: 644), ranging from 213 to 3,414 words (or seven 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 [332] 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 [71]. The estimated

time to skim a data breach notification ranged from 0.85 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 [299, 332], a 6-7 minute anticipated reading time, paired with the need for advanced reading skills, still creates a considerable burden for consumers.

**Structural headings are common and consistent.** Headings structure the text and guide a reader’s attention, helping them identify key information quickly in the long text. 106 (67%) notifications used headings to separate their main text into sections. Among them, 72 (68%) put the heading in a separate line; 34 (32%) included the heading in a paragraph’s first line, reducing its salience (see Figure 5.2). Only two 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 finding suggests that, unlike privacy policies, data breach notifications generally follow a consistent structure, thus facilitating the learning of that structure over time [366]. However, the disparities regarding heading formatting, paired with poor readability, indicate that inconsistency remains on the content level.

### 5.3.2 Risk Communication

Risk communication for data breaches requires companies to explain the situation clearly and openly acknowledge negative consequences [548]. Maryland law [176] 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.

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

**Varied specificity regarding cause and compromised data** As required by Maryland law, the types of information exposed were 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.”* In a few cases, the notification referred to the recipient’s own breached information, such as a credit card number’s last four digits. Such an individually tailored message provides clear evidence that the recipient was personally affected, which might alarm consumers and motivate them to take actions [344]. However, it may also pose identity theft risks if the notification falls into the wrong hands [583].

The cause of the breach was reported in 150 (93%) notifications. Example causes include unauthorized access (38), phishing (33), malware (19), inadvertent human error (16), and a few others. Eleven 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 happened 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 how long their data has been exposed before they become aware [48], which is an important metric to decide how urgently actions are needed. About two-thirds of the notifications (105, 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** Data breach notifications should clarify that the recipient’s information has been breached and describe associated risks [548]. 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.*” 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.*” As shown in previous work [417, 602], the use of hedge terms is detrimental to consumers’ ability to assess risks accurately and usually leads to confusion and misconception.

A positive example is as follows: “*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 that 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, e.g., “*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. The “no evidence” argument appeared in 64 (40%) notifications, e.g., “*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.*” In other cases, the hedge terms appeared in describing potential access and misuse of consumers’ information 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.”* While it might be factually accurate that companies do not have evidence of data access and misuse, lack of evidence does not support an 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 action.

On the contrary, nine notifications used an effective risk communication strategy—connecting the types of breach information with potential misuse scenarios. For instance, the company may highlight that multiple types of exposed data in combination can 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 aware that because of the Incident, there is a possibility of (1) identity theft, and (2) fraudulent filing of your tax information.”* The second point helps the recipient prioritize filing their tax return early.

### **5.3.3 Presentation of Recommended Actions**

Consumers are urged to take protective actions when a data breach occurs, and the required actions usually depend on the types of breached information. Options include enrolling in a provided protection service, placing credit freezes, changing compromised passwords, among others. Several states and federal agencies offer templates for describing available measures. These templates usually include definitions of fraud alert and credit freeze, the contact information of major credit bureaus and the state AG, and enrollment instructions of complimentary identity theft protection services if offered.



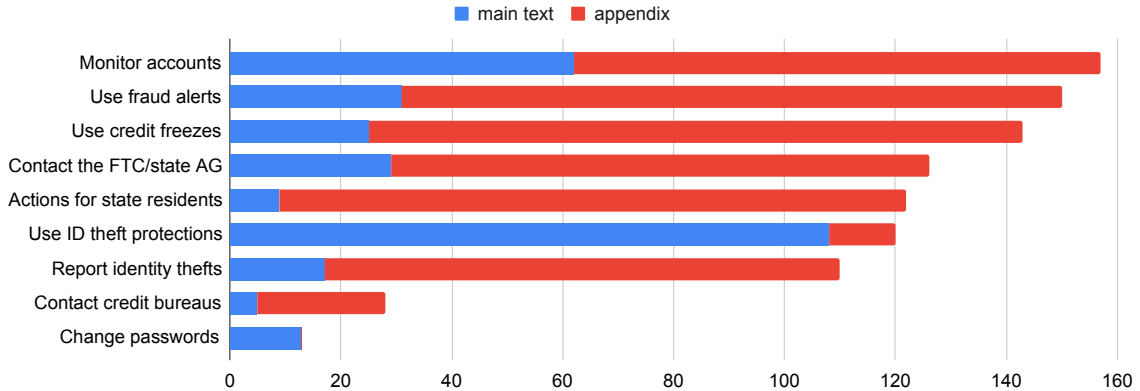


Figure 5.3: Frequency of recommended protective measures.

**Choice overload with no priorities.** Figure 5.3 shows the frequency distribution of commonly recommended actions. All notifications recommended at least one action, with a median of eight actions (mean: 7.19, sd: 2.24). 35 notifications (22%) recommended more than nine protective measures, and the most comprehensive included 16. The presence of so many options, compounded by lengthy explanations and poor readability, suggests the possibility of “choice overload” [270], meaning that the reader might delay the decision-making process, pick a random option under pressure, or even avoid all options.

We expected vital information in the main text and supplemental information in appendices. For each notification, we counted if there was an appendix; for each action, we coded whether it first appeared in the main text or as an appendix. Appendices were prevalent: 97 (60%) notifications used one appendix, 38 (24%) used two appendices, and one included three. 25 (16%) notifications had main text only, ending with the sender’s contact information. Often, effective actions were hidden in long texts in appendices with no indication of their high priority. A credit freeze, for example, is considered an effective measure to limit identity theft. 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 deserves a higher priority.

**Formatting focused on sub-level text.** Formatting, such as using lists and capitalizing important information, can highlight critical details and reduce the reader’s cognitive burden in processing text. Prior research suggests that formatting techniques in a data breach notification could enhance consumers’ perception of the affected company’s reputation [244].

We coded the presence of list and text formatting (e.g., bold, italicized, underlined, capitalized, or colored text) when presenting actions. Overall, lists were scarce in the top-level text (e.g., to describe different actions) but became more prevalent in sub-level text (e.g., to explain 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) than numbered lists (20 vs. 51). This finding indicates that while consumers are walked through details of each specific action, they may struggle to form a holistic view of the major actions due to the lack of list formatting on the top level.

Conversely, text formatting was common in both top-level (149, 93%) and sub-level text (90, 56%). However, text formatting was rare in describing important details of offered identity protection services, such as the enrollment deadline (after which the service is no longer free) and the duration of benefits (after which the protection is no longer valid). 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 lose track of the remaining time. When these crucial timings blend in with other 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.

## 5.4 Discussion

Our analysis contributes novel insights into the readability and usability of recent data breach notifications. Building on Bisogni’s study [48], we found that our sample’s notifications tended to include mandatory elements by Maryland law: over 90% specified types of compromised information, whereas only 65% reported the breach’s occurrence or discovery date. Consistent with Veltsos’s analysis [548], we observed ‘no evidence of data misuse’ claims and hedge terms as an additional strategy to downplay risks. Contrary to Jenkins et al.’s findings [244], formatting techniques were common in our sample, but substantial differences emerged between top-level and sub-level text, and crucial information was not highlighted.

Findings in Chapters III to V indicate that more efforts should go toward supporting consumers in protecting themselves rather than merely informing them of the breach. Next, we discuss the limitations of this study, before highlighting opportunities for improving the utility and usability of data breach notifications to make them an effective mechanism for helping consumers mitigate potential risks.

**Limitations.** Our study has certain limitations. We only analyzed breach notifications from Maryland Attorney General’s public database, which may differ from those sent to consumers in other states. From a cursory comparison with the databases of other states, we are confident that our findings are not Maryland-specific. 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.

Our content analysis does not provide direct evidence of how text and formatting variations in data breach notifications impact consumer behaviors. Nonetheless, given existing research on privacy policies [432, 540], 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 similarly contribute to consumer inaction. 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.

Lastly, the HCI contributions of a content analysis of notification letters may not be immediately apparent. 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. However, our findings' implications are relevant for general data breach notifications regardless of the delivery medium. Our recommendations can inform the design of more effective and actionable breach notifications in letter form and online contexts.

**Readability expectations beyond “plain language.”** GDPR and California's breach law both have the “plain language” requirement for breach notifications. However, according to our readability evaluation results, most of the breach notifications we analyzed failed to meet this requirement. A potential reason may be that these laws do not clearly define how to assess whether something is written in plain language. Regulators should provide more explicit guidelines on how this plain language requirement can be achieved, including recommended tips such as using short sentences, common words, and active voice. Furthermore, we can learn from practices in the insurance industry, where the FRES test is required as a readability assessment of insurance policies in some US states [222].

**Delivering notices through multiple channels.** Currently, most US state laws require written notices sent to affected consumers after a data breach—96% of our

analyzed notifications were mailed letters. Electronic notices (e.g., emails, website announcements, and notices to statewide media) are treated as substitutes when the cost of delivering mail is too expensive or the physical addresses of affected individuals are unavailable. However, the slow speed of mailed letters might increase the uninformed exposure time to potential risks for consumers [48]. It might also explain why many consumers learn about a data breach before receiving direct notifications from companies [1]. Conversely, electronic notices are faster and can provide consumers with direct links to actions, thus reducing barriers in moving from intention to action. The nature of electronic methods may also incentivize companies to shorten the text and increase aesthetics. That being said, electronic notices need to be compounded with precise readability requirements to prevent companies from sending lengthy and unreadable electronic notices.

**Consistent standards for style and format.** Even though our primary data source pertained to Maryland, most analyzed notifications with section headings adhered to wording required by California’s breach notification law [481]. This finding indicates a promising avenue for standardizing style and format expectations for data breach notifications. Legislators and regulators should provide specific content and style requirements, potentially templates with good readability and usability based on rigorous user testing. The provisions of California’s data breach notification law and the GLBA model privacy notice [498] demonstrate the reach and influence of official templates—on the premise that provided templates are usable and actionable.

**Using visual emphasis to enhance user experience.** Formatting makes information visually accessible and enhances the overall user experience. We suggest text formatting to highlight crucial information and consistent use of list formats to lay out major actions. When lists are used, each point should be followed by short and concise sentences instead of long paragraphs to keep the cognitive burden low

for readers. Furthermore, it is crucial to consider the needs of special groups [559], such as visually impaired people, which means the content should be displayed in a sufficiently large font size, be accessible to screen reader devices, and contain required metadata and text descriptions.

**Communicating risks clearly and concisely.** Risk communication is critical to data breach notifications since risk perception is the precursor for forming the intention to take action. Risk communication is also challenging. While companies should help consumers correctly assess risks and determine the necessity to take action, they also would like to avoid overstating risks—an act that would likely harm their business interests. Prior work on privacy nudges [4] has provided valuable insights for improving risk communication in data breach notifications. For instance, optimism bias could be addressed by removing hedge terms to clarify that the reader is personally affected by the breach. Loss aversion theory (i.e., people hate loss more than liking the equivalent gain) can be leveraged when framing the outcome of recommended actions by emphasizing the negative consequences of inaction. In Chapter III, we found that people with low socioeconomic status, due to their limited money or assets, may subscribe to an “I’ve got nothing to lose” attitude, lacking motivation to react [602]. This fallacy could be addressed by describing how the reader’s acquaintances or those they relate to have been affected, such as showing evidence of susceptibility to identity theft and scams. Essentially, companies should be transparent about whether the recipient has been affected and avoid “no evidence of data misuse” claims, or at least combine them with clear warnings of potential future misuse.

**Supporting consumers in prioritizing and executing actions.** Lengthy, confusing, and full-of-jargon descriptions of protective actions likely hamper consumers’ ability to act. When making recommendations, companies should identify and highlight those most relevant to the specific breach. Leveraging the anchoring effect [4],

actions of high priority should be listed first to receive the most attention from readers. Moreover, companies should clearly explain why a particular action is important. To deal with the choice overload problem, companies should tailor the recommended actions to the specific breach rather than blindly adopting a given template. For instance, in analyzing notifications to Maryland consumers, we often observed long lists of contact information for other state attorney general offices, which are unnecessary details – at least for Maryland residents – that should be removed.

In conclusion, research on privacy policies has identified their deficiencies in communicating privacy risks: most are written in lengthy paragraphs filled with jargon and ambiguity [417], leading readers to struggle with comprehending the content and forming accurate mental models [407]. Unfortunately, our research reveals that data breach notifications suffer from similar issues. Yet, we have a limited understanding of how these issues may impact consumers’ comprehension and reactions in a moment when they are most vulnerable. While data breaches are recognized as severe threats, the design of corresponding mandatory notifications has received little attention. Poor readability and actionability, compounded by ambiguous risk communication, are possible explanations for the “data breach fatigue.” We outline directions for more effective data breach notifications that can help consumers overcome hurdles in dealing with risk and take action to protect themselves. More research is needed to develop and validate best practices to guide consumers towards safety after a data breach.

## CHAPTER VI

# Examining the Adoption and Abandonment of Security, Privacy, and Identity Theft Protection Practices<sup>1</sup>

Chapters III to V have provided a bleak picture of consumers' struggles with reacting to data breaches and issues with breach notifications regarding readability, risk communication, and actionability. These findings beg the question: is it prohibitively difficult for individuals to take action in other security and privacy domains? After all, advice related to data breaches is only a small component of the plethora of advice about protecting oneself in navigating the online world, from choosing passwords and using two-factor authentication (2FA) to disabling third-party cookies and many more. Moreover, end-users have limited time and effort to spend on protective behaviors [218], and recommended behaviors might be unrealistic, contradictory, or economically irrational [416]. Therefore, it is critical to consider the cumulative ecosystem of security- and privacy-protective behaviors [415] and how users adopt or do not adopt them rather than view each behavior independently.

In this chapter, I document an online survey with 902 US Internet users, cov-

---

<sup>1</sup>This chapter is based on: Yixin Zou, Kevin Roundy, Acar Tamersoy, Saurabh Shintre, Johann Roturier, & Florian Schaub. 2020. Examining the Adoption and Abandonment of Security, Privacy, and Identity Theft Protection Practices. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (pp. 443:1–443:15). [603]



ering 30 expert-recommended security, privacy, and identity theft protection practices suggested by prior work [64, 242, 311, 507]. While prior work has investigated why users adopt or reject expert advice, most studies focused on security practices [146, 242, 411, 414]. A handful examined privacy practices in specific contexts (e.g., [2, 172]). Hardly any work has looked into the adoption of identity theft protection practices despite respective risks following a rising number of privacy scandals, data breaches, and financial fraud in recent years [279, 503]. Moreover, most prior work on advice adherence has focused on motivations and hurdles for initial advice adoption [146, 242]. Reasons for inconsistent implementation or abandonment of advice after initial adoption have not yet been examined systematically, despite potential risks generated from such behavior (e.g., data breach victims who do not re-freeze their credit reports after a loan application would still be at risk of identity theft). Aiming to provide a holistic understanding of how and why people adopt, partially adopt, or abandon expert advice, our study seeks to answer the following research questions:

**RQ1: Which security, privacy, and identity theft protection practices are commonly adopted fully, adopted partially, or abandoned?**

**RQ2: What are predictive factors for a practice’s level of adoption?**

**RQ3: Why are certain practices partially adopted or abandoned?**

We found that security practices received wider adoption than privacy and identity theft protection practices. Both manual practices (i.e., users need to remember to adhere to the practice) and automated practices (i.e., no user effort required after initial adoption) were more popular than practices requiring recurring user interaction. Participants’ gender, education, technical background, and prior negative

experience were correlated with their levels of adoption. Practices were abandoned when they were perceived as low-value, inconvenient, or when users overrode them with subjective judgment, such as ignoring phishing warnings. Notably, participants sometimes made exceptions to practices that should be adopted consistently to be effective. Based on our findings, we discuss how expert recommendations can be improved to better align with end-users' needs and encourage continuous and consistent adherence. We further identify opportunities for designing security, privacy, and identity theft protection tools to promote such adherence, especially when recurring user interaction is required.

## **6.1 Method**

We conducted an online survey with 902 participants in August 2019. The survey aimed to investigate (1) which security, privacy, and identity theft protection practices are adopted, partially adopted, abandoned, considered, or rejected by Internet users, (2) what factors influence levels of adoption, and (3) the reasons for partial adoption or abandonment. A survey allows us to quantitatively analyze adoption and abandonment differences between individual practices and domains, draw inferences between user behavior and related factors, and quantify the reasons behind partial adoption and abandonment at scale. This study was approved by the University of Michigan's Institutional Review Boards (IRB).

### **6.1.1 Taxonomy of Expert-Recommended Practices**

We conducted an extensive literature review to determine practices to include (see Table 6.1). Prior work mostly associates online safety with security measures [242], but privacy and identity theft risks are increasing, making it important to contrast and characterize user adherence to expert advice regarding these adjacent domains.

The chosen security practices ( $n=12$ ) were primarily based on Ion et al.'s 2015

**Practice** (Prefixed with **[Abbreviation]**)

---

- S1. **[2FA]** Opt-in to 2FA for online accounts
- S2. **[Antivirus]** Use antivirus software
- S3. **[Attachment-clicking]** Beware of attachments sent by unknown people
- S4. **[Automatic-update]** Keep automatic software updates turned on
- S5. **[Check-URL]** Check the URL when visiting a website
- S6. **[HTTPS]** Check if the website visited uses HTTPS
- S7. **[Install-software]** Only install software from trusted sources
- S8. **[Link-clicking]** Avoid clicking links sent by unknown people
- S9. **[Password-manager, Assisted]** Use a password manager
- S10. **[Strong-password]** Use strong passwords for online accounts
- S11. **[Unique-password]** Use different passwords for each account
- S12. **[Update-software]** Install OS and software updates immediately
- P1. **[Anonymity-system]** Use anonymity systems, such as Tor and VPN
- P2. **[Cookies-clean]** Clear web browser cookies and history
- P3. **[Cookies-disable]** Disable or turn off third-party browser cookies
- P4. **[Encryption]** Encrypt phone calls, text messages or emails
- P5. **[Extension]** Use browser extensions that block ads, scripts or tracking
- P6. **[Hide-info]** Refuse to provide info that is not essential to transactions
- P7. **[Incognito]** Use private browsing mode
- P8. **[Public-comp]** Use a public computer to browse anonymously
- P9. **[Real-name]** Avoid using websites that ask for real names
- P10. **[Search-engine]** Use search engines that do not track search history
- P11. **[Temporary-credential]** Use fake identities for online activities
- P12. **[Facial-recognition]** Opt out of facial recognition when possible
- I1. **[Credit-freeze]** Place a credit freeze
- I2. **[Credit-monitoring]** Use a credit monitoring service
- I3. **[Credit-report]** Obtain free copies of credit reports
- I4. **[Fraud-alert]** Place a fraud alert \*
- I5. **[Identity-monitoring]** Use an identity monitoring service
- I6. **[Statements]** Check for fraudulent charges on account statements

Table 6.1: Security, privacy, and identity theft protection practices included in our study.

study on security advice [242]. They surveyed >200 experts (5+ years of computer security work experience) about their top three pieces of online security advice for non-tech-savvy users. Most expert advice remained constant in Busse et al.’s 2019 replication study [64]. We studied the 11 most-mentioned practices (of 152 total) in our survey, as they are likely to be agreed on by most experts [416]. Following the authors’ recommendation [416], we replaced one of those practices (“be careful/think before you click”) with two (“don’t click links in email from unknown sender” and “check URL for expected site”), resulting in 12 security practices in total.

Because no comparable systematic elicitation of expert advice existed for privacy and identity theft protection practices, we broadened our search to online articles, reports, and blog posts by experts from industry, government, and NGOs. Our chosen privacy practices ( $n=12$ ) were primarily based on a census-representative 2015 Pew survey examining Americans’ attitudes and behaviors about privacy [311], which asked whether respondents had engaged in any of 13 privacy-enhancing practices. We included all but two practices (“delete/edit something posted in the past” and “ask someone to remove something posted about you”), for which consistent and frequent full adoption might not be applicable or practical. We added “opting out of facial recognition” to unpack users’ behaviors given its substantial privacy implications [77, 470].

Our chosen identity theft protection practices ( $n=6$ ) came from the US Federal Trade Commission [507]. We included practices focused on identity theft protection and excluded more general security/privacy practices (e.g., “don’t overshare on social networking sites”) and practices that only apply to victimized individuals (e.g., identity recovery services). Some practices like credit freeze (restricting access to one’s credit report at a credit bureau) are only available to US consumers. As such, we only recruited US participants to control for cultural differences.

In developing our practice taxonomy, we noticed that practices varied in the level

of required user involvement, which may explain differences in adoption and abandonment. *Manual* practices require users to implement it on their own consciously (e.g., avoiding clicking links sent by unknown people)—the success of the practice solely relies on users’ manual application and cognitive assessment. *Automatic* practices require adopting a particular tool or service that, after initial setup, provides automatic protection with minimal user involvement (e.g., using an ad-blocking extension). *Assisted* practices like 2FA require the adoption of a tool or service, but users also need to interact with them recurrently for full protection.

### 6.1.2 Survey Protocol and Recruitment

We conducted our study on Prolific, a crowdsourcing platform similar to Amazon’s Mechanical Turk but provides more demographically diverse participants [382, 391]. We described the topic as “risk management when using the Internet” to reduce self-selection bias by avoiding priming about security, privacy, or identity theft. We recruited US participants who were 18 years or older with a >90% approval rate. Participants were compensated \$1.20 for work that took 5-10 minutes (mean: 9.68, median: 7.34), in line with Prolific’s required minimum hourly pay.

Upon accepting the task, participants were directed to our survey hosted on Qualtrics. After agreeing to the consent form, we showed each participant ten practices (four security, four privacy, two identity theft) randomly selected from our list of 30 expert-recommended practices to minimize respondent fatigue. The practices were displayed in randomized order. An attention check question was randomly placed among the ten practices.

We used the question format “Have you ever...?” for all practices. We provided definitions of terms, tools, or services involved for practices that might not be immediately understandable to the general public. We also provided screenshots of relevant UI elements for some practices to reduce the chances of misconception and confusion

<i>Full adoption</i>	I am ALWAYS doing this.
<i>Partial adoption</i>	I am doing this but there are exceptions. Please describe it further: [text-entry box]
<i>Abandonment</i>	I am NOT doing this anymore, but I have done this before. Please describe it further: [text-entry box]
<i>Consideration</i>	I have NEVER done this before, but I EXPECT to do this in the near future.
<i>Rejection</i>	I have NEVER done this before, and I DO NOT EXPECT to do this in the near future.
<i>Unawareness</i>	I have NEVER heard of this/I do not understand.
<i>Other</i>	Other (please specify): [text-entry box]

Table 6.2: Response options relating to adoption for our survey questions.

(denoted by \* in Table 6.1). For each practice, we asked participants if they have *fully adopted*, *partially adopted*, *abandoned*, *considered*, *rejected*, *not understood* the given practice, or something else (*other*). See Table 6.2 for the full response option texts. For four practices, we further clarified response choices to help participants distinguish between full and partial adoption (e.g., defining “full adoption” as “making multiple requests throughout the year” for obtaining free credit reports), or when partial adoption did not apply to the practice (e.g., one either signs up for credit monitoring service or not).

After going through the ten practices, participants were asked about prior experiences with unauthorized account access, data breaches, and identity theft. The survey concluded with demographic questions about age, gender, income, education, employment, and background in computer science (CS)/information technology (IT), and security/privacy. A “prefer not to answer” choice was offered for potentially sensitive topics. The full survey is included in Appendix D.1.

Table 6.3 compares our sample to the US population demographics. Our participants are evenly distributed between men and women but are more educated and skew younger. Their income levels cover a wide range, but fewer participants live in a household with more than \$100k annual income. Of our participants, 66.6% had no background in CS/IT or security/privacy; 11.6% only in CS/IT, 8.0% only in security/privacy, and 11.0% in both. Furthermore, 67.0% have been victims of

<b>Metric</b>	<b>Sample</b>	<b>Census</b>
Women, Men, Non-binary	50.0%, 48.0%, 1.8%	51.0%, 49.0%, N/A
High school, Some college	10.9%, 25.9%	28.6%, 19.0%
Trade/vocational, Associate	2.9%, 10.4%	4.1%, 5.5%
Bachelor’s, Master’s	34.4%, 11.7%	20.6%, 8.5%
Doctoral, Professional	1.3%, 1.3%	1.8%, 1.3%
18-24, 25-34 years	22.3%, 29.6%	9.3%, 14.0%
35-44, 45-54 years	22.6%, 13.6%	12.6%, 12.7%
55-64, 65-74 years	7.9%, 3.6%	12.9%, 9.3%
75 years or older	<1%	6.7%
<\$20k	16.5%	[10.2%, 19.1%]
\$20k-\$35k	17.2%	[8.8%, 17.7%]
\$35k-\$50k, \$50k-\$75k	15.3%, 21.6%	12.0%, 17.2%
\$75k-\$100k, >\$100k	12.2%, 14.5%	12.5%, 30.4%

Table 6.3: Gender, education, age and income demographics of survey participants. Census statistics from [527, 528, 529, 530] as of 2018.

a data breach, 35.0% have been victims of unauthorized account access, and 11.3% have been victims of identity theft.

### 6.1.3 Data Analysis

After removing 17 participants who failed the attention check question, we received 902 complete survey responses. The sample size followed the rule of thumb for linear mixed-effect models – at least 1,600 observations per condition in designs with repeated measures [58].

**Qualitative analysis.** Participants provided 1,728 open-ended responses in total. Among these, 69% were explanations for partial adoption, 25% for abandonment, and 6% for others. We developed a codebook to analyze reasons for partial adoption and abandonment. The first author read all responses and developed codes using inductive coding [286]. Two co-authors then independently analyzed 150 (8.7%) randomly sampled responses, reconciling codes and revising the codebook iteratively until reaching high inter-coder reliability (Cohen’s  $\kappa=0.82$ ). The two co-authors then split the dataset and single-coded all responses. Our codebook is included in Appendix D.2.

In going through participants’ open-ended responses about partial adoption and abandonment, we realized some responses pointed to other options on the list. For instance, one participant selected “other” for placing a fraud alert and said, “I have heard of this, but I have never done it before. It’s possible I could do it in the future,” which was an exact match for *consideration*. Two authors re-coded these responses to minimize report biases and inconsistencies in the data. In total, 171 responses were re-coded, of which 75 were originally *abandonment*, 71 were *other*, and 25 were *partial adoption*.

**Statistical analysis.** Using the re-coded dataset, we calculated descriptive statistics for each practice’s rates of full adoption, partial adoption, abandonment, etc. Motivated by prior work suggesting the influence of user characteristics and tool usability issues on user behavior, we constructed mixed-effect regression models. For fixed-effect factors, we included user characteristics (i.e., demographics, technical background, prior negative experience) and practice characteristics (domain and nature of protection). We further included random effects resulting from differences between individual participants and practices when fixed-effect factors are under control. The intraclass correlation coefficient (ICC) [337] for all models are below 0.20, indicating that random differences between individual participants or practices contributed little to variances in the adoption level.

To understand what factors influence users’ current levels of adoption, we performed linear regressions on an adjusted scale of response options, from 0 as no adoption (combining *abandonment*, *consideration*, and *rejection*), 1 as *partial adoption*, and 2 as *full adoption*, excluding rare cases of *unawareness* or *other*. Since the response options are only quasi-linear, we also ran ordinal logistic regressions to validate linear regression results, which produced the same findings with only minor variations in numeric outputs of effect size. Thus, we report linear regression results



since they are more informative.

To know how the effects of different predictors vary across domains, we ran a series of models with interaction terms between the practice domain and another predictor (e.g., interaction terms between the practice domain and gender show how gender effects on adoption vary across domains). Post-hoc power analyses suggest that our study was sufficiently powered. Based on 1k simulations of the likelihood ratio test, the power to detect the overall effect of the domain variable on adoption is 97.20%, CI (95.98%, 98.13%). For demographic predictors, we binned the demographic data in Table 6.3 into fewer categories in the regression analyses to avoid model fit problems caused by too few observations in a category: gender (men, women, binary), age (three levels: 18-34, 35-54, 55+), educational attainment (three levels: no Bachelor’s degree, Bachelor’s degree, Graduate degree), and annual household income (three levels: <\$50k, \$50-100k, >\$100k).

To understand what factors influence a practice being abandoned, we tried running logistic regressions on a binary variable with “yes” meaning *abandonment*, and “no” meaning *partial adoption* or *full adoption*, excluding other response options. However, due to the small number of abandonment cases in our dataset (534 “yes” vs. 5,325 “no”) the model expectedly failed to converge. Similarly, multinomial logistic regressions on the full spectrum of response options failed to converge because response options like “unaware” and “other” were much less frequent than others. Therefore, our regression analysis only focuses on adoption; we refrain from making statements about which variables are correlated with abandonment, consideration, or other response options.

#### **6.1.4 Limitations**

While our scope of investigated practices exceeds most prior work, there might be other relevant practices related to security, privacy, identity theft protection, or other online safety topics, such as harassment and cyberbullying [410]. As with any survey,

participants may over-report their behavior due to social desirability bias [154], which might be particularly salient for *full adoption* when participants think they consistently implement a practice while forgetting exceptions they make. To mitigate this bias, we provided instructions to encourage honest answers and guarantee responses would be anonymized. Our main goal was not to provide empirical field measurements about actual user behavior, but rather to understand, in participants' own opinion, what practices they think they fully adopt and what practices are adopted only in certain situations.

Another point concerning consistency is the removal of partial adoption as a response option for credit monitoring, identity monitoring, credit freeze, and fraud alert. While this act makes the results of partial adoption for identity theft protection practices less comparable to those for security or privacy protection, we considered it an important measure to reduce confusion, as partial adoption does not apply to these practices.

## 6.2 Results

Below we discuss the most adopted and abandoned practices, before presenting predictors of and reasons behind adoption and abandonment behavior.

### 6.2.1 RQ1: Commonly Adopted and Abandoned Practices

Figure 6.1 shows the percentage distribution of response options for each practice. Overall, security practices had the highest full adoption rates, while partial adoption and abandonment mostly occurred to privacy practices. Most identity theft mitigation practices had low adoption rates, and many participants reported they would not consider them.




Figure 6.1: Distribution of response options for each practice. The practices are sorted by full adoption rates in descending order.

**High adherence to security practices.** Of the ten practices with the highest full adoption rate, seven are security practices, with the top two reflecting the importance of cautious clicking behavior (94.6% for links, 92.6% for attachments). Two privacy practices were also fully adopted at high rates, namely hiding information that is not essential to transactions (69.0%) and using a privacy-enhancing browser extension (68.0%). Checking account statements was the only identity theft protection practice fully adopted by over half of our participants (76.2%). Except for antivirus software and privacy extensions, these commonly fully adopted practices are manual, situated in people’s everyday digital experiences, and not overly technical.

**Partial adoption exists in both security and privacy practices.** As we did not provide partial adoption as a response option for four out of six identity theft protection practices, we only report partial adoption results for security and privacy practices. Overall, practices with high partial adoption rates were evenly split between security and privacy, with the top three being privacy risk management (49.7% for cleaning cookies, 39.9% for incognito, 39.1% for avoiding websites asking for real names). Consistent with prior work [388, 483], a substantial proportion of participants did not fully follow expert-recommended password management practices (29.4% for unique passwords, 21.4% for strong passwords).

**Abandonment mostly occurred for privacy practices.** Abandonment rates were below 20% for all practices and less common than full or partial adoption overall. Seven of the ten practices with the highest abandonment rates were privacy practices, with using an anonymity system being the most commonly abandoned practice (16.8%). Using automatic updates for software (13.3%) and antivirus (11.0%) were the most abandoned security practices. 7.2% had abandoned using a credit monitoring service. Some abandonment decisions appear rational since they seem more realistic for one-time use than long-term implementation (e.g., using a public computer for

anonymous browsing). Yet, some other abandoned practices, such as cleaning cookies, require consistent implementation for adequate protection.

**Low acceptance of identity theft protection practices.** Among practices that most participants had not yet adopted, many pertain to identity theft risk mitigation. The top practices considered for future implementation were opting out of facial recognition (29.2%), using an identity monitoring service (28.9%), and placing a fraud alert (24.6%). Most of these require adopting automated tools or services (e.g., credit/identity monitoring) or require frequent user interaction (e.g., password managers). Nonetheless, automated practices like credit freeze and fraud alert are among the top rejected practices (54.5% and 51.2%, respectively). This finding is concerning given that 66% of our participants reported being data breach victims and that these practices are among the most commonly recommended measures in data breach notifications [600].

### 6.2.2 RQ2: Factors Affecting Levels of Adoption

Our mixed-effect linear regression models show that different levels of adoption are related to the practice's domain and its type of user interaction. We further found significant effects on adoption from demographics, technical background, and prior negative experiences. Results of the main regression model are shown in Table 6.4.

**Levels of adoption: security > privacy ≥ identity theft.** Confirming the descriptive analysis, the adoption level of security practices was significantly higher than those for privacy practices ( $b=0.57$ ,  $CI=[0.20, 0.94]$ ,  $p<.01$ ) or identity theft protection practices ( $b=0.62$ ,  $CI=[0.31, 0.94]$ ,  $p<.001$ ). While privacy practices exhibit higher levels of adoption than identity theft protection practices, the difference is not significant ( $b=0.20$ ,  $CI=[-0.11, 0.50]$ ,  $p=.21$ ).

Category	Variable	<i>b</i>	CI	Original p-value	Adjusted p-value
Age	18-34	-	-	-	-
	35-54	0.00006	[-0.05, 0.05]	.99	.99
	>55	-0.04	[-0.12, 0.03]	.26	.33
Gender	Women	-	-	-	-
	Men	0.08	[0.04, 0.13]	<.01 (**)	<.01 (**)
	Non-binary	0.08	[-0.09, 0.25]	.36	.43
Income	<\$50,000	-	-	-	-
	\$50,000-\$100,000	-0.03	[-0.08, 0.02]	.23	.33
	>\$100,000	-0.06	[-0.13, 0.004]	.07	.13
Education	No Bachelor's degree	-	-	-	-
	Bachelor's degree	0.05	[0.002, 0.10]	<.05 (*)	.09
	Graduate degree	0.02	[-0.06, 0.09]	.66	.74
Tech background	Neither IT nor S&P	-	-	-	-
	Only IT	0.09	[0.02, 0.16]	<.05 (*)	<.05 (*)
	Only S&P	0.07	[-0.01, 0.15]	.09	.16
	Both IT and S&P	0.15	[0.08, 0.22]	<.001 (***)	<.001 (***)
Prior experience	Unauth. access: No	-	-	-	-
	Unauth. access: Yes	-0.01	[-0.06, 0.04]	.70	.74
	Data breach: No	-	-	-	-
	Data breach: Yes	0.05	[0.002, 0.10]	<.05 (*)	.09
	Identity theft: Yes	0.16	[0.10, 0.23]	<.001 (***)	<.001 (***)
Privacy domain	Identity	-	-	-	-
	Privacy	0.20	[-.11, .50]	.21	.33
	Security	0.62	[0.31, 0.94]	<.01 (**)	<.01 (**)
Practice nature	Assisted	-	-	-	-
	Automatic	0.53	[0.21, 0.85]	<.01 (**)	<.01 (**)
	Manual	0.64	[0.37, 0.90]	<.001 (***)	<.001 (***)

“-” means the variable is set as the baseline in the model. Comparisons between any pairs of non-baseline variables in this table were also made, and results are reported in text. The regression coefficient (*b*) shows to what extent the variable, compared to the baseline, brings the outcome (level of adoption) up or down on a scale from 0 to 2. CI is the 95% confidence interval. Statistically significant factors (adjusted  $p < .05$  after applying the Bonferroni-Holm correction) are denoted with \*.

Table 6.4: Results of the main regression model, excluding interaction terms between the practice domain and other variables.

**Low adoption for assisted practices.** We were interested in whether a practice’s degree of user interaction (manual, assisted, automated) affects adoption. We expected practices relying on manual effort to be least often adopted due to higher cognitive demand leading to errors or inconsistent behavior. Our results show the opposite: assisted practices, which require recurring user interaction, were adopted the least, with manual ( $b=0.64$ ,  $CI=[0.37, 0.90]$ ,  $p<.001$ ) and automated ( $b=0.53$ ,  $CI=[0.21, 0.85]$ ,  $p<.01$ ) practices exhibiting significantly higher levels of adoption. The difference between manual and assisted practices was particularly salient for identity theft practices ( $b=1.02$ ,  $CI=[0.40, 1.64]$ ,  $p<.001$ ), but such a significant difference does not persist for security or privacy practices alone.

**Gender and age differences.** We further identify significant effects on adoption levels from certain user characteristics. Men had significantly higher levels of practice adoption compared to women ( $b=0.08$ ,  $CI=[0.04, 0.13]$ ,  $p<.001$ ). Such gender difference applies to security and privacy practices in particular ( $b=.11$ ,  $CI=[0.04, 0.17]$ ,  $p<.01$  for both). This finding confirms prior work showing similar gender differences for phishing susceptibility [200, 454] and extends it to a wider range of practices.

Mapping age to the following categories: 18-34, 35-54, and 55+, we find significant age effects for security and privacy practices, but not overall. Middle-aged participants (35-54) adopted more security practices than younger participants ( $b=0.07$ ,  $CI=[0.01, 0.14]$ ,  $p<.05$ ). This finding aligns with prior finding [331] that older people demonstrate higher information security awareness. The opposite trend emerged for privacy practices, for which younger participants had significantly higher levels of adoption than middle-aged ( $b=0.14$ ,  $CI=[0.07, 0.20]$ ,  $p<.001$ ) and older participants ( $b=0.23$ ,  $CI=[0.13, 0.33]$ ,  $p<.001$ ). As prior work has shown that young adults are more likely to engage in privacy-protective behaviors on Facebook [266], our finding extends this age difference to other privacy practices.

**Higher adoption among lower-income participants.** Mapping household income to the categories <\$50k, \$50-100k, and >\$100k, we find the overall trend that participants with lower incomes exhibit higher levels of practice adoption, though there were no significant differences between any two groups. When looking at individual domains, those earning <\$50k had significantly higher levels of privacy practice adoption than those earning >\$100k ( $b=0.13$ ,  $CI=[0.04, 0.22]$ ,  $p<.01$ ). This finding may seem counter-intuitive, as higher-income people should have a stronger motivation and more resources to protect their privacy and assets. Nonetheless, prior work has found that people with lower incomes have heightened informational and physical privacy and security concerns [310], which might translate into adopting protective practices that are accessible and affordable to them.

**More education contributed to higher adoption.** Mapping educational background to the categories less than Bachelor's degree, Bachelor's degree or equivalent, and graduate degree, we find that more educated participants exhibited higher levels of practice adoption overall. In particular, participants with a Bachelor's degree ( $b=0.24$ ,  $CI=[0.15, 0.33]$ ,  $p<.001$ ) or a graduate degree ( $b=0.34$ ,  $CI=[0.21, 0.45]$ ,  $p<.001$ ) had significantly higher adoption of identity protection practices than those without. Compared to prior findings that more educated people tend to take less security precaution [565], our work suggests that the trend might be different for mitigating identity theft risks.

11% of participants reported a background in both CS/IT and security/privacy, and could therefore be considered experts. Their levels of practice adoption were significantly higher than those of the 67% who had no background in either field. Interestingly, this difference between experts and non-experts holds when considering CS/IT only, but not for participants who reported a background only in security/privacy (not CS/IT). This finding suggests that technology experience and



<b>Partial Adoption</b>	<b>Count</b>	<b>Abandonment</b>	<b>Count</b>
site-specific	179 (15%)	not-needed	68 (20%)
only-sensitive	129 (11%)	because-of-risk	50 (14%)
impractical	124 (10%)	impractical	41 (12%)
own-judgment-sufficient	111 (9%)	usage-interference	23 (7%)
because-of-risk	95 (8%)	own-judgment-sufficient	21 (6%)
usage-interference	80 (7%)	using-substitute	21 (6%)
only-finance	74 (6%)	platform-specific	17 (5%)

Table 6.5: Top coded reasons for partial adoption and abandonment.

expertise might significantly influence practice adoption more than security/privacy knowledge alone, which, as our participants reported in open-ended responses, was mainly based on university courses or employer-mandated training.

**Experiencing identity theft contributes to higher adoption.** Overall, participants with prior experience with identity theft incidents adopted more protection practices. This trend also holds true when looking at security, privacy, or identity theft practices individually, suggesting it is a robust trigger for pro-safety behaviors ( $b=0.14$ ,  $CI=[0.05, 0.23]$ ,  $p<.01$  for security;  $b=0.11$ ,  $CI=[0.01, 0.20]$ ,  $p<.05$  for privacy;  $b=0.33$ ,  $CI=[0.21, 0.45]$ ,  $p<.001$  for identity). Experience with being a victim of data breaches is also correlated with higher levels of adoption, though the effect is non-significant. By contrast, experience with unauthorized access to online accounts has little impact on adoption levels.

### 6.2.3 RQ3: Reasons for Partial Adoption and Abandonment

Participants were asked to provide explanations when indicating partial adoption or abandonment of a practice. The most prevalent reasons for each are shown in Table 6.5. Tables 6.6 to 6.8 provide the top three partial adoption and abandonment reasons for individual practices. To provide more informative results, we do not report reasons coded as “unclear” (i.e., unintelligible or irrelevant) or reasons that only describe adoption “frequency” (e.g., “*I do this sometimes*”).

Security Practice	<i>n</i>	Top Three Reasons for Partial Adoption	<i>n</i>	Top Three Reasons for Abandonment
Update-software	95	usage-interference (23), impractical (22), own-judgment-sufficient (11)	9	usage-interference (4), impractical (3), performance-issues (2)
Unique-password	93	impractical (25), site-specific (17), because-of-risk (14)	2	because-of-risk (1), forgetting (1)
2FA	68	only-finance (20), only-sensitive (16), impractical (6)	10	impractical (6), distrust-service (1), not-needed (1)
HTTPS	62	only-sensitive (23), only-finance (15), forgetting (11)	3	not-needed (1), practice-by-default (1), using-substitute (1)
Strong-password	59	because-of-risk (11), impractical (10), only-finance (10)	3	not-needed (1), only-required (1), site-specific (1)
Install-software	51	own-judgment-sufficient (20), usage-interference (13), impractical (10)	4	impractical (2), own-judgment-sufficient (1), using-substitute (1)
Check-URL	44	using-substitute (11), forgetting (8), only-suspicious (8)	6	only-suspicious (3), because-of-risk (1), unrelated-reason (1)
Automatic-update	37	own-judgment-sufficient (10), platform-specific (8), site-specific (5)	37	own-judgment-sufficient (13), impractical (9), usage-interference (8)
Password-manager	24	site-specific (7), using-substitute (6), impractical (3)	7	platform-specific (3), distrust-service (1), impractical (1)
Antivirus	12	platform-specific (3), because-of-risk (2), only-required (2)	30	platform-specific (12), own-judgment-sufficient (4), distrust-service (3)
Link-clicking	8	only-suspicious (2), own-judgment-sufficient (2), using-substitute (2)	1	impractical
Attachm.-clicking	6	own-judgment-sufficient (5), impractical (1)	1	impractical

Table 6.6: Participants’ most frequent reasons for incomplete adoption and abandonment of security practices.

Priv. Practice	<i>n</i>	Top Three Reasons for Partial Adoption	<i>n</i>	Top Three Reasons for Abandonment
Real-name	116	site-specific (57), own-judgment-sufficient (30), using-substitute (14)	6	not-needed (2), own-judgment-sufficient (1), unapplicable (1)
Incognito	110	only-sensitive (47), impractical (11), site-specific (11)	20	not-needed (5), using-substitute (4), account-or-device-sharing (2)
Cookies-clean	86	unrelated-reason (36), forgetting (14), impractical (12)	13	impractical (4), forgetting (2), not-needed (2)
Temp.-credential	71	site-specific (31), because-of-risk (11), only-suspicious (11)	22	because-of-risk (9), only-suspicious (4), not-needed (3)
Search-engine	45	only-sensitive (8), own-judgment-sufficient (7), because-of-risk (5)	14	not-needed (9), impractical (2), unrelated-reason (2)
Cookies-disable	37	usage-interference (12), site-specific (6), own-judgment-sufficient (5)	12	using-substitute (3), impractical (2), unrelated-reason (2)
Extension	32	usage-interference (17), site-specific (6), own-judgment-sufficient (5)	11	usage-interference (5), not-needed (4), performance-issues (1)
Anon.-system	30	because-of-risk (9), only-sensitive (8), usage-interference (4)	42	not-needed (13), only-blocking (10), only-sensitive (4)
Hide-info	27	own-judgment-sufficient (9), site-specific (6), impractical (4)	1	site-specific
Public-comp	17	not-needed (4), distrust-service (3), using-substitute (3)	18	not-needed (10), unrelated-reason (4), because-of-risk (2)
Encryption	10	only-sensitive (4), as-needed (2), platform-specific (2)	7	because-of-risk (2), only-required (2), when-offered-free (2)
Facial-recog.	7	only-social-media (3), platform-specific (2), forgetting (1)	1	unapplicable

Table 6.7: Participants’ most frequent reasons for incomplete adoption and abandonment of privacy practices.

<b>Id. Protection Practice</b>	<b><i>n</i></b>	<b>Top Three Reasons for Partial Adoption</b>	<b><i>n</i></b>	<b>Top Three Reasons for Abandonment</b>
Statements	28	because-of-risk (21), using-substitute (4), not-needed (2)	5	unapplicable (3), not-needed (2)
Credit-report	14	because-of-risk (6), unrelated-reason (5), when-offered-free (2)	12	not-needed (6), using-substitute (2), because-of-risk (1)
Id.-monitoring	N/A		10	when-offered-free (3), because-of-risk (2), using-substitute (2)
Credit-monitoring	N/A		12	not-needed (4), when-offered-free (4), because-of-risk (1)
Fraud-alert	N/A		14	because-of-risk (11), impractical (1), unapplicable (1)
Credit-freeze	N/A		15	because-of-risk (14), usage-interference (1)

Table 6.8: Participants’ most frequent reasons for incomplete adoption and abandonment of identity protection practices.

**Reasons for partial adoption.** As shown in Table 6.5, 179 participants (15%) who selected “partial adoption” described selectively using the practice for specific sites, apps, accounts, or software (coded as “site-specific”). This was the most common reason for privacy practices like avoiding websites that ask for real names (57 participants) and using temporary credentials for online activities (31). Unfortunately, most participants did not specify where they applied the practice selectively. For those who did, 129 (11%) did so for sensitive sites (“only-sensitive”), 74 (6%) for finance-related sites (“only-finance”), 41 (3%) for suspicious or odd sites (“only-suspicious”), 15 (1%) for social media services (“only-social-media”), and 5 (<1%) for gaming services (“only-gaming”).

For practices adopted when visiting sensitive sites, 47 participants reported using incognito mode to interact with sensitive websites such as adult sites and the dark web, in line with prior work [195]. Other privacy practices were also adopted for this reason, though less frequently, such as using an anonymity system like VPN (8) or a search engine that does not track search history (8). In addition, some participants mentioned taking extra precautions for finance-related sensitive information: using 2FA (20), checking for HTTPS (15), and using unique passwords (10).

Another prominent reason for partial adoption cited by 124 participants (10%) was the practice being inconvenient or unusable, resulting in difficulty for consistent

adherence (“impracticality”). Many security practices were described as causing inconvenience, including 2FA (“*very annoying*”), updating software immediately (“*if I am in the middle of something I will not [do it]*”), and using unique passwords (“*it’s hard to keep track*”). Inconvenience extended to privacy practices, including cleaning cookies (12, e.g., “*it kills all my passwords*”) and using incognito mode (11, e.g., “*I like to be able to have a list of the places I visited if I need to go back*”). A few participants mentioned the practice was too hard to follow consistently. They referred to “*rare occasions where I slip up*” despite best intentions. Such failures might be more common in real life than reflected in our self-reports due to social desirability bias and difficulties in recognizing when mistakes have been made.

111 participants (9%) reported relying on their judgment to determine when it is safe to depart from best practices (“own-judgment-sufficient”). For security practices, this usually means installing software from suspicious sources (20), disabling automatic updates for software at times (10), and clicking unknown attachments (5). For example, in talking about clicking attachments, one participant said: “*I don’t click on obvious spam emails, but I am willing to open emails that seem legitimate even if I don’t know the senders.*” However, even trained individuals routinely fall for phishing emails [454]. Similar trends manifested for privacy practices, with nine participants disclosing non-essential information when they trusted the service, e.g., “*I do play this by ear depending on the website and my familiarity with it.*”

95 participants (8%) reported adopting practices only when motivated by a perceived risk (“because-of-risk”), particularly for identity protection practices, such as checking account statements for fraudulent charges (21) and obtaining credit reports (6). The at-risk feeling also motivates using strong passwords (11) and anonymity systems (9). For identity theft protection practices, adoption typically occurred after a data breach, a lost credit card, or when suspicious activity appears on a bank/credit statement. Security risks revolved mostly around account hacking due to weak pass-

words. The most common privacy practice in this category was using a VPN when “*connected to untrustworthy or unsafe networks.*”

Finally, 80 participants (7%) reported struggling with practices that broke existing functionality or disrupted regular use of the device or service (“usage-interference”), such as updating software (23), using privacy-enhancing browser extensions (17), and disabling third-party cookies (12). Participants selectively abandoned updates when buggy updates had “*broken drivers, programs, or the OS itself,*” whitelisted sites on which browser extensions “*blocked things I didn’t want it to block,*” and allowed cookies when needed for the functionality of a site.

**Reasons for abandonment.** Top reasons for abandonment are summarized in Table 6.5. We primarily discuss cases in which abandonment reasons differ from partial adoption justifications.

The most common reason for abandonment, cited by 68 participants (20%) was that they did not need the practice anymore (“not-needed”). These users generally did not see sufficient value to continue its use, e.g., “*I decided it was useless.*” While this reason appeared across domains, it was particularly salient for privacy practices, with five out of ten privacy practices abandoned most likely because their value was not recognized. Four of these practices pertained to browsing activities, such as using incognito mode: “*I have used it but don’t find it that helpful,*” and “*I did it once, just to see how it worked, but found it awkward.*”

In 50 cases (14%), participants abandoned a practice after perceiving that risk levels had diminished (“because-of-risk”). This was a dominant justification for abandoning a fraud alert or credit freeze, which were commonly adopted after experiencing fraud or having credit cards stolen, and dropped soon afterward. Similarly, 11 participants had used temporary credentials for online activities when engaging with risky services. Still, they abandoned it either because of its negative repercussions, (e.g.,

*“when I made friends it was embarrassing to have to admit I lied about my name”*) or because their online social interaction habits changed (e.g., *“I’ve done this before when I used to have fights with people online, but I don’t anymore”*).

Participants abandoned practices due to their impracticality in 41 instances (12%), providing complaints similar to those for partial adoption. 23 participants (7%) reported abandoning practices when they caused usage interference, mainly citing the same set of partially adopted practices.

In 21 cases (6%), participants abandoned the practice in favor of relying on their own judgment (“own-judgment-sufficient”). This attitude was most prominent for abandoning automatic update (10) to regain control over the *“what and when”* of software updates, e.g., *“I used to have them on because that was the default setting. Now I am more mindful of what software updates I actually want.”*

Another 21 participants (6%) abandoned a practice after adopting a service that served a similar purpose (“using-substitute”). This reason applied to practices across all three domains. We noted a trend of switching to tools that offer automated protection from relying on manual effort, as in the case of disabling third-party cookies (3), e.g., *“I run programs to clear my cookies frequently.”* Most participants made sensible decisions when supplanting recommended practices with their own solutions. For instance, *“If I visit a website I have bookmarked I don’t check [the URL] as I already verified it before I bookmarked the site.”* Password managers were the rare case where substitutes appeared to be less effective, e.g., *“I use a password manager, but only to store passwords I create. I do not use the password generator. I usually create long, difficult passwords that are more memorable to me than what a generator produces.”* However, prior research suggests that users’ self-generated passwords are typically weaker than random passwords generated by password managers [388].

### 6.3 Discussion

Our findings provide insights into how well security, privacy, and identity theft protection practices are adopted, particularly why certain practices are only partially adopted or abandoned. Users struggle to adhere to experts' online safety advice [107, 227]. Expert advice is often vague, inactionable, and contradictory [219, 415, 416]. Our findings suggest ways to develop better expert advice and effectively convey it to consumers. Moreover, our study indicates that usability issues exist widely across different domains and are a key contributor to partial adoption and abandonment. While usability research has largely focused on security practices, the usability of privacy and identity protection practices requires more attention. Additionally, tools and services that demand consistent user interactions were adopted the least, indicating the need for improvement.

**Bridge the gap between security and other safety practices.** While security practices exhibited relatively high adoption rates in our survey, most privacy practices were often used selectively or abandoned; many identity theft protection practices were not even considered. This finding is concerning, given that practices from different domains often intersect. For example, phishing is a common attack vector for identity theft [383]. Manual security practices (e.g, avoid clicking unknown links) are prone to cognitive errors and inconsistent application. In this case, assisted security such as 2FA and credit freezes can help prevent account compromise and identity theft; identity monitoring services can further facilitate mitigation and recovery after attacks occur. Thus, adopting multiple practices across domains can create additional security layers and synergistic effects.

Security is usually conceived as something related to passwords, antivirus, or cautious interactions with websites and emails [64, 242]. However, security advice and education also need to cover related privacy and identity protection practices to

help people achieve a more holistic online safety posture. Rather than overburden users with too much advice, experts should identify the most effective and actionable recommendations from each area, articulate how they complement each other, and create safety gains beyond those from adopting a single practice.

**Leverage at-risk situations for communicating advice.** Prior work has identified triggers for adopting security and privacy practices [107]. We find that experiencing security incidents, especially identity theft, drives the adoption of protective measures across all three domains. As such, opportunities to convey advice more effectively might exist in post-incident guidance, when people are highly motivated to resolve the situation and mitigate future risks. Required security and privacy notices such as data breach notifications could be leveraged accordingly [600]. Similar to phishing training materials [564], for people who are not direct victims of security incidents, vivid and detailed stories recounting the negative experiences (e.g., on being an identity theft victim [578]) might be more effective than merely listing factual harms. Such stories should further come with actionable preventative advice.

Nevertheless, practice adoption triggered by negative experiences might not be long-term. From our qualitative analysis, some participants reported following certain practices only in high-risk situations and abandoned the practice soon after the perceived risk decreased. The abandonment of risk-triggered behaviors should be assessed case-by-case. Some practices might not be relevant anymore due to changes in circumstances (e.g., abandoning incognito mode when the device is not shared). Yet interventions are needed when perceptions of decreased risk are misaligned with objective risks. For instance, some participants abandoned credit freezes and fraud alerts soon after data breaches, even though the objective identity theft risks may not change over time once sensitive information has been exposed. Thus, expert advice to users needs to communicate the risk’s persistence, i.e., what practices can be



used selectively (and in which situations) and what other practices require consistent long-term adoption to be effective.

**Tailor advice to audience characteristics.** Prior research suggests a possible “digital divide” in security and privacy self-protection: people with less education and lower socioeconomic status may have access to fewer resources, exposing them to further vulnerability [188, 264]. Our findings are more nuanced. More technology knowledge is correlated with higher levels of practice adoption. Interestingly, security/privacy expertise alone had no effect. A lower income is also correlated with higher adoption, especially for privacy practices, possibly because people with lower incomes might be more acutely aware of digital privacy harms [310]. Notably, most of our investigated privacy practices are free or have free options (e.g., anonymity systems such as VPN). These results suggest that expert advice needs to tailor to specific audiences to be effective [310]. For instance, the use of personas [131] and scenarios reflecting different audiences and their needs could help users identify solutions most suitable to them.

**Usability issues prevent full adoption across practices** In line with prior work [388, 546, 567], we identify usability issues as a critical contributor to partial adoption and abandonment across different practices, such as updating software, using a password manager, and using unique passwords. Users may partially or fully abandon a practice when it is difficult and inconvenient to implement, sometimes reaching the level of disrupting the user experience even when they recognize the practice’s value. While prior work has primarily advocated for improving the usability of assisted security practices such as 2FA [106], more usability research is needed for frequently abandoned or rejected privacy and identity protection practices to lower their barriers to adoption. Browsing-related privacy practices, in particular, show significant usability issues and deserve more attention. For instance, cleaning browser

cookies was considered impractical as it removes desired cookies (e.g., session and login cookies). Similar to purpose-oriented cookie consent banners [539], browsers and web standards could support cookie management controls that distinguish different types of cookies to let users set more meaningful preferences.

**More support for practices requiring recurring user interactions.** Practices requiring recurring interactions have significantly lower adoption rates than manual and automated ones. While the manual practices we investigated are primarily instructive rules of thumb (e.g., avoid clicking unknown links), they are prone to slip-ups and easily overruled by users' judgment. Conversely, most assisted practices (e.g., anonymity systems and password managers) generally require some level of expertise for initial setup, which may scare non-tech-savvy users away [389], or have known usability issues that significantly impact user experience [546].

For tool-based practices such as using a password manager, their features and functionality need to be better communicated to prospective users to dispel identified misconceptions. For instance, most participants who adopted password managers chose those built into their browsers due to direct integration into the browsing experience. By contrast, dedicated password managers often require extra steps to retrieve passwords. Even eliminating a few clicks can make a big difference as users' compliance budgets are extremely limited [219]. Lastly, minor tweaks to mechanisms can have diminishing returns compared to paradigm changes. For instance, despite its flaws and weaknesses, biometric authentication can be used with password managers to ease the adoption and usability of multiple practices at once [226].

## CHAPTER VII

# Icons and Link Texts for Conveying Privacy Choices<sup>1</sup>

Chapters III to VI constitute the first half of my dissertation about the types of privacy-protective behaviors that consumers adopt or do not adopt and the reasons why. Regarding data breaches involving account credentials, changing passwords and reviewing financial statements were the most popular [329]. Beyond data breaches, security behaviors received wider adoption than privacy or identity theft protective behaviors; behaviors requiring cognitive adherence or with automated protection were favored over behaviors requiring recurring interactions with a tool or interface [603]. Nonetheless, consumers' adoption of protective behaviors is subject to psychological factors (e.g., optimism bias and present bias in response to the Equifax data breach [602]) and issues with expert advice, which is often burdensome and inactionable [600, 603, 604]. Starting from this chapter, the second half of my dissertation seeks to find ways to better support consumers' adoption of privacy-protective behaviors.

In this chapter, I present a study on how to effectively convey to consumers the presence of privacy choices on websites through icons and accompanying text descriptions (which we refer to as link texts). The study is motivated by the reality that

---

<sup>1</sup>This chapter is based on: Hana Habib\*, Yixin Zou\*, Yaxing Yao, Alessandro Acquisti, Lorrie Cranor, Joel Reidenberg, Norman Sadeh, & Florian Schaub. 2021. Toggles, Dollar Signs, and Triangles: How to (In)Effectively Convey Privacy Choices with Icons and Link Texts. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (pp. 63:1–63:25). [199]

\*The first two authors contributed equally to this research.

notice and choice, while often challenged as a broken mechanism [72, 100, 418, 466], is still the primary component of privacy regulation and consumer protection guidelines. Privacy advocates, legal experts, and academic researchers have argued for standardized mechanisms of providing privacy notices and choices [29, 263], as they facilitate the learning process for consumers and enable comparisons of data practices across companies for regulatory purposes. We chose to focus on icons as a promising mechanism for conveying privacy choices since icons (in principle) can communicate concepts quickly and concisely across linguistic and cultural differences [428] among other advantages. Prior privacy icons were mostly about data flows or specific data practices (e.g., [124, 353, 428]), but few have explored icons for privacy choices — a key component of consumer privacy regulation [100, 144, 373].

For this study, we considered the presence of generic privacy choices and opt-outs for the sale of personal information, the latter mandated by the California Consumer Privacy Act (CCPA). Through an iterative process, we identified that a blue stylized toggle icon best conveyed the idea of choices. By contrast, icons focused on the sale of personal information created misconceptions about what would happen after clicking the icon. The Digital Advertising Alliance (DAA)’s Privacy Rights icon [494] and the older AdChoices icon [495], as comparison points for our newly designed icons, suggested “more information” but not “choice.” For icon-text combinations, “Privacy Options” paired with the blue stylized toggle icon best conveyed the presence of privacy choices. The link texts mandated by CCPA (“Do Not Sell My Personal Information” and “Do Not Sell My Info”) effectively conveyed the expectation of choices related to the sale of personal information in combination with most icons. Our follow-up study of an icon proposed by the California Office of Attorney General (OAG) revealed that even minor design changes could severely reduce an icon’s comprehension and increase misconceptions.

## 7.1 Background

**Requirements for privacy choices.** GDPR requires businesses to provide privacy choices to European consumers, including an option to request the erasure of personal data about them (Art. 17) and opt-outs for data processing for direct marketing purposes (Art. 21) [144]. GDPR also emphasizes the usability of privacy notices and choices, requiring that notices be provided in “a concise, transparent, intelligible and easily accessible form, using clear and plain language” (Art. 12). In the US, privacy choices are regulated by sector-specific federal and state privacy laws, such as the CAN-SPAM Act’s requirement of marketing email opt-outs [497], and COPPA’s requirement of honoring parental requests about data collection and deletion for children under 13 [500]. Since January 2020, CCPA has provided California residents the right to opt out of the sale of their personal information by companies [373]. The first draft of the proposed CCPA regulations specified that this opt-out should be provided, at a minimum, through “an interactive form accessible via a clear and conspicuous link titled ‘Do Not Sell My Personal Information,’ or ‘Do Not Sell My Info,’<sup>2</sup> on the business’s website or mobile application” as well as an unspecified optional “opt-out button or logo” [373]. In November 2020, Californians voted to pass the California Privacy Rights Act (CPRA) [374], which amends and expands consumer privacy rights stipulated by CCPA.

Self-regulatory requirements exist for online advertising practices [492, 510]. Since 2010, DAA has required its member companies to provide opt-outs for tracking-based targeted advertising by placing the AdChoices icon and an approved text above an ad [492]. In 2021, DAA introduced a Privacy Rights icon (a green variant of the AdChoices icon) to address CCPA’s opt-out requirements [494]. Additionally, the Interactive Advertising Bureau Europe has published the Transparency and Consent

---

<sup>2</sup>“Do Not Sell My Info” was mandated in the proposed CCPA regulations [373] but got eliminated in the final version [375] after we completed our study.

Framework for obtaining consumer consent under GDPR [508].

Companies' compliance with legal and self-regulatory requirements varies. Opt-outs for email communications are common due to the CAN-SPAM Act, with most US companies offering links to unsubscribe within email messages and privacy policies [118, 198]. However, privacy choices related to targeted advertising or data deletion are less common [198, 511]. Even when they exist, privacy choices appear at inconsistent locations and often include unhelpful information or broken links [198]. Furthermore, research on CCPA's do-not-sell provision has shown that consumers struggle to locate the required links to opt out, and the opt-out processes are permeated with dark patterns [312]. As such, consumers face considerable barriers to exercising privacy choices [118, 197, 198].

**Privacy icons.** Researchers have proposed various privacy icons as succinct indicators of complex privacy concepts. Some privacy icons represent specific data practices, such as Disconnect.me's icons for different types of tracking [124] and Mozilla's icons for retention periods and third-party data sharing and use [353]. Some only serve specific application domains, such as social media [234], web links [259], or webcams [134, 404], while others can apply across contexts [223]. Icons are also commonly used as security indicators (e.g., a lock in a browser's URL bar that indicates HTTPS [148]). However, prior work has found that users tend to ignore or misunderstand these indicators [300, 439]. Fewer privacy icons are designed to convey privacy choice, consent, or opt-outs. The Stanford Legal Design Lab has proposed icons [480] that could potentially indicate privacy choices, but they have not been empirically evaluated. While the Data Protection Icon Set (DaPIS) [428] has been user-tested, it is specific to GDPR's consumer privacy rights.

Icons have several advantages that can address the limitations of traditional privacy notices. First, icons can visually communicate information concisely while cir-

cumventing language and cultural barriers [320]. Second, icons can be helpful information markers since they are easy to recognize [225]. Finally, when placed next to lengthy privacy statements, icons can enhance readability by helping users navigate the text [428]. In a review of iconography guidelines, Bühler et al. [61] summarized principles for designing effective icons—they should be based on users’ knowledge and needs, utilize well-known concepts, and closely mimic real-world objects. However, designing comprehensible icons is challenging. Icons alone sometimes perform worse than text-only or icon-text interfaces in assisting learning [580]. Fischer-Hübner et al. [153] therefore argue that icons should be used alongside text to illustrate data practices in privacy policies and aid user comprehension. Beyond an icon’s comprehensibility, the focus of our study, discoverability is another challenge. For instance, the size, position, state, and color all impacted how visible the AdChoices icon was to users on a mobile device [173].

Privacy icons explored in prior work have primarily focused on communicating data practices, but few proposed privacy icons have received wide adoption. Even widely adopted icons, such as DAA’s AdChoices icon, are problematic [173, 537]. In addition, not much work has focused on using icons to convey privacy choices effectively to consumers. We fill this gap by iteratively designing and evaluating privacy choice icons and associated link texts. Complementing prior research on icons for GDPR-specific user rights [428], we focus on conveying the presence of general privacy choices and CCPA-mandated do-not-sell opt-out.

## 7.2 Study Overview

Between November 2019 and February 2020, we conducted a series of studies to iteratively design and evaluate two types of icons and associated link texts: one indicating the presence of generic privacy controls on websites, and the other indicating choices related to the sale of personal information, as required by the CCPA. Our re-

search involved two pre-studies (one focusing on icons and the other on link texts), a large-scale online experiment to evaluate icon-link text combinations, and a follow-up evaluation of an icon that the Office of the California Attorney General (OAG) had proposed based on our initial findings.

**Icon Pre-Study (Section 7.3,  $n=520$ ).** We developed 11 privacy icons that center on three choice-related concepts: the broad idea of *choice*, the action of *opting out*, and choices regarding *the sale of personal information*. We iteratively refined and tested these icons to identify which to include in our main experiment. Our icon pre-study suggests that a stylized toggle switch was promising for conveying the presence of choice; three icons including elements of dollar signs, slashes, stop signs, and ID cards were good candidates for conveying the CCPA do-not-sell opt-out.

**Link Text Pre-Study (Section 7.4,  $n=540$ ).** We tested 16 textual descriptions, or link texts, to accompany our icons. We analyzed how each link text, when displayed alone, was interpreted by participants; and identified three link texts (“Privacy Options,” “Privacy Choices,” and “Personal Info Choices”) with mostly correct interpretations. The two CCPA link texts (“Do Not Sell My Personal Information” and “Do Not Sell My Info”) effectively indicated choices related to the sale of personal information, but they did not generalize to broader privacy-related choices.

**Icon-Text Combinations Evaluation (Section 7.5,  $n=1,468$ ).** We conducted a large-scale, nearly full-factorial online experiment to evaluate how well 23 combinations of icons and link texts, selected from our pre-studies, communicated the presence of privacy choices and do-not-sell choices. We showed participants one icon-text combination on a screenshot of a fictitious online shoe retailer webpage, mimicking how users may see such privacy choice indicators in the real world. A blue stylized toggle icon paired with the link text “Privacy Options” best conveyed the presence of privacy choices. The two CCPA link texts effectively conveyed the presence of do-not-sell opt-outs when paired with most icons.



**OAG Icon Evaluation (Section 7.6,  $n=421$ ).** After we shared our results with the OAG, they proposed an icon for the CCPA’s do-not-sell opt-out, which was similar to our stylized toggle icon but with notable deviations. We conducted a follow-up experiment to explore the impact of the icon’s toggle style and color on expectations for do-not-sell choices. Compared to our stylized toggle icon, participants were much more likely to perceive the OAG’s proposed icon as a toggle switch rather than a static icon.<sup>3</sup>

## 7.3 Icon Pre-Study

We developed 11 icons related to privacy choices and evaluated how users interpreted the icons with and without a text description. We found that a stylized toggle icon effectively communicated the concept of choice, but communicating the concept of “privacy choice” was difficult without text. While icons with arrows to depict removal were mostly unsuccessful, icon elements focusing on “do not” and “sell” could communicate an opt-out for the sale of personal information. However, participants often misunderstood an icon without a text description.

### 7.3.1 Icon Development

**Icon ideation.** To explore potential icon candidates, we leveraged existing privacy iconography to generate three key concepts in line with our objectives: the broad concept of *choice*, the action of *opting out*, and a specific opt-out related to the *sale of personal information* for CCPA. We did not attempt to design an icon that visualizes privacy since privacy is a broad concept with many interpretations [369]. Additionally, we did not test existing privacy and security icons since some of them

---

<sup>3</sup>In December 2020, the OAG published the fourth set of modifications to the CCPA regulations [377], recommending that businesses use our blue stylized toggle icon next to the CCPA link text when notifying consumers of their right to opt out of the sale of personal information. The OAG maintains a website that includes documents relevant to CCPA rulemaking [379].

are already known for representing other concepts unrelated to privacy choices (e.g., lock or shield for HTTPS indicator [148]), and others focus more on specific data practices [428].

To capture a wide range of icon ideas embodying the three choice-related concepts we identified, we conducted design ideation activities at our institutions with colleagues interested in privacy and security research. During the activities, participants drew ideas on sticky notes and discussed themes with the group. We then conducted affinity diagramming [286] of the sketches by grouping similar ideas and identifying themes in the visual elements participants used to represent the three concepts (see Figure 7.1). In selecting themes to iterate upon further, we eliminated those focusing on privacy more than choice due to our goal of conveying choice. We also eliminated themes that seemed too abstract from privacy choice (e.g., leaving or refusing something) or difficult to graphically depict (e.g., third parties). Considering that web icons are generally small, we further eliminated themes that would produce unrecognizable icons when shrunk down in size due to complexity (e.g., exchange/trade-off of data for money). In the end, we identified five themes (see Table 7.1) that had the potential to represent our three choice-related concepts effectively.

**Refinement with graphic designers.** Next, we worked with three graphic designers to develop icons for the five themes. The graphic designers worked individually with sketches from our brainstorming sessions as a starting reference. They were encouraged to produce variants and alternative designs, such as varying the shape or size of icon elements. Then, the research team jointly reviewed the graphic designers' work and selected 11 icon designs as candidates for user testing in the icon pre-study.

Table 7.1 shows all 11 candidate icons. Three icons intended to convey the broad idea of *choice*: one featured a toggle—a standard UI element for turning on or off settings [35], and two featured checkboxes (transitioning from a checked to an unchecked

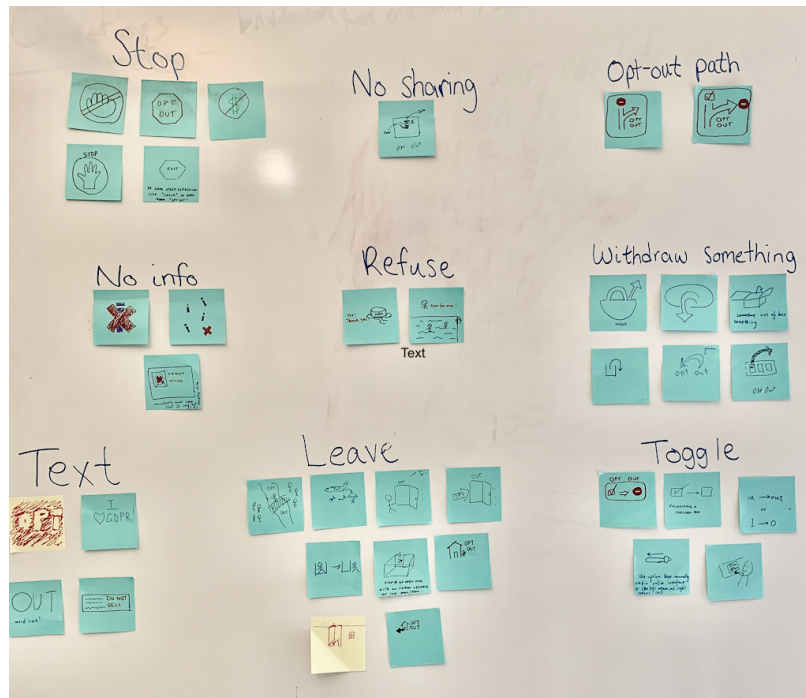


Figure 7.1: Common themes that emerged in one of the brainstorming sessions for an icon that conveyed *opting-out*.

Choice Concept	Icon Themes	Preliminary Icons
Choice/consent	<ul style="list-style-type: none"> <li>toggle switch</li> <li>change toggle or checkbox choice</li> </ul>	<i>Stylized-Toggle</i>
		<i>Changed-Choice</i>
		<i>DoNot-Checked</i>
Opting Out	<ul style="list-style-type: none"> <li>withdrawing something from a basket or box</li> </ul>	<i>Box-Arrow</i>
		<i>Circle-Arrow</i>
		<i>Folder-Arrow</i>
Do-Not-Sell Choices	<ul style="list-style-type: none"> <li>no money/selling</li> <li>stop selling personal info</li> </ul>	<i>DoNot-Dollar</i>
		<i>Slash-Dollar</i>
		<i>Stop-Dollar</i>
		<i>ID-Card</i>
		<i>Profile</i>
Existing icons		<i>DAA Privacy Rights</i>
		<i>DAA AdChoices</i>

Table 7.1: Icon themes that emerged in ideation sessions for each choice-related concept, and the corresponding icons included in our preliminary testing.

box or negating a checkbox) since checkboxes are common in online forms and consent interfaces [35]. Three icons intended to convey the action of *opting out*, which is analogous to withdrawing consent: two had an arrow coming out of simple shapes (a circle and a box), and the third used a file folder to represent personal data. Five icons intended to convey *do-not-sell* choices: three used different negations of a dollar sign to represent stopping a sale, and two further included a “person” element to represent personal data. To minimize the potential biasing effects of color in our pre-study, we created the initial versions of our icons in black and white. Additionally, we included DAA’s AdChoices [495] and Privacy Rights [494] icons in our icon pre-study as a benchmark for industry practices.

### 7.3.2 Preliminary Icon Testing

We conducted an initial round of user testing on all 11 candidate icons to decide which to test in subsequent studies. We developed an online survey to capture qualitative and quantitative responses that would help us identify feasible icons for indicating the presence of generic privacy choices and do-not-sell choices.

**Study protocol.** Our initial testing sought to identify difficult-to-interpret icons and specific icon elements that help indicate privacy or do-not-sell choices. We implemented a between-subjects design, in which we randomly showed each participant one of the icon candidates without context. To examine the impact of placing a link text next to the icon (as required by CCPA), half of the participants saw the icon displayed with the text “Do Not Sell My Personal Information.” We hypothesized that this text would aid the comprehension of icons intended to convey do-not-sell choices.

After presenting the icon, we asked participants to provide open-ended responses regarding their interpretation of the icon and their expectations of what would happen

if they clicked on it—this was to capture their unprimed impressions of the icon. As a complementary quantitative data point, we next showed participants all icons, asked them to select which would best convey the presence of privacy choices and do-not-sell choices respectively, and explained the rationale behind their selection.<sup>4</sup> We then asked participants about their familiarity and expectations regarding DAA’s AdChoices icon [495] to evaluate the recognizability and comprehension of an already widely deployed privacy choice icon. Lastly, we collected participants’ demographic information and asked about awareness of a US law that required companies to provide a “do not sell” option. Appendix E.1 includes the full set of survey questions.

For this and all subsequent studies, we did not collect personal data from participants, and we instructed participants to avoid revealing personal information in their open-ended responses. The Institutional Review Boards at Carnegie Mellon University and the University of Michigan approved all study protocols.

**Recruitment and sample demographics.** We recruited 240 participants from Amazon’s Mechanical Turk (MTurk) to ensure roughly 20 responses per condition—a sufficient number for capturing a variety of opinions for descriptive analysis. We set the recruitment filter as US residents over 18 years old, with a 95% or higher approval rate. Before answering survey questions, participants reviewed a consent form and confirmed their age and residency eligibility. The average study completion time was 5.25 minutes, and participants were compensated \$1.00 (average \$11.43/hour).

In line with demographic characteristics of MTurk workers [230], our samples for this and the subsequent studies were diverse but not representative of the US general population: they skewed younger, more male, and more educated. We summarize participant demographics here once as they were fairly uniform across all studies, and we provide detailed demographics for each study in Table E.1. Participants were

---

<sup>4</sup>The Privacy Rights icon was green when presented alone but black-and-white when presented with other icons to eliminate the impact of color on participants’ selection.

residing in most US states (with 10-20% living in California) and somewhat tech-savvy (with 23-48% reporting education or job experience in computer science, computer engineering, or IT). In addition, 3-10% of participants reported awareness of a US law that required companies to provide a “do not sell” option, with relatively higher percentages in the icon-text combinations and OAG toggle evaluations, indicating a potential increase of awareness after CCPA went into effect. Once participants completed one of our studies, they were excluded from participating in subsequent studies evaluating icons and link texts.

**Data analysis.** We conducted a thematic analysis [434] of participants’ qualitative responses. One author examined a subset of the qualitative data to identify common themes and developed an initial codebook. The team then discussed the initial codebook, adding and modifying codes as necessary. To ensure high consistency in coding, two authors coded 20% of all responses and additional responses if needed until reaching a Cohen’s  $\kappa$  of at least 0.7, which is considered sufficient agreement [156] (average  $\kappa=.81$  across all questions).<sup>5</sup> Most responses mapped clearly to a code, and multiple researchers discussed ambiguous responses before coding them. After we achieved high inter-coder reliability, one researcher coded the remaining responses. We calculated descriptive statistics of coded qualitative data rather than conducted hypothesis testing, since our primary objective for this pre-study was to eliminate icons that appeared confusing or did not effectively convey intended concepts from further consideration. Eleven responses were excluded from the analysis, as they only included text that did not respond to the open-ended questions. We note the number of responses excluded from the analysis for this and subsequent studies in Table E.1.

---

<sup>5</sup>Responses to the AdChoices interpretation lacked variations, meaning that a single disagreement between coders would cause a significant drop in Cohen’s  $\kappa$ . For this question, we used inter-coder percentage agreement instead to measure inter-coder reliability and ensured the percentage agreement was at least 75%.

**Findings.** As shown in Table E.2, most icons did not lead to their intended interpretations when shown alone. Participants did not exhibit a clear preference for which icon best represented generic privacy choices, but most chose *Slash-Dollar* as the icon for representing do-not-sell choices.

**A stylized toggle icon best conveyed “choice.”** Among the three icons that were intended to convey choice, participants commonly associated *Stylized-Toggle* with the notion of choosing or selecting something. Participants thought of “completion” (i.e., marking something as completed or completed downloads), rather than choice, upon seeing *DoNot-Checked*. *Changed-Choice* received a variety of interpretations, suggesting that it would not work well for indicating privacy choices either.

**Icons for conveying “opting out” were confusing.** Though two participants interpreted *Box-Arrow* as “removing something” (as intended), other participants interpreted it differently. Participants mostly interpreted *Circle-Arrow* as something related to motion, and they focused on the folder element rather than the arrow in *Folder-Arrow*; neither prompted participants to think of opting out.

**Dollar signs suggested payment rather than selling.** All icons intended for do-not-sell choices conveyed a sense of payment or money, but not selling. Interpretations included “cash or American dollars are not accepted,” “something is free,” “something requires payment,” and “something related to an account balance.” Promisingly, three participants connected *ID-Card* with a person and money, which aligns with its intended purpose of signaling do-not-sell choices.

**No clear preference for the privacy choices icon.** Participants were divergent in their opinions of which icon best represented choices about the use of personal information (see Figure 7.2). *Stylized-Toggle* was selected most frequently, though *ID-Card*, *DAA*, and *Folder-Arrow* were not far behind. In open-ended responses, participants identified certain icon elements that conveyed privacy choices to them, including “select/choose” (32.3%), “money/selling” (21.0%), “personal infor-

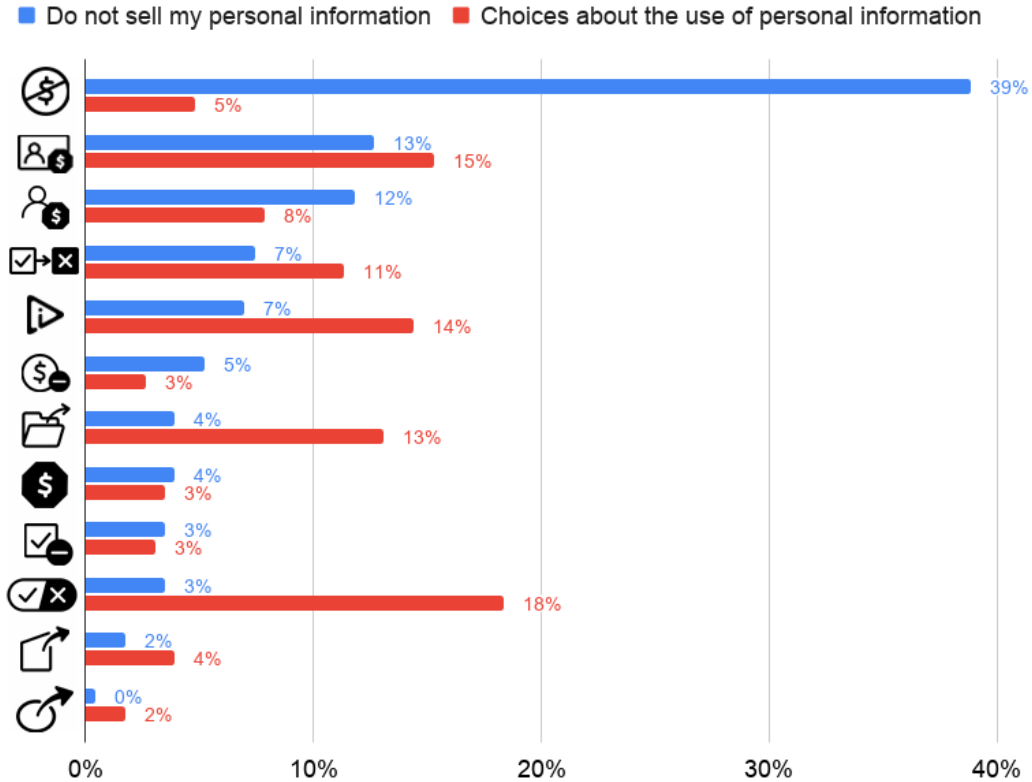


Figure 7.2: The percentage of participants selecting an icon that best conveys there’s an option to (1) “tell websites ‘do not sell my personal information’” (blue); and (2) “make choices about the use of my personal information” (red) in the preliminary testing.

mation” (19.2%), and “stop/do not” (16.6%). The mentioning of “money/selling” and “stop/do not” suggests potential priming effects from the question about the best icon for do-not-sell choices or the “Do Not Sell My Personal Information” link text when presented.

**Slash-Dollar preferred as “do-not-sell” icon.** Participants exhibited a clear preference for which icon best represented do-not-sell choices as 38.9% selected *Slash-Dollar* (see Figure 7.2). In open-ended responses, participants mentioned “money/selling” (48.9%), “stop/do not” (46.7%) and “personal information” (21.0%) as important icon elements for conveying do-not-sell choices. Participants preferred “stop/do not” to be represented by a circle with a slash rather than an octagonal stop sign or a do-not-enter sign, as indicated by the stark difference between *Slash-Dollar* and *DoNot-*





Figure 7.3: Promising icons from preliminary testing in their refined versions.

*Dollar/Stop-Dollar*. This finding suggests that the octagon shape in *Stop-Dollar* may be difficult to recognize as a stop sign without color, and the “do not enter” sign in *DoNot-Dollar* was not widely recognized or was misidentified as a minus sign.

### 7.3.3 Refined Icon Testing

Our preliminary testing suggested comprehension issues with most icons but surfaced some promising candidates. In selecting icons for further testing, we included *Stylized-Toggle* and *ID-Card* as candidates for privacy choices: the former appeared to communicate “choice” well, and the latter was ranked highly by participants in preliminary testing. For do-not-sell icon candidates, we included *Slash-Dollar* due to participants’ preferences and *Stop-Dollar* to explore whether the color would increase recognition of the stop sign.

We evaluated refined versions of the four icons mentioned above and the DAA’s Privacy Rights icon (see Figure 7.3) to further narrow down icon selections for the larger-scale icon-text evaluation. Specifically, we colored the stop sign and slash red in *ID-Card*, *Stop-Dollar*, and *Slash-Dollar*, and made the dollar sign in *Slash-Dollar* more readable. We colored *Stylized-Toggle* blue—a neutral color that does not convey a particular state, unlike green or red.

**Study protocol.** We followed the same protocol as before to evaluate the five icons. We randomized the order of the “best icon” questions for privacy/do-not-sell choices to mitigate potential priming effects. We recruited 280 participants (roughly 28 per condition) to detect a medium effect size (0.3) [85] with at least 80% power for our

planned statistical analysis. We aimed for a medium effect size due to the study’s exploratory nature and to save the budget for oversampling in the icon-text evaluation. The average study completion time was 4.50 minutes, and each participant received \$1.00 (average \$13.30/hour).

**Data analysis.** We followed the same qualitative data analysis approach as before ( $\kappa=0.79$ ). Additionally, we collaboratively categorized the codes used to analyze open-ended responses to “What does this symbol communicate to you?” as *correct* or *incorrect* interpretations regarding the icon’s intended purpose. We then used these binary labels as the dependent variable of Chi-squared tests (or Fisher’s exact tests when applicable) to determine whether the overall difference in study conditions was statistically significant. Follow-up pairwise comparisons were adjusted with Holm-Bonferroni corrections.

### 7.3.3.1 Findings

Participants interpreted *Stylized-Toggle* as an indicator of some form of choice and preferred it over other candidates for conveying generic privacy choices. Consistent with the preliminary testing, participants preferred *Slash-Dollar* for communicating do-not-sell choices. The CCPA link text’s presence made participants more likely to expect an icon to lead to do-not-sell choices.

**Stylized-Toggle was interpreted as intended.** Table E.3 provides common interpretations of each icon when displayed without the CCPA link text. A Fisher’s exact test showed significant differences between icons, when presented alone, in generating correct interpretations that align with the icon’s intended meaning ( $p<.001$ ,  $V=0.58$ ). Pairwise comparisons found that *Stylized-Toggle* was more likely to be interpreted correctly compared to other icons (all  $p<.001$ ). Open-ended responses suggested that *Stylized-Toggle* was primarily interpreted as an option to “accept/decline”

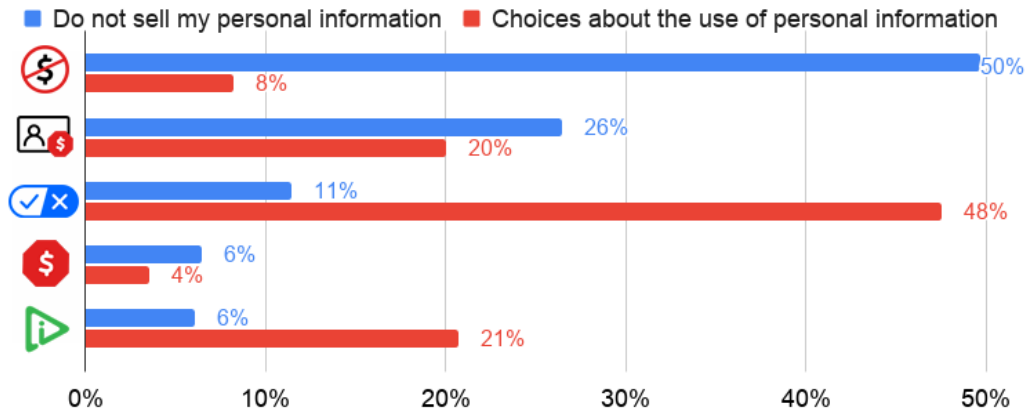


Figure 7.4: The percentage of participants selecting for an icon that best conveys that there’s an option to “tell websites ‘do not sell my personal information’” (blue); and “make choices about the use of my personal information” (red) in the refined testing.

or “activate/deactivate” something. In contrast, the interpretations of other icons were often misaligned with their intended meanings. DAA’s Privacy Rights icon conveyed an option to “get more information” but did not suggest a choice or opt-out. Common interpretations of *Slash-Dollar* were “something is free or does not require money” or “cash or American dollars were not accepted.” *ID-Card* was mostly interpreted as “something costs money.” *Stop-Dollar* was similarly associated with money but not selling.

**Clear icon preference for privacy choices and do-not-sell choices.** As shown in Figure 7.4, when the icons were colorized, participants exhibited a clear preference for *Stylized-Toggle* to represent choices about the use of personal information. 16.8% of participants explicitly stated that a toggle “with a checkmark and an X in it” nicely conveyed choice. Similar to the preliminary testing, *Slash-Dollar* was selected most frequently as the icon for conveying do-not-sell choices; *ID-Card* ranked second (see Figure 7.4).

**CCPA link text led to expectations of do-not-sell choices.** A Chi-squared test showed that participants who saw the CCPA link text were significantly more likely to interpret the icon as its intended meaning ( $p < .001$ ,  $\phi = 0.38$ ). Of the 139

participants who saw an icon with the CCPA link text, 43.2% (60) expected some form of choice to stop websites from selling their personal information. 13.7% (19) expected the ability to configure the types of personal information they could prevent from being sold or entities to which information is sold. 31.7% (44) expected being immediately opted out of the sale of personal information after clicking. There was no significant difference between icons in creating any of these expectations, suggesting that the link text impacted participants' expectations rather than the icon. Notably, the CCPA link text's presence did not eliminate misconceptions, such as expecting a different type of privacy choice (e.g., opting out of data collection on the website) or interpreting the link text as a warning not to give out their personal information to websites.

**DAA's AdChoices icon is still mostly unknown.** Even though the DAA launched its AdChoices icon in 2010, only 40 (14.3%) participants recalled seeing this icon before. The most common expectation of the AdChoices icon was that it provided more information about something, as indicated by 152 (54.3%) participants. Only six participants expected it would lead them to choices related to targeted advertising. Our results confirm Leon et al.'s 2011 findings that there is little recognition of the AdChoices icon [294]—time and widespread adoption do not seem to have increased consumer awareness of this icon.

## 7.4 Link Text Pre-Study

We developed and iteratively evaluated potential link texts to accompany our icons and aid comprehension. In addition to the CCPA-mandated link text, which was a possible candidate for conveying the do-not-sell controls, we also wanted to brainstorm and test link texts for conveying generic privacy controls. We found that “Privacy Choices” was the best candidate for conveying generic privacy controls with few misconceptions, closely followed by “Privacy Options.” The CCPA link text

- |                                       |  |
|---------------------------------------|--|
| • Do Not Sell My Personal Information | • Privacy Options                          |
| • Do Not Sell My Info                 | • Privacy Opt-Outs                         |
| • Don't Sell My Info                  | • Privacy Choices                          |
| • Do Not Sell <sup>a</sup>            | • Personal Info Choices                    |
| • Don't Sell <sup>a</sup>             | • Personal Info Options                    |
| • Do-Not-Sell Choices <sup>a</sup>    | • Personal Info Opt-Outs                   |
| • Do-Not-Sell Options                 | • Do Not Sell My Info Choices <sup>b</sup> |
| • Do-Not-Sell Opt-Outs <sup>a</sup>   | • Do Not Sell My Info Options <sup>b</sup> |

<sup>a</sup> Preliminary link text testing only, <sup>b</sup> Refined link text testing only.

Table 7.2: Link texts tested in the link text pre-study.

variants performed well in conveying do-not-sell opt-outs but did not generalize to other types of privacy controls.

#### 7.4.1 Link Text Development

We generated link text candidates by identifying words or phrases corresponding to the three icon concepts we focused on (*choice*, *opting-out*, and *do-not-sell*). During our ideation, we observed that link texts could follow a pattern of two components: a privacy-focused prefix and, optionally, a choice-focused suffix. We wanted to explore whether the general prefix “privacy” or the more specific prefix “personal info” would convey the type of choices more clearly. For the suffix, we hypothesized that the broad terms “choices” and “options” would create different expectations compared to “opt-out,” a more specific type of choice. We also included the two CCPA do-not-sell opt-out texts [373] and their variants—including an abbreviated version (“Don’t Sell My Info”) and versions emphasizing *choice* rather than information (e.g., “Do-Not-Sell Choices”)—to control for confounds and explore potential alternatives to the CCPA link texts.

Our initial set included 14 link texts revolving around six words or phrases: personal info/privacy/do-not-sell for the prefix and choices/options/opt-outs for the suffix. After preliminary testing, we eliminated four with poor comprehension and added two for further testing. Table 7.2 shows the full set of link texts we evaluated.

### 7.4.2 Preliminary Link Text Testing

We tested the initial link text set using a similar protocol as the icon pre-study. Based on the findings, we eliminated four candidates from subsequent testing and added two more variants of the CCPA link texts.

**Study protocol.** We showed each participant one of the 14 candidate link texts at random, styled as a hypertext link but non-clickable, without an icon or other context. We asked participants to describe their expectations of what would happen if they clicked on the link and interpretations of specific text components. Then, we presented eight scenarios constructed from open-ended responses from the icon pre-study and asked participants to rate the likelihood that clicking on the link would lead to each scenario. Two scenarios were accurate expectations related to privacy notices and choices, three were accurate expectations related to do-not-sell, and three were misconceptions. Lastly, participants were asked to provide demographic information and indicate their familiarity with the CCPA. We recruited 140 participants on MTurk (roughly ten responses per condition) to have a diverse set of qualitative responses for descriptive analysis. The average study completion time was 4.20 minutes, and each participant received \$1.00 (average \$14.29/hour).

**Data analysis.** We coded participants' open-ended responses using the same thematic analysis approach as in the icon pre-study ( $\kappa=0.89$ ). The coded data was used for descriptive analysis only, as our primary goal was to identify link texts with high rates of misconceptions and eliminate them from further consideration.

**Findings.** Our preliminary testing of link texts suggested a more significant influence of the prefix, rather than the suffix, on expectations of what happens after clicking the link. "Personal information" was understood as personally-identifiable information, and its absence led to misconceptions about the word "sell."

**“Personal information” was primarily interpreted as PII (personally identifiable information).** When asked to interpret the phrase “personal information,” “personal info,” or “info,” 33 of the 57 participants (57.9%) who saw a corresponding link text listed examples of PII, such as name and birthday. 11 participants interpreted the phrase as demographic information, such as age or gender. In addition, nine participants thought it referred to their IP address or location, and another nine believed it referred to cookies or past activities on the website or elsewhere.

**“Sell” on its own was often misunderstood.** Without an explicit reference to personal information, participants struggled to identify the subject to which “sell” referred. Among the 45 participants who saw one of the “do not sell” variants without “personal information” or “my info,” 18 (40.0%) thought the sale referred to a physical product. Four thought the sale was related to stocks or money, and five did not know what the sale was about. Given that participants saw the link text with no further context, it is not surprising that such misconceptions occurred.

### 7.4.3 Refined Link Text Testing

Our preliminary testing showed that link texts containing the word “sell” without “info” did not convey privacy choices or do-not-sell choices well. Therefore, we eliminated four corresponding link texts from further testing but retained “Do-Not-Sell Options,” which conveyed a control/choice related to personal information about as frequently as “Privacy Opt-Outs” and “Personal Info Options.” In addition, we included two new link texts (“Do Not Sell My Info Choices” and “Do Not Sell My Info Options”) to assess how adding choice-related suffixes would affect the interpretation of the CCPA-mandated link texts. We did not test “Do Not Sell My Info Opt-Outs,” as our preliminary testing suggested “opt-outs” might be less intuitive than “choices” or “options.”

**Study Protocol.** We recruited 400 additional participants, roughly 33 per condition, to detect a medium effect size (0.3) with at least 80% power for our planned statistical analysis comparing expectations generated by the candidate link texts. The average study completion time was 4.1 minutes, and participants were compensated \$1.00 (average \$14.63/hour). Since we used the same protocol and survey instrument, we aggregated participant responses with those collected from the preliminary testing for the analysis.

**Data analysis.** We followed the same qualitative data analysis approach as in previous studies; two authors coded 20% of the data ( $\kappa=0.81$ ), and one author coded the remainder. For this and the following studies, we structured the codebook hierarchically by grouping codes into four categories (high-level codes) for category-level analysis. Specifically, we labeled “yes” or “no” for whether a code conveyed (1) the concept of choice, (2) the ability to opt out of the sale of personal information, (3) the concept of privacy broadly, and (4) misconceptions.<sup>6</sup> Three authors completed the mapping for all codes together and resolved any disagreements. We then used the values of these categorizations as the dependent variables in Pearson chi-square or Fisher’s exact tests, with link text conditions as the independent variable. Pairwise comparisons were Holm-Bonferroni corrected.

### 7.4.3.1 Findings

As seen in Figure 7.5, participants’ expectations significantly varied across link texts. “Privacy Choices” created the least misconceptions. The CCPA link texts and their variants successfully led to expectations of do-not-sell choices.

**Link text suffix did not impact expectations of choices.** 47.9% of partic-

---

<sup>6</sup>For example, the response “It would give you the option not to have your personal information given, shared, or sold to someone else” was coded as “choices: do not sell.” For high-level categories, the code was labeled as “yes” for conveying choice and do-not-sell, and “no” for conveying privacy or a misconception.



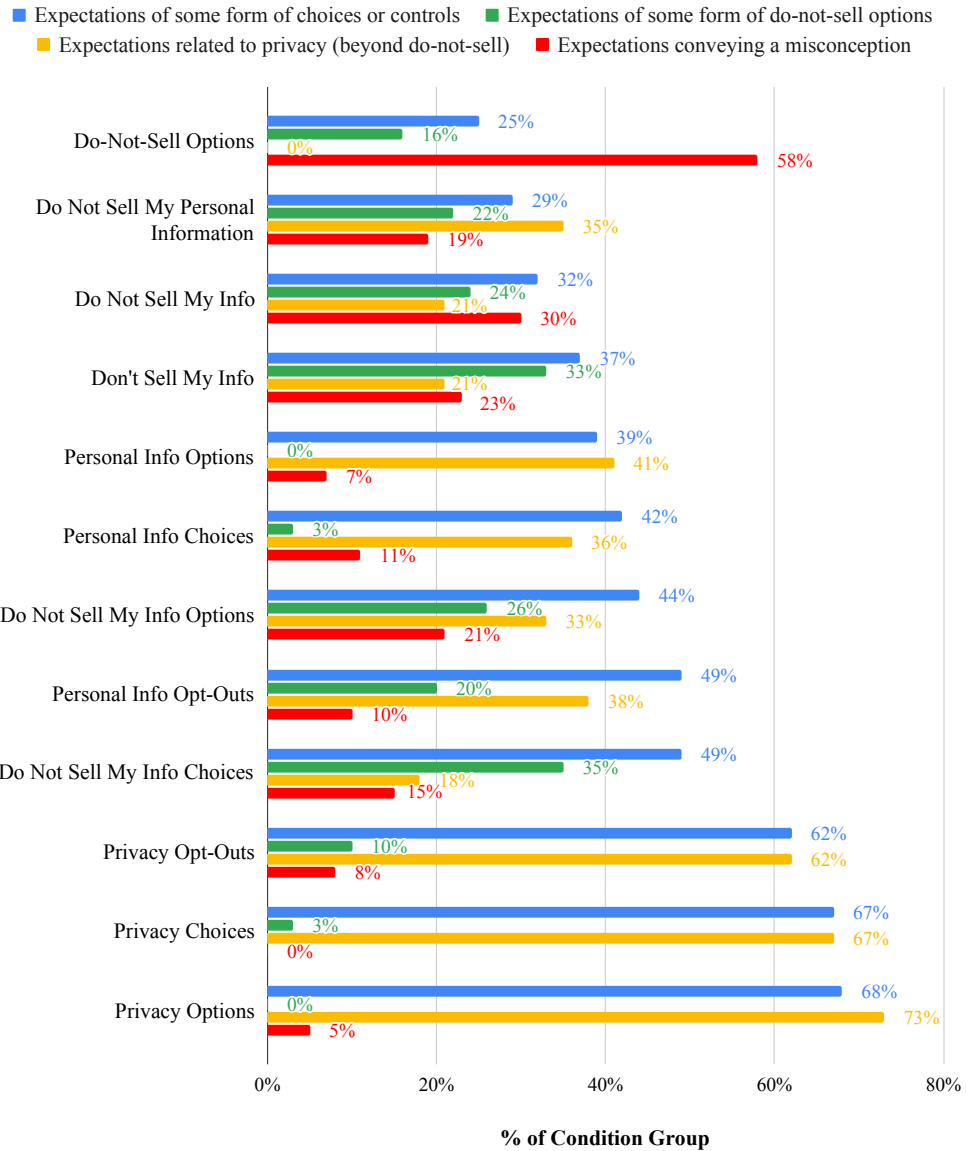


Figure 7.5: Distribution of expectations in response to “What do you think would happen if you clicked on this [link]?” in our link text pre-study.

ipants expected to see some form of choices, including those related to privacy and do-not-sell. As seen in Figure 7.5, there was a significant overall difference between conditions ( $p < .001$ ,  $V = 0.27$ ). Pairwise comparisons revealed that the only significant difference was between “Privacy Options” and “Do-Not-Sell Options” ( $p = .04$ ); 67.6% and 25.0% of participants in those conditions expressed expectations of choices, respectively. The choice-related suffixes (i.e., “choices,” “options,” or “opt-outs”) did not appear to impact participant expectations of choices, given the small differences between link texts with the same privacy-related prefix.

**CCPA link text variants led to expectations of do-not-sell choices but did not generalize.** As seen in Figure 7.5, there was a significant difference between conditions in generating expectations of do-not-sell choices ( $p < .001$ ,  $V = 0.34$ ), or something more broadly related to privacy ( $p < .001$ ,  $V = 0.42$ ). Link texts beginning with “Do Not Sell” most often led to expectations of do-not-sell choices, with “Do Not Sell My Info Choices” performing significantly better than “Personal Info Options” ( $p = .005$ ), “Privacy Options” ( $p = .008$ ), and “Privacy Choices” ( $p = .04$ ) in this regard. 35.0% of participants who saw “Do Not Sell My Info Choices” expected do-not-sell choices, whereas no participants who saw “Personal Info Options” or “Privacy Options” expressed the same expectation. However, link texts beginning with “Do Not Sell” did not effectively convey broader privacy-related information or options. “Privacy Options,” “Privacy Choices,” and “Privacy Opt-Outs” were all significantly better than “Do-Not-Sell Options” (all  $p < .001$ ), “Do Not Sell My Info Choices” ( $.0003 < p < .012$ ), “Don’t Sell My Info” ( $.001 < p < .04$ ), and “Do Not Sell My Info” ( $.002 < p < .05$ ) for this purpose. 67.1% of participants who saw a “Privacy” prefixed link text described a privacy-related expectation, compared to 21.4% who saw a “Do Not Sell” prefixed link text.

**“Privacy Choices” generated the least misconceptions.** As seen in Figure 7.5, the distribution of misconceptions were not even across conditions ( $p < .001$ ,

$V=0.39$ ). Pairwise comparisons revealed that “Privacy Choices” created significantly fewer misconceptions than “Do Not Sell My Info” ( $p=.04$ ). Among the 63 participants who saw one of the link texts beginning with “Do Not Sell,” some thought the link would lead to phishing/malware risks (16), investment advice (8), the site’s policy on selling items (8), and ads for privacy products or other services (6).

**Some link texts might apply to both privacy choices and do-not-sell choices.** In examining participants’ Likert responses to the predefined scenarios, five link texts were rated as “definitely” or “probably” likely to lead to choices about how personal information is used and shared by over three-quarters of participants. Among them, “Personal info Choices,” “Privacy Opt-Outs,” “Do Not Sell My info Options,” and “Privacy Options” were also among the top five link texts rated as “definitely” or “probably” likely to lead to the scenario describing choices about the sale of personal information. This finding suggests that these four link texts had the potential to convey both generic privacy choices and do-not-sell choices relatively well.

## 7.5 Icon-Text Combinations Evaluation

Our pre-studies suggested a need for combining icons with link texts, consistent with prior research and recommendations [153, 580]. Icons alone do not necessarily translate to correct expectations even with a certain degree of familiarity [243, 428], as reflected by our findings on DAA’s AdChoices icon. Similarly, link text alone might not stand out, but pairing the two together can attract user attention and aid comprehension [225]. We conducted a large-scale evaluation to find icon-text combinations that accurately convey privacy choices and do-not-sell choices.

### 7.5.1 Method

For icons, we selected *Stylized-Toggle* and *Slash-Dollar* since they were the most preferred for indicating privacy choices and do-not-sell choices respectively. We also included DAA’s Privacy Rights icon because of its potential for widespread adoption by DAA member companies. For link texts, we selected “Privacy Options” and “Privacy Choices” since they best generated expectations of choices/controls and expectations related to privacy (see Figure 7.5). We also included the two CCPA-mandated link texts since they conveyed do-not-sell choices well. We did not include any variants of the CCPA link texts since the choice-related suffix did not influence participant expectations. Additionally, we included “Personal Info Choices” since Likert responses to predefined scenarios suggested it worked well to communicate both do-not-sell choices and broader privacy controls.

**Study protocol.** To measure to what extent icons and link texts interact with each other in shaping participant expectations, we used a nearly full-factorial experimental design including four icon conditions and six link text conditions (23 conditions in total). The four icon conditions were DAA’s Privacy Rights icon, *Slash-Dollar*, *Stylized-Toggle*, and no icon. The six link text conditions were “Do Not Sell My Personal Information,” “Do Not Sell My Info,” “Privacy Choices,” “Privacy Options,” “Personal Info Choices,” and no link text. We excluded the combination of no icon and no link text since participants would not see any information. Our examination of icon-text combinations was exploratory—even though the pre-studies indicated that some icons and link texts perform better than others for certain purposes, interaction effects might exist between the icon and text, making it difficult to generate specific hypotheses.

We followed a between-subjects design, showing each participant an icon-text combination at random. While we presented icons and link texts with no context in




Figure 7.6: Icon and link text presented on a fictitious online shoe retailer webpage used in the icon-text combination evaluation. The icon and link text were highlighted with an orange rectangle to attract participants’ attention. Shown is the condition combining *Stylized-Toggle* (icon) and “Privacy Options” (link text).

the pre-studies, here we showed the icon and link text together on a fictitious online shoe retailer website (see Figure 7.6) to emulate how consumers might encounter them in the wild. We modified the eight scenarios for Likert questions based on common expectations uncovered in the link text pre-study; two were correct expectations, two were semi-correct expectations, and the rest were misconceptions about unwanted outcomes. We recruited 1,468 MTurk participants (roughly 64 per condition) based on heuristics that would allow us to run planned regressions [390]. The average study completion time was 4.55 minutes, and participants were compensated \$1.00 (average \$13.19/hour).

**Data analysis.** We followed the same qualitative analysis approach as in the link text pre-study ( $\kappa=0.83$ ) before using the data for quantification.<sup>7</sup> We coded partic-

<sup>7</sup>There was little diversity in responses to the question regarding the meaning of “sell” in the link text. Thus, we used percentage agreement rather than Cohen’s  $\kappa$  to measure inter-coder reliability

ipants’ responses about expectations to identify common themes, then categorized individual codes based on whether they convey the idea of choice, do-not-sell choices, privacy broadly, or misconceptions. We then ran logistic regressions using these high-level code categories as the dependent variable, the icon-text combination condition as the main independent variable, and participant demographics as control independent variables. We ran additional logistic regressions with the same independent variables on a binary variable that represented participants’ expected likelihood of each predefined scenario.<sup>8</sup> We applied Holm-Bonferroni corrections to  $p$ -values in all regressions since we conducted multiple tests without preplanned hypotheses [32]. Detailed regression results are provided in Tables E.4 and E.5.

### 7.5.2 Findings

We found significant differences between icon-text conditions in creating expectations of privacy choices or do-not-sell choices; link texts impacted participant expectations more than icons in this regard. Furthermore, *Slash-Dollar* and “Personal Info Choices” generated more misconceptions than the other icons or link texts.

**Conveying privacy choices.** Regressions of participants’ categorized open-ended expectations (Table E.4) compared how well different icon-text combinations conveyed the concepts of choice (e.g., “My choices would pop up on the screen”) and privacy (e.g., “It will enable a more private experience”). Compared to *Toggle-Privacy Options* as the baseline, combinations including the “Privacy Options” or “Privacy Choices” link text, as well as *Stylized-Toggle* by itself, performed similarly in generating privacy-related expectations; participants in all other combinations were significantly less likely to expect something related to privacy ( $0.005 < OR < 0.13$ , all  $p < .001$ ). Furthermore, participants were significantly less likely to expect some form

---

and ensured the percentage agreement was at least 75%.

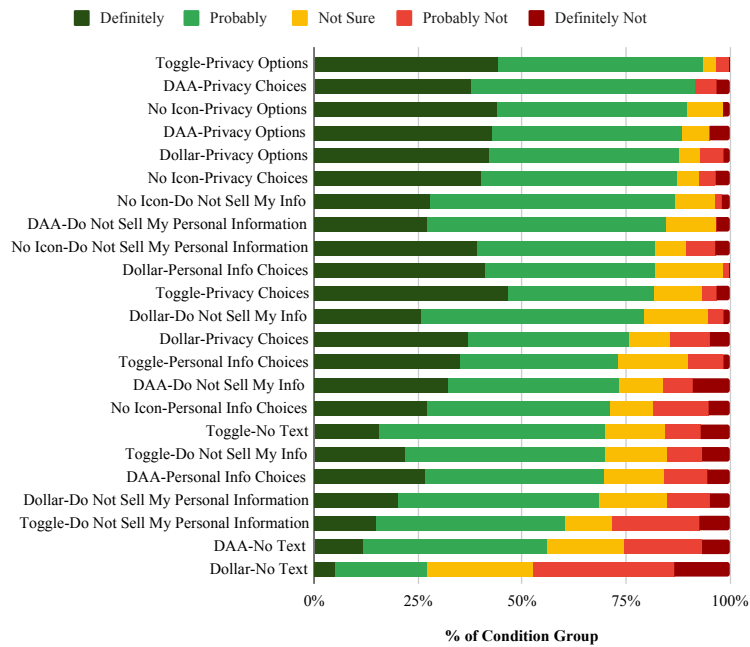
<sup>8</sup>“Definitely” and “probably” were coded as “expected” (expecting the scenario would happen). The other answer options were coded as “unexpected.”

of choice when seeing the link text “Personal Info Choices” without *Stylized-Toggle*, or *DAA/Dollar* without an accompanying link text ( $0.03 < OR < 0.27$ ,  $.001 < p < .03$ ).

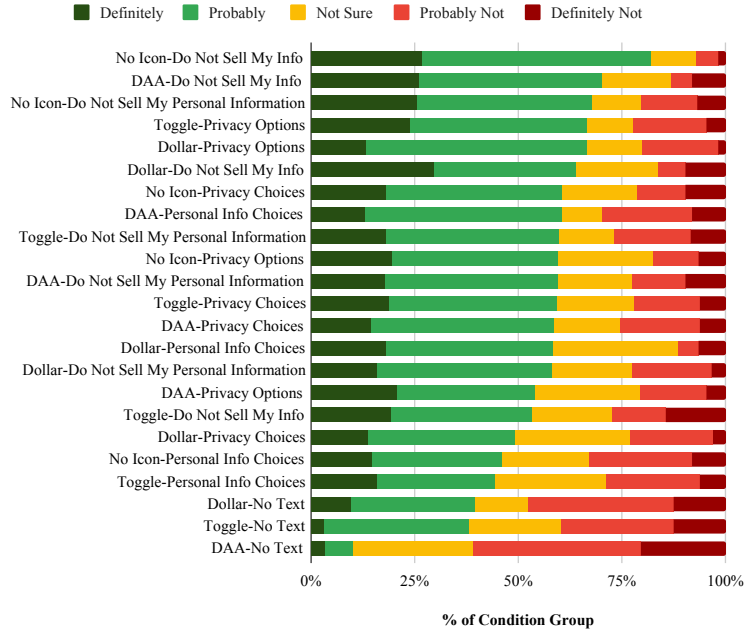
Figure 7.7a shows participants’ Likert responses to the generic privacy choice scenario. Overall, *Toggle-Privacy Options* was the best candidate for conveying “choices about how personal information is used or shared”: 93.4% of participants who saw this combination thought they would definitely or probably be led to privacy choices. Regressions of Likert responses (Table E.5) further showed that participants were significantly more likely to expect privacy choices when seeing *Toggle-Privacy Options*, compared to *Toggle-Do Not Sell My Personal Information*, *Slash-Dollar* icon alone, and *DAA* icon alone ( $.03 < OR < .17$ ,  $.001 < p < .009$ ). However, the differences between *Toggle-Privacy Options* and other conditions with “Privacy Options” as the link text were minimal and not significant in regressions. Most combinations involving the “Privacy Options” and “Privacy Choices” link texts effectively conveyed privacy choices.

**Conveying do-not-sell choices.** Regressions of participants’ categorized open-ended expectations indicated that the two CCPA-mandated link texts significantly outperformed other link texts in creating the expectation of do-not-sell choices (e.g., “It would let you opt out of them selling your information”). Relative to “Do Not Sell My Personal Information” with no icon, all conditions with the link texts “Privacy Options,” “Personal Info Choices,” and “Privacy Choices” performed significantly worse in generating expectations of do-not-sell choices ( $0.01 < OR < 0.13$ , all  $p \leq .001$ ). There were no significant differences between “Do Not Sell My Personal Information” or “Do Not Sell My Info” in this regard.

Figure 7.7b shows participants’ Likert responses to the do-not-sell choices scenario. The three conditions with the highest percentage of definitely/probably responses all included one of the CCPA link texts: *No Icon-Do Not Sell My Info* (82.1%), *DAA-Do Not Sell My Info* (70.5%), and *No Icon-Do Not Sell My Personal Information*



(a) “It [the symbol/phrase] will take me to a page with choices about how my personal information is used and shared by the website.”



(b) “It [the symbol/phrase] will take me to a page with choices about the sale of my personal information.”

Figure 7.7: Distribution of Likert responses across conditions in icon-text combinations evaluation.



(67.8%). Regressions on Likert responses further showed that *No Icon-Do Not Sell My Personal Information* performed significantly better than the *DAA* ( $OR=0.06$ ,  $p<.001$ ) and *Slash-Dollar* icons alone ( $OR=0.28$ ,  $p=.04$ ) in conveying do-not-sell choices, suggesting effectiveness of the CCPA link texts in this regard.

**Stylized-Toggle was occasionally perceived as an actual control button.** While *Toggle-Privacy Options* conveyed privacy choices well and the two CCPA mandated link texts conveyed do-not-sell choices well, putting *Stylized-Toggle* next to the CCPA link texts led to an unintended consequence. 40.0% of participants who saw *Toggle-Do Not Sell My Personal Information* expected that clicking on them would definitely or probably “give the website permission to sell my personal information.” *Stylized-Toggle* significantly increased the likelihood of this misconception compared to no icon ( $OR=5.25$ ,  $p=.02$ ) when combined with the “Do Not Sell My Personal Information” link text. This suggests that participants might perceive *Stylized-Toggle* as an actual control switch for the sale of one’s personal information on the website when the icon was next to the CCPA link texts. However, we did not observe a similar pattern in participants’ open-ended expectations—this expectation only emerged when we explicitly asked participants whether clicking the icon would give the website permission to sell their personal information, indicating a potential priming effect.

**Misconceptions with Slash-Dollar icon and “Personal Info Choices.”** Regressions of participants’ categorized open-ended expectations revealed that *Slash-Dollar* without a link text significantly increased the likelihood of misconceptions relative to *Toggle-Privacy Options* ( $OR=67.2$ ,  $p<.001$ ). Among the 371 participants who saw *Slash-Dollar*, 33 (8.9%) expressed expectations of payment options, particularly related to secure or encrypted payment (e.g., “It would present your rights to pay through secure links”). These findings indicate that the *Slash-Dollar* icon, even when paired with a link text, might be too suggestive of payment, transaction, or other financial concepts that do not concern personal information.



Figure 7.8: Our stylized toggle, OAG’s proposed opt-out button, its variant, and the iOS switch button.

Also relative to *Toggle-Privacy Options*, all conditions with “Personal Info Choices” increased the likelihood of misconceptions ( $11.9 < OR < 18.1$ ,  $.005 < p < .04$ ). Only 42.0% of participants who saw “Personal Info Choices” accurately interpreted choices as controls related to the collection, processing, and sharing of their personal data or broader privacy choices, compared to 66.5% of those who saw “Privacy Choices.” Misinterpretations of choices most frequently included profile settings related to purchasing shoes (16.7%; e.g., “Probably it would let you input your shoe size, height, favorite styles, etc. for a more customized look”). Other misconceptions included that the link would lead to choices about shoe styles or sizes available on the website (13.1%) and choices related to payment methods (1.6%). The remaining participants were either unsure about or did not specify the types of choices they expected.

## 7.6 OAG Icon Evaluation

In February 2020, the California OAG released the first set of modifications to the CCPA regulations [376] after we had shared our results with them. The proposed modifications included an opt-out icon (*CalAG-Toggle*) that was similar, but not identical to our *Stylized-Toggle* icon (see Figure 7.8).

Our icon-text combinations evaluation suggested that *Stylized-Toggle* might occasionally be perceived as an actual control switch rather than an icon when paired with the CCPA-mandated link texts. We were concerned that *CalAG-Toggle* would make this misconception even more likely for two reasons. First, *CalAG-Toggle* closely resembled the toggle switch in iOS (see Figure 7.8). By contrast, *Stylized-Toggle* used a

checkmark and “X” to visually convey the availability of options and a dividing line to differentiate it from a real toggle control. Second, *CalAG-Toggle* being in red created a potentially confusing double negative when paired with “Do Not Sell My Personal Information.” One could interpret it as either “my data is currently being sold” (because red indicates the setting “Do Not Sell My Personal Information” being off), or “my data is currently not being sold” (because red indicates the sale of personal information is prohibited). In contrast, *Stylized-Toggle* used blue, a neutral color that does not convey a particular state. We conducted a follow-up study to examine whether the style and color of *CalAG-Toggle* might diminish icon comprehension compared to *Stylized-Toggle*.

### 7.6.1 Method

We used the method already employed in our icon-text combinations evaluation to test the OAG’s proposed icon.

**Study protocol.** To understand to what extent icon style and color jointly shape participant interpretations, we implemented a full factorial design that included two color conditions (red and blue) and three style conditions (six conditions total). In addition to *Stylized-Toggle* and *CalAG-Toggle*, we created a third style condition, *CalAGX-Toggle* (see Figure 7.8), which seeks to improve the visual aesthetics of *CalAG-Toggle* by enlarging the “X” to make it visually equivalent to the circle.

As before, we used a between-subjects design, showing participants one of the six icons at random next to “Do Not Sell My Personal Information” on a fictitious online shoe retailer website. In addition to their open-ended expectations, we asked participants about the likelihood of eight scenarios occurring as a 5-point Likert item. To understand whether participants viewed the toggle as an actual control switch, we included two misconception scenarios: clicking the toggle will immediately change

the setting. We recruited 421 MTurk participants (roughly 70 per condition) for this study based on heuristics for running our planned regressions [390]. The average study completion time was 4.6 minutes, and participants were compensated \$1.00 (average \$13.04/hour).

**Data analysis.** We used the same approach in our previous studies to analyze qualitative data ( $\kappa=0.90$ ). Additionally, we grouped codes into high-level categories as to whether the code conveyed (1) any misconceptions or (2) the icon was perceived as an actual control switch. We then ran logistic regressions on these coded expectations and Likert responses (converted into a binary variable) to scenarios. We treated the interaction term [96] between icon color and style as the key independent variable and participant demographics as the control independent variables.<sup>9</sup> Detailed regression results are provided in Tables E.6 and E.7. We did not apply corrections to  $p$ -values since we ran a small number (2) of regressions with preplanned hypotheses (i.e., *Stylized-Toggle* would perform better than *CalAG/CalAGX-Toggles*) [32].

### 7.6.2 Findings

We found that *Stylized-Toggle* better conveyed do-not-sell choices than OAG’s proposed opt-out icon and its variant with fewer toggle-related misconceptions. The icon’s color (red or blue) did not significantly alter participant expectations in most cases.

***Stylized-Toggle* better created expectations of do-not-sell choices.** Figure 7.9 shows expectations of what would happen after clicking an icon. The most frequent expectation regarding *Stylized-Toggle* (29, 21.2%) was to be directed to a page with choices about the sale of personal information, a correct and desired inter-

---

<sup>9</sup>Following statistical analysis guidelines [430], for any model in which the interaction effect between style and color was not significant, we compared its performance with another model without the interaction term (i.e., style and color were examined in isolation as main effects). If the “interaction model” provided a much better fit to the data than the “main effect only model,” we report results from the first model; otherwise, we report results from the latter model.




Figure 7.9: Common expectations of what would happen after clicking based on open-ended responses in conditions with *Stylized-Toggle* ( $n=137$ ), *CalAG-Toggle* ( $n=134$ ) and *CalAGX-Toggle* ( $n=132$ ).

pretation according to the CCPA [373]. This expectation, however, was mentioned much less often in conditions involving *CalAG-Toggle* (16, 11.9%) and *CalAGX-Toggle* (10, 7.6%). The regression results on Likert responses to the do-not-sell choices scenario confirmed the significant differences: participants who saw *Stylized-Toggle* were significantly more likely to expect “it will lead me to a page where I can choose whether or not the website can sell my personal information” compared to *CalAG-Toggle* ( $OR=0.40$ ,  $p<.001$ ) and *CalAGX-Toggle* ( $OR=0.41$ ,  $p=.001$ ).

***Stylized-Toggle* led to fewer toggle-related misconceptions.** Regressions on participants’ categorized open-ended expectations revealed that *CalAG-Toggle* and *CalAGX-Toggle* were significantly more likely to generate misconceptions compared to *Stylized-Toggle* ( $OR=2.3$ ,  $OR=2.4$ ; both  $p=.003$ ). Example misconceptions include perceiving the toggle icon as an actual switch, expecting a negative outcome (e.g., more tracking), or believing that nothing would happen. Specifically, participants

who saw *CalAG-Toggle* and *CalAGX-Toggle* were significantly more likely to perceive the toggle as an actual control switch compared to *Stylized-Toggle* ( $OR=2.4, p=.003$ ;  $OR=2.4, p=.004$ ). A participant quote that conveyed this misconception is, “It would change between red and green depending on if I wanted to allow it.”

As shown in Figure 7.9, the most frequent expectation in conditions involving *CalAG-Toggle* (38, 28.4%) and *CalAGX-Toggle* (30, 22.7%) was that the icon was an actual toggle switch currently set to “Do Not Sell My Personal Information”—clicking would give the website permission to sell the user’s personal information, which is the opposite of the intended meaning. Users with this notion might avoid clicking the icon or link text for fear of losing their privacy and thus lose the opportunity to exercise the do-not-sell opt-out. In contrast, only 10 (7.3%) participants who saw *Stylized-Toggle* mentioned this misconception.

Another misconception that occurred for all three icon styles (9, 6.6% for *Stylized-Toggle*; 8, 6.0% for *CalAG-Toggle* and 12, 9.1% for *CalAGX-Toggle*) was that the website is currently selling the user’s personal information, and that clicking the toggle would stop it. Participants who held this misconception understood the icon’s purpose but misinterpreted the icon’s functionality—according to the CCPA [373], the icon should take users to respective settings but is unlikely to result in immediate changes. Regressions on the Likert responses for the respective scenario revealed interaction effects between toggle style and color; *Stylized-Toggle* in blue significantly decreased the likelihood of this misconception compared to *Stylized-Toggle* in red ( $OR=2.78, p=.006$ ) and *CalAGX-Toggle* in blue ( $OR=2.75, p=.009$ ). This misconception is not particularly problematic as it is less likely to discourage users from clicking. However, a privacy choice icon should ideally communicate its intention and function accurately.

## 7.7 Discussion

Our findings provide insights into the design and effectiveness of icons and link text in conveying privacy choices. Below we discuss our study’s limitations and outline implications for design practice and privacy regulations.

**Limitations.** Our research has several limitations. First, we recruited all participants from MTurk, and they were more educated and tech-savvy than the US general population. Nonetheless, prior work has shown that MTurkers are more demographically diverse than student samples [44] and that they offer similar responses to security and privacy surveys as traditional participant pools [413]. Second, our experiments focused on one application scenario (a fictitious online shoe retailer), which might have primed participants (e.g., to associate the dollar sign with payment and “sell” with shoe discounts). That noted, participants’ responses for our best-performing icons/link texts did not indicate that the website context affected their interpretations. Third, we measured the perception and comprehension of the icon/text by presenting them in a static screenshot; we did not measure whether participants would notice the icon/text on their own or how they would interact with the provided choices, as that was not the focus of this study. Fourth, we did not investigate accessibility issues or evaluate the use of icons with screen readers. Lastly, we did not directly compare our privacy choice icons with icons focusing on different privacy-related aspects (e.g., those that seek to visualize the concept of privacy itself or specific data practices [428]), which could be a direction for future work.

**Icons for privacy choices should be rooted in simple and familiar concepts.** *Stylized-Toggle* was participants’ favorite privacy choice icon in the pre-study and performed best in conveying privacy choices when paired with “Privacy Options” in the icon-text combinations evaluation. *Stylized-Toggle* adopts a minimalistic design

and conveys the notion of choice using a toggle — a familiar and common UI element representing the ability to make selections [35]. Nonetheless, the OAG icon evaluation shows the importance of an icon *taking inspirations from* rather than *copying* other familiar UI elements to convey the intended concept without creating confusion. Conversely, the icons that were comprehended poorly and thus excluded after the icon pre-study either attempted to convey a more abstract concept (e.g., the three icons that intended to convey “opt out”) or appeared too complicated as they combined multiple concepts (e.g., *ID-Card* and *Profile* combined elements representing “do not,” “personal information,” and “money/selling”).

Our findings suggest that an icon for privacy choices should focus on a simple and familiar concept, like choice, instead of abstract or complex concepts. For the same reason, we hypothesize that a choice-focused icon would work better than an icon attempting to convey “privacy” in indicating privacy choices. Future work is needed to validate this hypothesis, as we did not test privacy-focused icons. While prior work has proposed graphical representations of privacy—such as sunglasses, keyholes, locks, and cameras—users’ mental models of privacy are diverse and nuanced [369]. Instead, we highlighted the notion of choice through the icon and used the word “privacy” in the accompanying text. As our findings show, this effectively clarified the type of choice the icon represents.

**Icons should be accompanied by link texts.** In line with prior work suggesting that icons and text information should appear in conjunction [141, 437], our findings show that link text has a significant impact on the icon’s comprehension. Participants who saw an icon without a link text exhibited more misconceptions. Even when participants correctly recognized the concept of choice, payment, or stopping, they often failed to connect those concepts to personal information without a text description. In our icon-text combinations evaluation, conditions without link text



performed comparatively worse. These findings suggest the importance of placing a descriptive link text next to an icon to aid comprehension and reduce misconceptions. This does not undermine the merits of icons: they still complement and reinforce a text description with a visual depiction, which aids recognition [225], enables textual descriptions to be more concise [153], and conveys concepts across language barriers [428]. Any icon should come with a text description when first introduced. Once it has been broadly adopted, further testing is needed to evaluate whether the text description can be removed.

**Usability issues of the AdChoices icon persist despite wide adoption.** Even though thousands of companies have adopted DAA’s AdChoices icon, our participants struggled to recognize or interpret it accurately. In the icon pre-study, only 14% of participants recalled seeing the icon before, and even fewer correctly associated it with advertising choices. This finding echoes prior work conducted nearly a decade ago [294, 537] and shows that comprehension of this icon has not improved much since then. Coloring the AdChoices icon in green—as done by DAA’s Privacy Rights icon—did not improve comprehension either. Most participants thought of “more information” upon seeing the lowercase “i” and perceived the triangle shape as an audio/video play button. Icons have the potential to acquire a universal communicative power after being used over time, even when their constitutive elements may not be intuitive, as demonstrated by the gear icon for settings [460] or the three arrow triangle for recycling [253]. However, our findings suggest that this is not the case for the two DAA icons, as our participants rarely associated them with privacy, do-not-sell, or other types of choices. Rather than adopting a problematic icon and expecting users will understand it over time, our findings demonstrate the importance of evaluating initial icon designs with user testing to ensure the icon is comprehensible.

**Privacy choice indicators are only one component of usable privacy choices.**

Prior work has shown that users struggle to find privacy choices on websites [197, 198]. Our research seeks to help users with this discovery problem. Our proposed icon-text combinations could serve as gateways leading users to website privacy choices, especially if a standard mechanism were to be adopted and used consistently. Nevertheless, privacy choice indicators alone are insufficient. Designing indicators to help users locate privacy choices is only the first step in improving end-to-end interactions with those choices. The indicators have to compete with many other UI elements for users’ attention, and they still place the burden of accessing, learning, and exercising privacy choices on users [100, 269]. Therefore, the interfaces after clicking on an icon/link text should be designed to minimize user effort. For instance, a web form for the CCPA do-not-sell opt-out could provide a conspicuous global “opt out” option on top, with more granular options presented below [155]. For a more substantial reduction in user burden, privacy choice indicators should be part of automated mechanisms [39], such as APIs that allow users to control privacy settings across websites in their web browsers, or personalized privacy assistants that learn users’ privacy preferences and semi-automatically configure settings for them [88, 302].

**Incorporate user testing into the policy-making process.** Researchers have argued that privacy interfaces should be developed through a user-centric and iterative design process involving user testing at early stages [437, 438]. Unfortunately, most existing privacy laws do not emphasize usability or include vague requirements for presenting privacy choices in UI design. For instance, the FTC advocates that any privacy notice or choice must be “clear and prominently displayed” [554] but does not provide specific guidance on achieving this [499]. By contrast, US financial institutions’ widely adopted model privacy notice was the product of an iterative design and testing process [498]. Another positive example is the guidance for GDPR compliance

from the United Kingdom Information Commissioner’s Office [240], which included visual examples to illustrate what constitutes valid consent. The OAG’s consideration of our research in the CCPA rule-making process further demonstrates that incorporating user-tested privacy interfaces into privacy laws is both necessary and feasible. The OAG removed their proposed opt-out icon from the CCPA regulations [375] after we shared our findings with them about how their icon could generate critical misconceptions. Subsequently, the fourth set of modifications to the CCPA regulations recommended businesses use our blue stylized toggle icon to convey the presence of do-not-sell opt-outs [377].

**Mandate unified privacy choices indicators.** Even though the CCPA has an optional icon for conveying do-not-sell opt-outs [377], we consider it unrealistic and inefficient for privacy laws to require a specific icon or UI element for each privacy choice that businesses might offer, voluntarily or to comply with regulations. A web page with many different indicators is likely to confuse or overwhelm consumers [275]. Instead, mandating a standardized privacy choices indicator that directs users to all privacy choices in one place (e.g., a centralized privacy dashboard, account settings, or dedicated privacy choices page) would provide numerous benefits. For lawmakers, this approach is more economical than the significant time and resources required to develop, test, and oversee the enforcement of individual privacy choice indicators. Consumers would also appreciate a consistent and thus learnable path to navigate and exercise privacy choices [366]. Our research shows that *Stylized-Toggle* paired with the link text “Privacy Options” could be a good candidate for such a unified privacy choices indicator.

**User-tested icons should be paired with public outreach and education.** User testing can identify poor privacy choice indicators with comprehension issues, such as DAA’s icon or the OAG’s proposed icon [376], that would require significantly

more effort in consumer education. However, even for icons that have undergone rigorous testing, consumer education is still needed to raise awareness, communicate the icon's purpose, and dispel misconceptions. In our research, even the best-performing *Stylized-Toggle* icon generated misconceptions occasionally. We find little documentation on associated education or public outreach efforts for most existing privacy icons. While there have been education campaigns for the AdChoices icon in the US and Europe [493, 519], consumer awareness remains low, as we and others have found [537]. Whether this is due to ineffective messaging or insufficient reach is unclear. We suggest that effective education campaigns for new privacy choice icons should address the misconceptions uncovered in initial user testing, create an active and engaging learning experience [277], and possibly use personalized education content tailoring toward individual users' characteristics [392, 561].

## CHAPTER VIII

# Computer Security Customer Support for Survivors of Intimate Partner Violence<sup>1</sup>

The study presented in Chapter VII illuminates how to increase the adoption of privacy controls—as a specific type of online privacy-protective behaviors—for an average consumer who is motivated to protect their privacy. Nonetheless, prior work on intersectionality [101] and intersectional HCI [442] has illuminated that the concept of an “average” consumer itself is problematic, as it might leave out the needs and preferences of marginalized individuals across gender, race, class, and all sorts of other identities. When it comes to privacy-protective behaviors, specific individuals, groups, or communities face more consequential ramifications for privacy loss, and they may already opt out of such protective behaviors for valid reasons. For example, using real names is a norm on many social networking sites. Yet for racial minorities and low-income individuals, this can result in threats of harassment [513] or self-censoring reinforcing racist and sexist notions of appropriate behavior [401]. While surveillance videos and facial recognition technology can be framed as privacy invasion to all data subjects involved, it creates disproportionate harm to immigrants when police use such systems for targeting and deportation [192]. To move toward a safer,

---

<sup>1</sup>This chapter is based on: Yixin Zou, Allison McDonald, Julia Narakornpichit, Nicola Dell, Thomas Ristenpart, Kevin Roundy, Florian Schaub, & Acar Tamersoy. 2021. The Role of Computer Security Customer Support in Helping Survivors of Intimate Partner Violence. In *30th USENIX Security Symposium (USENIX Security '21)* (pp. 429–446). [601]

more accessible, and more equitable digital future, we must consider and center the needs of marginalized and otherwise vulnerable communities in developing privacy tools and advice on protective strategies [334, 461].

In this chapter, I present how to reimagine computer security customer support as a novel approach to helping survivors of intimate partner violence (IPV) as a specific at-risk population [325]. Prior research has demonstrated how to support IPV survivors via couples therapists, online resources, peer support networks, computer security clinics, and more [31, 42, 140, 166, 212, 358, 521]. Beyond these approaches, we see the promise of involving customer support agents at computer security companies in this support ecosystem for several reasons. Customers turn to these agents for a wide range of security-related problems beyond products [447], making them a possible point of contact for survivors experiencing tech-enabled IPV. Computer security companies also provide products that can catch spyware or warn users about data leaks, meaning that support agents can offer timely and impactful tech assistance when contacted by a survivor. Additionally, several leading computer security companies have expressed interest in supporting IPV survivors [479]. Despite these potentials, any targeted help customer support offers to IPV survivors should be carefully designed considering that inappropriate responses may re-traumatize survivors [308] or even escalate violence as the abuser seeks to regain control [595]. We investigate the opportunities and challenges for computer security customer support to help IPV survivors via three steps:

1. To discover if customer support agents already encountered IPV cases, we searched customer support cases at a large computer security company. Our search surfaced at least 53 tech-enabled IPV cases. In these cases, survivors described the attacks' severity and resulting distress; support agents typically focused on technical solutions without expressing empathy or awareness of IPV.
2. Having established that support agents encounter tech-enabled IPV cases, we ex-

plore how customer support could better serve IPV survivors by engaging 17 IPV professionals from five support organizations in focus groups. IPV professionals provided numerous suggestions, such as using trauma-informed language, avoiding promises to solve problems, and making referrals to external resources for support beyond the immediate technical issue.

3. To gauge the practicality of IPV professionals' recommendations, we conducted focus groups with 11 customer support practitioners from four major computer security companies. These practitioners agreed on the importance of training agents for IPV cases while noting challenges for implementation, such as frontline agents' limited capacity and uncertainty in identifying whether a customer may need IPV-related help.

Chapter 7.3 shows how the design of privacy icons could be informed by input elicited from consumers directly. This study presents a complementary approach—we did not engage with survivors directly in the research because doing so might cause unnecessary traumatization by requiring survivors to remember and recount traumatic events, and because we have identified stakeholders who have in-depth knowledge of their situations and can provide relevant insights for answering our research questions. The process described above enabled us to synthesize stakeholders' insights into novel recommendations that both cater to the needs of IPV survivors and consider the real-world constraints of customer support. By reflecting on common ground and conflicting viewpoints between IPV professionals and support practitioners, we identify a path forward for computer security companies to better address tech-enabled IPV in customer support. Computer security companies should train their support agents to be aware of IPV, the limitations of security software in combating IPV, and when and how to provide additional help. Tracking the frequency and nature of relevant cases would help companies assess their current practices, coordinate efforts between agents, and help larger companies determine whether more

costly initiatives (e.g., a specialized team) are worth pursuing. We further note the opportunity for computer security companies and IPV professionals to join forces in helping IPV survivors combat tech-enabled abuse.

## 8.1 Background

Intimate partner violence (IPV)—abuse or aggression that occurs in a romantic relationship [74]—is a pervasive societal phenomenon that causes physical and psychological harm to victims [399]. In the US, nearly 20 people per minute are physically abused by an intimate partner, which equates to more than 10 million abuse victims annually [509]. A growing body of literature on tech-enabled IPV has documented the many ways abusers maliciously use technology [165, 167, 215, 291, 292, 296, 305, 478, 585] and how IPV survivors struggle to protect their privacy and security [122, 127, 168, 325]. The complex socio-technical and relational factors due to the intimacy of the relationship differentiate tech-enabled IPV from abuse in other contexts such as online harassment [513], doxxing [541], and cyberbullying [569].

**Malicious apps in IPV.** One way that tech-enabled IPV occurs is through surveillance apps installed on survivors’ devices [76, 168, 209, 325]. mSpy, one of the largest spyware vendors, allegedly had around two million users as of 2014 [95]. In NortonLifeLock’s 2020 survey, 10% of respondents admitted using an app to monitor a former or current partner’s messages, calls, emails, or photos [518].

Most spyware apps are, in fact, dual-use, i.e., they have a legitimate purpose (e.g., “Find My Phone” for anti-theft) but can be repurposed for spying on an intimate partner [76]. Growing awareness of the spyware problem has motivated improved detection features and related research [164, 431]. Security companies have also joined forces with one another and with IPV advocacy groups through the Coalition Against Stalkerware [97]. Regulators are also strengthening their oversight of spyware, such



as the FTC’s settlement with Retina-X in 2019 [504].

**Protecting survivors against tech-enabled IPV.** Survivors of technology-enabled intimate partner violence face unique threats and challenges in protecting their digital privacy and security. Survivors may still share residences, parental responsibilities, and social networks with their abusers, meaning that abusers could have physical access to the survivor and their devices while possessing an intimate knowledge of the survivor’s routines and preferences [168, 291]. Using privacy controls to block abusers on social media or phones could escalate the abuse from online harassment to physical violence [167, 168]. Stopping using specific technologies could improve the survivor’s privacy but simultaneously prevents evidence collection for legal procedures [168].

In addition to spyware detection, prior work has proposed apps that help IPV survivors erase their browser history [140], record evidence of abuse [31], and engage in safety planning [358]. However, few of these efforts have led to practical and widely-used solutions. Support organizations such as the National Network to End Domestic Violence (NNEDV) [517] and Safe Horizon [224] have provided tech-focused resources for survivors. Still, these resources are often outdated or lack detailed guidance [168]. Computer security clinics are a recent approach for helping IPV survivors through one-on-one consultations with trained technologists, who analyze survivors’ digital assets and provide personalized advice on resisting tech-enabled attacks [84, 191]. However, these clinics currently have limited geographic reach and predominantly serve survivors who have already physically separated from their abusive partners [166, 212]. Our focus—computer security companies’ customer support—has the potential to reach a broader audience, but this approach requires careful attention to the nuances and unique risks in IPV so as not to inadvertently exacerbate harm.

**Customer support.** Customer support exists in many companies to provide information and assistance for products or services [161], with a crucial role in strengthening customer engagement and increasing revenue [125, 179, 182]. According to the SERVQUAL measurement [384], the quality of customer support is evaluated by its reliability, assurance, tangibles, empathy, and responsiveness [384]. To create a positive experience, support agents need to make customers feel heard, respected, and appreciated [34]. For customer support in information technology in particular, agents tend to be technical thinkers with limited soft skills, so it is critical to train agents to use phrases that build rapport, show empathy, and communicate effectively [571]. We investigate how computer security customer support might better support customers experiencing tech-enabled IPV.

**Interacting with IPV survivors.** Training materials for IPV professionals [89, 367] point to specific considerations needed in interacting with IPV survivors beyond demonstrating empathy and good communication skills. These materials typically explain what domestic violence means, who the perpetrators are, and how to provide support such as safety planning, education, advocacy, and referrals. Some materials further emphasize *empowerment*—supporting survivors in finding their inner strength [271], and *trauma-informed communication*—taking into account that clients likely have a history of trauma [367]. Others discuss secondary trauma on IPV professionals and respective coping strategies, acknowledging that bearing witness to abuse and trauma is emotionally taxing [464].

However, most training for IPV professionals does not cover tech-enabled abuse [168]. IPV professionals currently do not have best practices for discovering, assessing, and mitigating tech issues [168]. Meanwhile, support agents at computer security companies provide complementary strength in delivering tech-related assistance but may not be sensitive to the nuances of IPV. Our work synthesizes perspectives from IPV

professionals and support practitioners to identify how computer security customer support could help IPV survivors and how this help should be provided.

## 8.2 Preliminary Analysis of Support Cases

As a starting point, we sought to discover if IPV survivors experiencing tech abuse seek assistance from computer security companies' customer support and what those interactions look like. We performed keyword searches on customer support records from a major computer security company and surfaced 53 cases in which the customer identified their attacker as an intimate partner. However, typical reactions from support agents indicated they did not recognize the complexity of IPV beyond tech issues.

### 8.2.1 Method

The computer security company we worked with provides customer support via phone, interactive chat, and self-service (e.g., FAQs, forums, tutorials). We analyzed chat records since they are anonymized, searchable, and represent a decent portion (40%) of support requests. All cases include customer-provided problem descriptions (255 characters maximum). Some cases also include chat transcripts and agents' notes.

To identify relevant cases, we searched a database of 18,900 customer support cases from January 1, 2017 to June 1, 2019, using keywords indicative of abusive relationships (e.g., variants of “creepy,” “restraining,” “ex-\*) and IPV-related attacks (e.g., “spy,” “stalk,” “monitor,” “violen\*”) [165, 167, 478, 585]. Our search surfaced 1,083 cases. We then went through these cases to exclude those irrelevant to our interest, such as users reporting generic malware or complaining about false positive warnings. This filtering resulted in 273 instances of victim-reported interpersonal attacks. Three researchers jointly coded the 273 customer-provided problem descrip-

tions for the attacker’s relationship to the victim (Fleiss  $\kappa=1.00$ ). In 53 cases, the attacker was clearly identified as an intimate partner (e.g., “my boyfriend” or “my ex-boyfriend”). The researchers also coded other dimensions related to the attack, such as attack type ( $\kappa=0.75$ ), attack mechanism ( $\kappa=0.59$ ),<sup>2</sup> and intimate partner relationship stage as defined by Matthews et al. [325] ( $\kappa=0.82$ ).

We focus on analyzing the 53 cases clearly identified as tech-enabled IPV. For these cases, we summarize the attacks based on the customer’s problem description. When available, we thematically analyze the agent-customer interaction based on chat transcripts and agents’ notes. However, we note that other cases in which the attacker’s identity was not specified (e.g., “*I’m being stalked*”) may still be IPV related. We also note that we did *not* intend to measure the prevalence of IPV cases within customer support data. Instead, our goal was to understand *if* such cases occur and qualitatively uncover scenarios customer support agents are dealing with.

**Ethical Considerations.** Our study received IRB approval. By agreeing to the company’s privacy policy, which is prominently featured when a chat session starts, customers consented to chat recordings and messages as examples of diagnostic information being shared with third parties. We acknowledge that this consent process is not ideal since prior research has documented it well that most consumers tend to consent to privacy policies without reading through the details [67, 370]. However, we were unable to seek explicit consent from individuals whose cases were analyzed or quoted in our study since the chat data was already anonymized. As additional measures to ensure participants’ privacy, we asked a company employee to review all chat records to verify anonymity and remove references to unique circumstances before handing the records to us.

---

<sup>2</sup>We did not pursue high inter-rater reliability for this dimension since multiple attack mechanisms were frequently at play.

### 8.2.2 Results

**Diverse attack types.** The most common attack types among the 53 cases were spying or surveillance of the survivor (23), account or device compromise, such as changing the account password to lock the survivor out (17), and interference with the account or device usage (12). Less frequently mentioned attacks included harassment (5), spoofing (2), financial fraud (2), phishing (2), and modifying content on the survivor’s account or device (2). Installing spyware or other malicious apps on the survivor’s device was the primary attack mechanism (23), though account compromise based on knowledge of credentials (10) and physical ownership-based attacks (6) also occurred. These attack types and mechanisms generally align with Freed et al.’s taxonomy of stalkerware [167].

**Attacks’ repercussions on survivors.** In 49 of the 53 cases, the survivor mentioned they were in the process of separation or had separated from their abusive partner. Though the survivor’s risks might appear lower for attacks after separation, there are signs of distress and helplessness with references to violence, ruined lives, and even contemplation of suicide. In 13 cases, the survivor mentioned multiple types of attack at play, e.g., *“my husband’s hobby is to hack my home network and try to track my email, calls, and whereabouts.”* The attacks caused apparent emotional distress to the survivor, e.g., *“I know that my ex-boyfriend is stalking me through my phone...He has ruined my life.”* Another survivor wrote: *“I found out my soon-to-be ex-wife hired a professional hacker to really mess me and my folk’s computers and phones up...just had a heart attack from the stress.”* In six cases, the survivor described that their abuser was *“a computer expert,” “worked at a top IT firm,” “can remote access most computers,”* or in other terms that indicate the abuser’s tech-savviness. Even though most attacks in IPV are technologically unsophisticated [167], survivors in these cases expressed fear and helplessness, especially when their own computing

skills were limited.

**Support agents focused on technical issues** Our thematic analysis of chat transcripts revealed that support agents were not well prepared for these tech-enabled IPV cases. Figure 8.1 shows three representative agent-customer interactions. A typical agent reaction was to scan the survivor’s device for malicious applications and launch a remote assistance session to investigate the problem further if needed. Agents might also receive out-of-scope requests. For instance, one survivor asked *“I am blocking my wife’s/future ex-wife’s messages. Is there any way I can have these sent to my email for presentation to my attorney?”* In these cases, the agent would refer the survivor to more experienced experts on the team, device manufacturers, or operating system vendors. For survivors who described traumatic attacks, agents generally expressed confidence in resolving the technical issue but rarely used empathetic language. When survivors suspected hacking, spyware, or malware, agents would express reassurance that the company’s security product would protect them well. Such claims might not be valid, as there were cases in which the survivor expressed skepticism or appeared to have contacted customer support multiple times.

### 8.3 Focus Groups with IPV Professionals

Our preliminary analysis of customer support cases indicates that agents received requests for help from IPV survivors but might not be sufficiently prepared to handle them. To explore how to improve customer support to serve survivors’ needs, we sought input from IPV professionals with extensive training and experience working with survivors. We conducted five focus groups (2–5 participants each) with 17 IPV professionals between November 2019 and February 2020. In line with prior work [168, 358], we chose focus groups to create a setting where participants can listen to others and feel comfortable expressing their opinions. The Institutional Review Boards at

Scenario A	Scenario B	Scenario C
<p><b>C:</b> My ex-husband hacked my phone. He keeps getting my account passwords. I have changed phones so many times and got a restraining order, but he still managed to do this. Please help.</p> <p><b>S:</b> Thank you for contacting us. I'm happy to help resolve the issue. I would recommend installing [product], which should prevent malware from being installed if you get a new phone.</p> <p><b>C:</b> I have already spent a lot of money trying to fix this problem and talked to my phone provider. No one has been able to fix it. I can't spend more time and effort on this. Please help, this problem has almost driven me to commit suicide.</p> <p><b>S:</b> Please do not worry about these devices if you have [product] installed. We will do everything we can to help you further.</p>	<p><b>C:</b> My husband is violent and keeps hacking my email and watching everything I do online. Could you help me get him off my network?</p> <p><b>S:</b> I'm sorry to hear what you are going through. How do you think he is watching your activity?</p> <p><b>C:</b> He doesn't live with me anymore, but he broke into my apartment last month and I think he hacked my router. I am afraid he can see everything I am doing.</p>	<p><b>C:</b> My ex used to share my computer and installed some programs, but I think she installed spyware. I think she is remotely accessing my computer. Can you help?</p> <p><b>S:</b> Thank you for contacting [company name], I will be happy to assist you. Let's set up a remote connection so I can scan your device for malware. Please visit this link: &lt;link&gt;.</p> <p><b>C:</b> I can't open it. My computer just restarted. I think she is monitoring this chat and trying to stop me from getting help.</p>

Figure 8.1: Portions of three representative customer support chats from our dataset (“C” is customer, “S” support agent). We removed specific details and slightly altered wording to protect anonymity.

Cornell Tech and the University of Michigan approved the study protocol.

### 8.3.1 Method

**Recruitment.** Our 17 participants came from five organizations that provide free and confidential civil, legal, counseling, and support services for IPV survivors in two major US cities. One organization primarily serves human trafficking survivors, but participants noted that many clients experienced sex trafficking by intimate partners and IPV. We explained our study to each organization’s director, who then advertised our study to their staff and assisted with recruitment and scheduling. Our participants included seven program directors or managers, five attorneys or paralegals, two administrative assistants, two counselors, and one case manager. Participants had one to 30 years of experience; 12 participants had worked in this sector for more than five years.

**Study protocol.** We conducted in-person focus groups at participants’ organizations. Sessions lasted one hour on average and were audio-recorded with participant consent. We did not compensate participants as the organization directors did not deem it necessary. We prepared a list of prompts to guide the discussion (see Appendix F.1) and encouraged participants to comment or ask questions at any time. We started by asking about participants’ experience working with IPV survivors, especially regarding tech-enabled abuse. Next, we presented the three scenarios in Figure 8.1, which represented common attack types (see Section 8.2) and reflected explicit threats from an intimate partner. After participants read the scenarios, we asked them to share their perspectives and recommendations for support agents’ role in providing advice, making referrals, and more. We also probed participants to consider adversarial situations in which the abuser might monitor the chat or impersonate the survivor.

**Qualitative data analysis** We used inductive coding [434] to analyze focus group transcripts. Two researchers independently reviewed and coded the first three transcripts before discussing discrepancies. After agreeing on a consistent codebook, they applied it independently to the remaining transcripts and added new codes that emerged. They then jointly reviewed all coded transcripts, reconciled disagreements, and clustered codes into themes. Our final codebook has 71 codes in five themes: (1) positive aspects of current practices, (2) negative aspects of current practices, (3) advice from IPV professionals, (4) challenges of given advice, and (5) other comments. We do not report inter-rater reliability since all data was double-coded and disagreements were reconciled [336]. Given the qualitative nature of our findings, we do not report the frequency of individual statements. Further, while we mention how many groups a topic came up with, we cannot assume consensus within the group, as participants may have silently agreed or disagreed with each other.



### 8.3.2 Suggestions for Interacting with Survivors

IPV professionals provided three key recommendations for interacting with customers who might be IPV survivors: using trauma-informed language, asking follow-up questions without judging, and avoiding overpromising.

**Use trauma-informed language.** IPV professionals reacted strongly to the language support agents used to respond to survivors' concerns. Four groups said that Scenario A included dismissive language that might mislead or re-traumatize the survivor. Professionals took issue with the phrase "please do not worry about these devices if you have [product] installed," noting that it is highly inappropriate to focus on the security software's functionality right after the survivor mentioned a restraining order on their abuser and suicidal thoughts. An attorney discussed how the agent's language might arise from the goal of making customers happy in their regular work:

*"I understand that the role of customer support is to make their customer feel better. But this is just a space where . . . they have a limited capacity to make [the survivor] feel better . . . I think the goal should be to hear and be honest about the limitations of what [product] can or cannot do in those moments."* (P11, attorney)

All groups highlighted the importance of trauma-informed language, a common element in their training and practices [367, 440] and in other fields serving trauma victims [3, 315]. Being trauma-informed means accounting for the pervasive nature of trauma and avoiding unintentional re-traumatization through careful language and interactions [60]. A counselor explained how to provide trauma-informed responses in customer support:

*"Just acknowledge that 'this is scary' and that 'it sounds like you're having a really hard time.' Even just the smallest little pieces of empathetic*

*language so [the survivor] knows that [the agent] is actually hearing them . . . and expressing concern for them.” (P2, counselor)*

Professionals provided concrete suggestions for training support agents to use trauma-informed language, such as using the Forensic Experiential Trauma Interview (FETI) [94], which is aimed at law enforcement but makes analogies for people who do not typically work with survivors. Another suggestion was incorporating trauma-informed responses into scripts so agents do not need to figure out what to say on the fly. Nonetheless, scripts alone were insufficient: part of the training should educate agents about the complexities of IPV and why trauma-informed responses are needed.

One group highlighted the need to address support agents’ own trauma. Due to the prevalence of IPV [518, 520], some agents may be survivors themselves. Agents may also feel distressed and helpless hearing survivors’ experiences:

*“Some of these calls will be really harmful to the people who receive them. They’ll be really traumatized by these experiences . . . Any company that’s recognizing that their front line employees are experiencing these phone calls needs to think about how to support employees through secondary trauma issues and process it.” (P13, attorney)*

**Ask follow-up questions without judgment.** Four groups suggested that agents could ask follow-up questions to surface additional risks that should be considered when giving advice and ensure that the customer is safe to receive and act on advice. The question, “How do you think he is watching your activity?” in Scenario B was identified as a good example: it is open-ended, non-judgmental, and might help the agent better diagnose the case by encouraging the customer to speculate about the source of the problem.

Professionals also provided additional examples of appropriate follow-up questions, e.g., asking about the customer’s immediate concern in the form of “What are you

most concerned about?” or “What is your goal of calling me today?” Professionals explained that such questions do not assume the survivor’s needs and might help identify other risks that warrant attention, such as risks related to immigration status, health, or economic situation. Another follow-up question could be, “What have you already tried?” to facilitate the troubleshooting process and make the conversation more productive since the survivor likely tried to address the problem before reaching out for help.

Professionals further discussed the need to account for the possibility that the abuser might be physically or remotely accessing the survivor’s devices and accounts. Four groups recommended a safety check-in with the customer by asking, “Do you think you’re on a secure line?” or, “Are you safe now?” If the response is no or unsure, the agent should offer to call back or initiate a chat from a different device, such as a friend’s phone. Three groups also recommended verifying the customer’s identity in case the abuser is impersonating the survivor to gain access to the security software or other accounts. The agent could verify the customer’s email, phone number, or account history (e.g., “I see in our records someone just called about this account. Is that you?”).

Nevertheless, professionals acknowledged that it is challenging to handle situations where the abuser is present: identity verification takes practice and can still go wrong; giving advice, such as switching to a different phone, might tip off the abuser. Yet, professionals noted that the risk does not undermine the importance of support agents providing necessary help and information. As a program director explained:

*“[The survivor] had to disclose the problem to begin with, so [the abuser] has already [been] tipped off. But . . . that’s why we need to connect [the survivor] to a safety clinic. It’s really tricky when the phone is the only way to communicate.”* (P1, program director)

**Avoid overpromising.** All groups took issue with the phrase, “I’m happy to help resolve the issue” in Scenario A, saying that “resolve” is an overpromise because one chat session is unlikely to solve the physical and digital complexities survivors face in IPV [167, 168, 325]. From their perspectives, agents might promise to solve problems instinctively or do so to comply with company policy. Yet, for IPV survivors, many of whom have traumatic experiences, such promises could be frustrating and misleading. Professionals noted that a better approach is being honest about the security software’s limitations while still providing support, such as saying, “I will help you as much as I can in this call today, and whatever we don’t take care of, we might have to keep working on it.” Doing so does not necessarily contradict the agent’s responsibility to help customers and does not let them down. As a legal advocate said:

*“[The agent] can still support the survivor while giving them a response that they don’t want . . . But do it in a way that still lets [the survivor] know that they are there, they understand, they are validating their experience, and they are . . . empathetic. They can still give [the survivor] the bad news without completely turning them down.”* (P16, legal advocate)

### 8.3.3 Responsibilities of Customer Support

Customer support’s typical role is to provide technical assistance related to the company’s products and services, engage with customers, and drive more sales. Professionals stressed that agents should only offer advice on topics within their expertise and refer the customer elsewhere for issues beyond that. Nevertheless, for IPV survivors, support agents could do more than troubleshooting technical problems or recommending the company’s products. For instance, agents could discuss the potential consequences of advice they give or share basic technology safety tips.

**Avoid making advice too product-oriented.** In Scenario A, the agent recommended installing one of the company’s software products. Professionals commented that this behavior is understandable, given that the agent represents the company and that the product might be helpful. Nevertheless, the line might read too product-oriented and convey the impression that the agent was following a script and making a sales pitch without actively listening. To make a product recommendation more helpful, a counselor suggested explaining how and why the software is going to help in the survivor’s situation:

*“[The survivor] didn’t call for that product. She called with a problem. [The agent] never explained how their product was going to solve the problem . . . So please give more explanation about that.”* (P2, counselor)

**Discuss consequences of given advice.** While professionals agreed that support agents could provide IPV survivors with vital assistance, they emphasized the caution required in providing such assistance. One suggestion was explaining potential negative consequences that might result from the advice to prompt the survivor to think about safety issues. A program director gave an example:

*“Ask [the survivor]: if this app were to be uninstalled, how would it affect you? . . . Do you use it often? Do you rely on it? Does the [abuser] have access to it? Will they notice if it’s uninstalled?”* (P9, program director)

However, another professional mentioned a potential issue with discussing negative consequences—it might trigger additional questions from the survivor that catch the agent off guard, which points to the importance of external referrals:

*“I feel it’s like a slippery slope because [the agents] are not domestic violence advocates. And so [the survivor] is going to just be like, ‘What do you mean? What do you think will happen?’ And they’re never going to be able to answer those questions.”* (P1, program director)

**Share resources for tech safety.** Three groups suggested sharing resources that might improve the survivor’s digital security and privacy, such as adjusting privacy settings on social media and using strong passwords. Prior work with IPV survivors [166] also indicates that survivors have many general tech safety questions and desire credible information on this topic, validating the need for sharing tech safety resources. One group noted that in addition to sharing existing resources, computer security companies could utilize their expertise to provide self-created content on tech safety. Such content could appear on the company website’s FAQ or “Contact Us” page to put such resources into a survivor’s pathway of seeking help.<sup>3</sup> Tech safety resources should be written in plain language and provided with non-technical support resources, such as information on domestic violence shelters, to ensure relevant resources are available in one place.

**Have a specialized team.** Three groups suggested a specialized team within the company’s customer support division for handling IPV cases transferred from frontline agents. A specialized team resolves the dilemma for frontline agents who are often pressured or incentivized to finish cases fast, whereas dealing with complex issues like tech-enabled IPV requires extensive efforts and patience. It could also reduce the company’s workload in training, as training a core group of specialists would be much easier than training all frontline agents. One group further noted that the company could track the number of potential IPV cases encountered by frontline agents to understand the issue’s prevalence and decide whether investing in a specialized team is warranted.

Given the potential that a survivor might face imminent danger, professionals emphasized that frontline agents should always conduct a safety check-in (e.g., “Do you think you can stay on the line with us?”) to determine whether the survivor

---

<sup>3</sup>Some security companies are already doing this (e.g., [239, 313]), although most content does not explicitly address IPV or tech-enabled abuse.

could tolerate a transfer to the specialized team. Additionally, many survivors might have experienced prior failures in obtaining assistance and could easily get frustrated when being transferred. A counselor gave an example of appropriate language taking this into account:

*“We, as a company, remain interested and committed to trying to help you and talk to you . . . But if you can hold on a minute, I’m going to get you connected with a colleague who knows our product but can [also] talk to you about some of these [safety] issues.”* (P2, counselor)

Without the pressure of completing cases in a limited time, professionals envisioned that these specialist agents could even build long-term relationships with survivors, such as following up with them if the problem does not get fixed in the initial chat session. Importantly, four groups cautioned that support agents should never provide advice beyond their expertise and training. Examples of out-of-scope advice included comprehensive IPV-related counseling, safety planning (e.g., maintaining physical safety in leaving an abuser), and legal advice. While professionals identified a handful of follow-up questions to ask or common advice to give, the extent to which agents can help customers think through potential consequences is dependent on the situation and individual needs. If the survivor needs support the agent cannot provide, the agent should refer them to external IPV professionals with expertise in the social, legal, or health aspects of IPV.

### **8.3.4 Suggestions for External Referrals**

In addition to technology-related problems, many IPV survivors are concurrently dealing with medical, legal, financial, and other complex problems [168]. With this in mind, professionals discussed the need to refer survivors to external support, including IPV advocates, legal experts, and law enforcement. We now discuss *where*, *when*, and *how* to refer survivors to external organizations.

**Where to refer.** All groups stressed the need to refer survivors to relevant hotlines (e.g., the National Domestic Violence Hotline, Safe Horizon, and Crisis Text Line) and organizations that provide resources for survivors (e.g., NNEDV). Four groups also suggested referrals to 911 or the National Suicide Prevention Lifeline if there are cues of physical danger or suicide contemplation. Two groups mentioned that survivors might also benefit from referrals to legal resources (e.g., WomensLaw.org) or sex trafficking resources (e.g., the National Human Trafficking Hotline).

One challenge in making referrals is that the resources available differ substantially across local, state, national, and global boundaries. In the US, there is the National Domestic Violence Hotline, but each state also has its own hotline. The referrals get more complicated for companies that operate globally. However, an attorney argued that figuring out the exact resource for referrals is not that important as long as any referral is given, as staff at these hotlines and organizations are sufficiently trained to refer onward if they are not in a position to help:

*“Most of these places that you call can handle any of these intakes and they’ll figure out the way ... If you get the company committed to giving out a suicide hotline and a collection of these numbers, honestly the distinctions don’t matter.”* (P13, attorney)

**When to refer.** All groups suggested monitoring for “red flag words” in the conversation to determine when an external referral is needed. For example, indications of adverse behaviors such as spying, stalking, and violence from an intimate partner generally point to the need for IPV-related resources. “This problem has almost driven me to commit suicide” in Scenario A or other indications of threatened physical safety are clear red flags that call for 911 and suicide prevention resources.

Three groups suggested that agents should be trained to understand and identify common types of tech abuse. One resource that could be part of such training is



NNEDV's Power and Control Wheel on Technology and Abuse [516]. For situations without clear indications of IPV (e.g., the customer mentions abusive behaviors but does not say they come from an intimate partner), professionals believed the agent should still share relevant resources not limited to IPV. An attorney gave an example:

*“If it’s a stranger, there would have to be some concerning conduct . . . So if [a customer is] calling, maybe it’s because [they] are getting creepy spoofed messages from an account [they] don’t recognize. Well that’s already raising flags, right?”* (P5, attorney)

In Scenario C, in which the customer believed their ex was monitoring the chat to prevent them from getting help, one group pointed out this was an example of controlling behavior that still warrants attention, as IPV can occur via coercive control without physical violence [103]. Professionals across all groups advocated for making referrals without worrying too much about verifying whether the customer is experiencing IPV: a referral is better than no referral, because not providing resources to someone in need can do more harm than providing resources to someone who does not need them. A paralegal explained:

*Let’s say [the customer] is actually safe . . . They Google the number, they see it’s . . . the domestic violence helpline. They’re going to be, ‘whatever, I’m not calling that’ . . . But for the person who really has the need, if they want it, they will follow up on that phone call.”* (P12, paralegal)

**How to refer.** Four groups mentioned that an important principle in making referrals was respecting the survivor’s decision-making agency. The idea that survivors should be able to decide if and how they want to utilize resources is common in IPV professionals’ training [358, 367]. As an example, an administrative assistant explained that agents should always ask survivors whether and how they would like to be transferred to external resources:

*“Maybe this survivor is not in a private space to have that conversation . . . Maybe transferring them directly to a domestic violence agency [is] too overwhelming at that moment and not what they are looking for . . . Give them resources to explore it on their own.”* (P15, admin. asst.)

Three groups discussed the potential harms of labeling the customer as an IPV survivor. Here, the harm does not come from the action of providing IPV-related resources but rather from the repeated mentioning of words like abuse, domestic violence, or survivor. Customers who are not survivors might find it offensive, and customers who are survivors might not be ready to be identified. Instead, agents should use the same language that the customer uses, e.g., if the customer describes abusive behaviors from an ex-partner, the agent should also use “ex-partner” in referring to the abuser. As a counselor described:

*“If [my clients] say something is going on, I am not going to say ‘you are a survivor of domestic violence’ . . . You don’t want them to think that the person has assumed and knows what you are . . . You want to give them the opportunity to call it in whatever ways they want.”* (P14, counselor)

## **8.4 Focus Groups with Support Practitioners**

IPV professionals provided many suggestions for how customer support could help IPV survivors. To assess the practicality of these suggestions, we conducted four focus groups (2-4 participants each) with 11 customer support practitioners between April and June 2020. We sought to learn how attuned support practitioners are to tech-enabled IPV and their opinions on the suggestions’ feasibility and implementation challenges. This study also received IRB approval.

### 8.4.1 Method

**Recruitment.** Our participants came from four major security companies affiliated with the Coalition Against Stalkerware [479].<sup>4</sup> All four companies offer consumer- and business-facing security software and services to millions of customers. Each had customer support divisions to answer product-related questions and concerns. Six participants were customer support managers; the other five worked as support agent, training consultant, liaison between support and engineering, content writer, and malware researcher. Eight participants had extensive experience in the industry (9+ years of experience).

**Study protocol.** We conducted focus groups remotely over video chat since participants were geographically dispersed. We synthesized our results from Section 8.3 into a presentation in five parts to guide the discussion (see Appendix F.2). In Part 1, we explored participants' backgrounds, their company's customer support organizational structures, and metrics for measuring success. We also asked if participants had encountered tech-enabled IPV cases in their roles (either personally or through a team member) and any company initiatives to support IPV survivors. In Parts 2–4, we presented summaries of IPV professionals' suggestions: how to interact with survivors (Section 8.3.2), the responsibilities of support agents (Section 8.3.3), and how to refer survivors (Section 8.3.4). In Part 5, we elicited feedback on IPV professionals' suggestions for training (e.g., common types of tech abuse, trauma-informed responses, secondary trauma). Each part contained specific examples and quotes from our focus groups with IPV professionals. We invited participants to freely share their reactions and thoughts on the value, cost, feasibility, and challenges of implementing the suggestions.

---

<sup>4</sup>Specific companies are not named per agreement with our participants.

**Qualitative data analysis.** We used inductive coding [434] to analyze focus group transcripts. Our coding process was similar to Section 8.3.1: two researchers independently coded two transcripts, compared differences, created a consistent codebook, applied the codebook to the remaining transcripts separately, and reviewed all coded transcripts together. Our final codebook has 54 codes in six themes: (1) values of the company or participants, (2) existing practices, (3) metrics for measuring agents’ success, (4) opinions on IPV professionals’ suggestions, (5) challenges of implementation, and (6) ideas for supporting IPV survivors.

#### 8.4.2 Well-Received Suggestions

Practitioners agreed on the importance of assisting survivors and training frontline agents. Practitioners also endorsed the idea of providing and sharing tech safety resources, which they had been doing to some extent.

**Existing practices to support survivors.** Practitioners across all groups reported having received tech-enabled IPV cases in their roles, confirming the need for customer support to assist survivors. Although no company had a dedicated protocol to respond to IPV cases, each company had a specialized team for handling complex cases transferred from frontline agents, such as complex malware-related issues that would demand more time and expertise. S9,<sup>5</sup> a customer support specialist, mentioned sharing a license key of their product’s premium version with customers experiencing IPV. Agents also ask each other for advice when encountering unfamiliar cases:

*“Even though we don’t have formal training or content around such issues  
... out of experience, we do share some information on how we can handle*

---

<sup>5</sup>We use “S[number]” as identifiers for support practitioners to differentiate them from IPV professionals.

*such customers . . . higher tier agents actually talk to [frontline agents] and guide them appropriately.” (S1, training consultant)*

**Train agents on tech-enabled IPV.** Three groups acknowledged the importance of training agents for cases of tech-enabled IPV, recognizing that these cases were happening and that agents did not have an established protocol to follow. A manager noted that even if a specialized team exists, frontline agents still need to receive training that covers the complexity of IPV and the role of technology in facilitating abuse:

*“We [can have] a specialized team which . . . knows exactly about next steps. But the first contact is regular support agents, who have no dedicated training on this, and therefore there must be at least the awareness that these kind of privacy issues, stalkerware . . . could be on the device.” (S10, manager)*

Another manager liked the idea of embedding trauma-informed responses in training, noting that such responses would benefit all customers, not just IPV survivors:

*“We do a lot of this already in terms of what we call the empathy phrases or scripting. I think this is something that could be done regardless of whether or not I’m interacting with someone that is dealing with trauma or IPV. This should be used across the board.” (S6, manager)*

Practitioners contributed ideas on training. S1, who created training content for their company’s support agents, suggested basing materials on stories or scenarios so that agents could quickly draw connections to cases they encountered and identify potential solutions. S10, a manager, emphasized that training should be offered regularly to keep up with the evolving spyware landscape.

**Address agents' secondary trauma.** Two groups reflected on the necessity of providing mental health support to agents who interact with IPV survivors and witness the tech-enabled abuse they are experiencing. The notion that support agents themselves might be survivors provoked reflection:

*“Didn’t even consider that. It’s funny that considering the stats . . . I got a hundred [agents] on the floor, odds are some of them have been affected by this.”* (S6, manager)

S8, who maintained their company’s blog on digital rights and anti-stalkerware initiatives, noted the psychological toll in dealing with IPV cases, especially for newcomers:

*“These stories add up. I think they take a toll on us, particularly for people who aren’t aware of them. For people who maybe this is the first time they’re learning about how prevalent this problem really is, it can be a bit of a shaky, shattering moment for them.”* (S8, content writer)

**Share tech safety resources.** In line with IPV professionals’ suggestions, practitioners from all groups reported that their company was already providing customers with general tech safety advice under certain circumstances. Such advice included performing a factory reset when getting a new phone and using a password manager if the customer reports account hijacking.

Practitioners further expressed interest in providing curated content to educate customers about security and privacy. Given that all companies already had a website with basic online safety advice, practitioners viewed adding articles about IPV and tech-enabled abuse as a low-hanging fruit of critical importance. A manager stressed that tech safety alone might be insufficient for survivors and should come with external resources, similar to the IPV professionals’ suggestions:

*“This could be quite easily done ... setting up this knowledge base article, help center ... and giving the guidance of ‘These could be potential steps to take in consideration of safety planning. Get in contact with ... organizations that can support you.’”* (S10, manager)

**Make referrals.** Practitioners considered referrals to external organizations achievable. Three groups said they already did this to some extent, e.g., by directing victims of online scams to a governmental fraud investigation team. A manager described a case where they directed a customer to law enforcement:

*“We’ve gotten requests in the past where people have said, ‘Hey, I think my husband is hacking my computer. Can you find their IP address and do all this stuff for us?’ I’m like, ‘Well, we can’t do that for you. If you suspect that something’s going on, first let’s make sure that the [product] is installed and running properly to protect any type of intrusions ... If you still have concerns, then contact the local police and report.’”* (S6, manager)

Practitioners commented that expanding the scope of their current list of external referrals would improve the process without negatively impacting agents’ capacity. Nevertheless, practitioners noted that referred resources should be up-to-date and relevant, which requires maintenance efforts. Moreover, sharing geographically applicable resources could be challenging for companies that operate on a global scale.

Regarding the idea of creating an internal specialized team to handle tech-enabled IPV cases, three groups mentioned budget and capacity barriers, particularly facing financial constraints due to the COVID-19 pandemic. Two groups further suggested tracking the number of relevant cases to inform this decision, echoing the IPV professionals’ suggestions. As a manager told us:

*“I think our founder would have a genuine interest but I think we’d also need to balance that with business needs too ... We need to get a better sense of how many calls we have coming in that ... go more towards violence and partners taking retaliatory behavior.”* (S3, manager)

### 8.4.3 Implementation Challenges

Practitioners discussed challenges in implementing some of IPV professionals’ suggestions. Some practitioners questioned whether customer support, as experts on products and technical issues, should intervene in IPV cases. Others worried that frontline agents have limited capacity to help and might struggle to identify survivors who need help.

**Uncertain role of customer support.** Two groups expressed uncertainty about the role of customer support in addressing tech-enabled IPV. From their perspectives, agents should play the traditional role of customer support—focusing on the product and making customers happy. They were hesitant to let agents “take sides” in IPV situations. A manager said:

*“The agent’s role is to focus on the product. Because we don’t know what’s going on in the customer’s life ... There’s the rights of the person that’s calling us as well as the rights of the individual being accused. It’s best not to take sides and just stay neutral.”* (S6, manager)

Other practitioners expressed confidence in their products, viewing them as the ultimate solution for most customers, including survivors. A training consultant considered increasing customers’ confidence in the product as the end goal:

*“[Customers] need to get confidence in [the agent] they talk to, that here, this person knows what technology is ... whatever workaround that person is providing, if they follow that, then they don’t have to worry any*



*further about ...being [the] victim of technological abuse.” (S1, training consultant)*

While a commitment to providing customers with high-quality technical solutions is essential, the confidence in security software’s ability to fully protect survivors contradicts the caution requested by IPV professionals who view overpromising as frustrating and dangerous for survivors. Importantly, not all practitioners shared this overconfidence. For example, a researcher agreed that agents should not overpromise and drew connections to a case in which the attacker was configuring the victim’s Google accounts for location tracking:

*“In this case, technically our detection could not help, because this was actually done through the official Google apps ... We are aware of what stuff can go on, and we are careful not to overpromise ...pushing [our] product or anything.” (S11, malware researcher)*

S10, who came from the same company as S11, similarly acknowledged their product’s limitations and the importance of safety planning in removing stalkerware from the survivor’s device. They further provided an example script for agents to explain the situation to a survivor:

*“We cannot support you in the full steps but we know organizations you can [get] in contact with ... If you discuss the safety planning [with] them ... then you can come back and discuss with us how we [can] remove the app from your device.” (S10, manager)*

**Identifying potential survivors is challenging.** IPV professionals argued that customer support should not be conservative in making referrals. By contrast, support practitioners focused more on correctly identifying survivors who might need referrals and saw challenges to this end. In response to IPV professionals’ suggestion

to familiarize agents with common types of tech abuse, a manager said this would not be effective without self-disclosure from the survivor:

*“That’s a good idea but in practice would be difficult . . . I think it’s really going to be the customers coming forward and saying that this is happening. That would trigger stuff on our end to handle it differently.”* (S6, manager)

Another manager noted that most customers might not have extensive technical knowledge and they might struggle to describe issues accurately, making it challenging to diagnose the problem:

*“The victims may be aware that something is wrong on [their] phone, but cannot really describe what the issue is about . . . or maybe [they] describe it [on] a high level.”* (S10, manager)

As one solution, a practitioner proposed probing questions to confirm the customer’s “survivor” identity. However, we caution that such questions, especially those on the history of abuse, might unintentionally re-traumatize the customer and differ from IPV professionals’ suggestion to consider additional risks and attack vectors rather than to verify the IPV situation:

*“We do some verification for customer contacts . . . where we collect basic information like name, email, address . . . But I don’t know, it’s not fool-proof to see if they were actually victims of abuse. Or by giving them some open questions like, how were they victimized? Having them quote some examples that can give us a sense?”* (S2, engineering & support liaison)

**Complexities of tech-enabled IPV.** Practitioners discussed the socio-technical challenges in IPV and the resulting problems for support agents. All groups listed

the dual-use nature of many apps used by abusers [76] as a challenge. A manager described training agents to watch out for dual-use apps:

*“Sometimes [agents] have to make some additional changes to . . . our software to categorize those types of gray applications as malicious so that it can be removed. Our agents are trained on that so that’s probably one of the first things they would do.”* (S6, manager)

Another manager considered the possibility that the abuser might be monitoring the conversation, and simply removing the stalkerware might put the survivor at further risk:

*“Just to say, ‘Hey, your device is infected’ and remove the stalkerware typically means a risk for the victim . . . We don’t see [an] ideal way of communication if we identify stalkerware on a device, because the victim most likely gets observed on all channels . . . if we shot them an email to their Google account . . . the attacker can see this communication. Just removing without notification, a victim could also be at risk because the attacker assumes that the victim is aware.”* (S10, manager)

**Frontline agents have limited capacity.** On top of challenges in identifying and addressing tech-enabled IPV, two groups pointed out that support agents already work hard and have little time or capacity to take on new and complex tasks. S7, a manager, described frontline agents as *“the Cinderella of companies”* with the least pay but the expectation of doing a perfect job. In response to IPV professionals’ suggestion that agents mention possible consequences of given advice, S8 was concerned that there might be too many consequences for frontline agents to foresee, pointing to the importance of external referrals for safety planning:

*“My answer is trust the National Domestic Violence Hotline. Call them from a safe device. But that’s it. There really isn’t a one-size-fits-all*

*answer on this. [Safety planning] is something that takes more than a couple of minutes ... I could not see that happening in under an hour."*

(S8, content writer)

## 8.5 Discussion

Our findings show that support agents already encounter cases of tech-enabled IPV. There are many ways customer support could help survivors and their challenges playing this role. We now discuss the limitations of our work, reconcile perspectives between the two sets of focus groups, and distill areas computer security companies can explore to improve their customer support for IPV survivors.

**Limitations.** Our research has several limitations. Our samples were relatively small, including 17 IPV professionals and 11 customer support practitioners. While the companies in our sample are global leaders in the consumer security market, the IPV organizations are all based in US metropolitan areas. As a result, there may be recommendations for other national and international audiences that we have not yet uncovered. Additionally, we recruited support practitioners from companies already involved in the Coalition Against Stalkerware [479]. While these companies are more likely to adopt our recommendations, given their commitment to fighting stalkerware and tech-enabled abuse, they might not represent the acceptance and feasibility of adoption at companies that have not demonstrated such commitment.

**Security software is not a silver bullet.** Existing anti-virus and anti-spyware tools struggle to detect dual-use apps in intimate partner surveillance [76]. Even with improved detection algorithms [431], security software cannot fully protect IPV survivors as they face complex social and legal challenges [168]. IPV professionals unanimously agreed that security software is not a silver bullet for addressing

tech-enabled IPV, and coordination with other stakeholders in the IPV ecosystem is vital to providing survivors with holistic support. Some customer support practitioners acknowledged their products' limitations and the importance of avoiding overpromises. By contrast, others sought to give customers confidence in provided solutions or believed their software would protect most customers by default. The divergent opinions between practitioners from different companies reflect that a mentality change in dealing with IPV cases must occur at the company level—pursuing perfect technical solutions might be reasonable for general customers but could be dangerous and misleading for IPV survivors. Agents should communicate the benefits of a technical solution while acknowledging that successfully resolving a tech issue at the moment is unlikely to resolve all of a survivor's problems.

Importantly, IPV survivors face risks of escalated abuse even for routine privacy-protective measures like turning off location tracking or changing passwords [167, 168, 265]. Therefore, for any provided technical solutions, agents should be equipped to recognize the potential consequences for survivors and recommend alternative solutions that account for an abuser who might control a survivor's devices or accounts. For instance, for survivors with suspected spyware on their phone, agents should highlight that any activity on the device may be seen by the abuser and ask the survivor to consider how to proceed instead of simply removing the spyware, as noted by both IPV professionals and support practitioners. Then, agents should be prepared to refer customers to IPV professionals and other external resources for in-depth safety planning, which is a highly personalized and comprehensive plan that helps survivors remain safe both physically and digitally, and outside of the scope of what a support agent can reasonably be trained to do. By recognizing their work's boundary and facilitating the connection to other experts, customer support agents increase the chance that a survivor gets the help they need with precaution.

**Make external referrals with care and caution.** IPV professionals and support practitioners both emphasized the importance of external referrals. All companies we spoke with were already referring customers to certain external resources such as law enforcement, so the infrastructure and general procedure for doing this are in place. The most immediate next step is adding domestic violence hotlines, human trafficking hotlines, suicide helplines, and others to the referred resources. Furthermore, the provided resources need to be up-to-date and geographically relevant, as noted by practitioners. Even though some regional organizations maintain lists of state and local domestic violence hotlines (e.g., the National Domestic Violence Hotline in the US and the Women Against Violence in Europe) and can refer survivors onward, many countries lack a national hotline for domestic violence [372], indicating the need of broad referrals for survivors in these areas.

Support practitioners and IPV professionals noted different challenges regarding the specific processes of making external referrals. Support practitioners highlighted challenges around *when to refer*: not only recognizing signs that someone might need a referral but also doing enough vetting to determine that the customer was experiencing abuse. IPV professionals did not consider the latter point necessary or advisable, as it could lead to presumptive labeling or traumatizing questions. Instead, they emphasized that agents should provide referrals whenever there are signs (e.g., red flag words) indicating a need for further assistance. They were mainly concerned with *how to refer*, suggesting that agents use respectful language in offering referrals to avoid labeling and give the customer enough agency to decide whether they need or want to act on it. For high-stakes situations like IPV, it is critically important to ensure that whoever needs resources can learn about the resources, whereas recommending resources with proper and non-judgmental language does not harm customers who do not need them. By offering referrals, support agents are not “taking sides,” but rather serve as crucial connectors to social workers, attorneys, law enforcement, and

other IPV experts.

Importantly, avoiding harmful labeling does not mean agents should be vague in describing the referral resources and associated risks. Agents should clearly note the nature of the referred organizations and account for any potential repercussions from the abuser in providing the advice. For instance, when sharing the number of a helpline, agents can use the same terms the survivor uses to avoid labeling while still being explicit about the audience it serves (e.g., domestic violence survivors). Agents should further caution that the number, if called, would be in the call history and might be seen by the abuser; a safer option may be to call from a friend's phone or a public phone. Additionally, agents should not treat all abuse victims as IPV survivors by default. Targeted digital attacks also occur to NGO employees [288], politicians [206], journalists [552] and many others; the victims bear similarities to IPV survivors but have distinct vulnerabilities. Ideally, trained agents can quickly recognize these situations, use trauma-informed responses, and make referrals to related resources if needed.

**Train customer support agents.** Both IPV professionals and support practitioners unanimously agreed that training frontline agents to be better prepared for tech-enabled IPV cases is both feasible and critical for supporting survivors. Support agents are already dealing with these cases, and survivors who contact computer security companies may not be aware of existing IPV-related resources. Some survivors may not even realize they are facing tech-enabled IPV—that a weirdly-behaving device might be the first indication of intimate partner surveillance at play. Therefore, having more potential contact points, including but not limited to support agents who receive training in identifying signs of tech-enabled IPV, is essential in raising survivors' awareness and providing them with the necessary help. Equipping agents with basic awareness of tech-enabled IPV and the caution needed for a proper response is

also vital to prevent inadvertent harm, such as escalating abuse by removing spyware without further precautions or misleading survivors by overpromising.

Based on our findings, we identify the following components as potential elements of such training. We have developed respective training materials, and pilot tested them with our partner companies, who all provided positive feedback.

1. *Introduce IPV to customer support agents.* Discuss the prevalence of IPV, including technical (e.g., how technology is misused to facilitate IPV) and non-technical aspects (e.g., the survivor’s and abuser’s social entanglements and the need for holistic safety planning). Explain why agents should be committed to learning how to support survivors.
2. *Describe common tech-enabled abuse and desired responses.* Present scenarios of how abusers exploit technologies in IPV and model how agents should respond. Define and give examples of trauma-informed language, and explain its importance. Frame the problem as an opportunity to offer help rather than a situation in need of carefully vetting or evaluating the customer’s victimhood.
3. *Explain how agents could provide support.* Introduce agents to methods for assisting survivors, such as asking questions that take into account broader risks beyond the immediate tech issue, sharing tech safety resources, internal coordination, and external referrals.
4. *Identify mental health resources for agents.* Provide internal and external resources (e.g., therapeutic sessions and peer support) for agents who might be experiencing IPV or suffering secondary trauma from handling such cases.

Ultimately, training could make agents aware of unique risks and nuances in IPV, help them pick up cues that indicate customers experiencing IPV, and how to safely and respectfully respond and share external resources. As discussed by support practitioners, training should be updated and provided periodically to strengthen recall, as



frontline agents might not encounter IPV cases frequently enough to practice applying the knowledge. As highlighted by both IPV professionals and support practitioners, some training components, such as trauma-informed language, provide benefits beyond IPV survivors. For example, ransomware or identity theft victims also deal with complex tech issues, distress, and repercussions affecting their lives. They would also benefit from interacting with agents that use trauma-informed language.

**Track IPV cases to inform decision-making.** Some IPV professionals proposed having an in-house specialized team for IPV cases to reduce the pressure on frontline agents and save effort in training. However, support practitioners responded that justifying the cost of building this specialized team is difficult when their company does not know how frequently their customers would need it. The idea of anonymously tracking tech-enabled IPV cases in support agents' daily work emerged in both sets of focus groups. Doing this would create meaningful data to guide companies in making business decisions, including but not limited to creating a specialized internal team to support survivors. It would also provide a better understanding of the frequency and types of attack mechanisms, how agents handle these cases, and the extent to which agents may experience secondary trauma. Such knowledge can inform computer security companies about future challenges and opportunities to help IPV survivors and support agents.

**Partnership between security companies and IPV advocates.** Tech-enabled IPV is likely to persist, indicating the need for coordinated expert support. Both computer security companies and IPV advocacy groups are vital to supporting survivors. Our research illuminates the first steps in synthesizing the expert advice from IPV professionals and support practitioners, who each have in-depth knowledge and awareness of constraints in their professions. As tech-enabled IPV grows in prevalence and changes its forms, new countermeasures are needed to protect survivors.

An enduring partnership between IPV support organizations and computer security companies provides learning pathways for both parties. Support agents can learn about guidelines for interacting with survivors and incorporating them into protocols and training. IPV professionals can receive guidance on recognizing signs of spyware and other abuse-enabling technology in their work.

We further envision coordinated approaches to help survivors via this partnership. For example, instead of sporadic referrals to domestic violence hotlines, computer security companies and IPV professionals could work together to design a remote version of a “security clinic” [166, 191, 212] that could eventually offer digital safety planning for individual survivors. An established partnership could increase IPV professionals’ confidence in referring their clients to computer security companies that are known to be committed to knowledgeably and compassionately assisting survivors. Notably, support agents and IPV professionals should reach a consensus about their own responsibilities in such a collaboration—support agents for technical issues and basic tech safety tips; IPV professionals for comprehensive safety planning and non-technical assistance—so that survivors do not end up being referred back and forth between these parties without getting help.

## CHAPTER IX

# Nudges to Encourage Password Changes After Data Breaches Using the Protection Motivation Theory<sup>1</sup>

The previous two chapters, while focusing on different audiences and contexts, both demonstrate the value of developing the approaches to motivating adoption by adopting a user-centered process, i.e., eliciting the needs and preferences of the user population from the subject directly (Chapter VII) or from other relevant stakeholders indirectly (Chapter VIII). Rather than rely on consumers entirely to find the right approach, in this chapter, I present a study in which we developed approaches that motivate consumers to change passwords after data breaches drawing inspiration from the Protection Motivation Theory (PMT) and findings in previous chapters of this dissertation about impediments to action. Specifically, we developed nudges that seek to help consumers assess threats and coping strategies regarding password breaches, pilot tested the nudges extensively to ensure comprehension, and evaluated their effectiveness at motivating password change intention and action in a longitudinal experiment.

My previous chapters have documented the numerous harms data breaches can introduce to affected consumers. Among all types of data breaches, the ones that involve login account credentials have hit some high-profile service providers such as

---

<sup>1</sup>Yixin Zou, Khue Le, Peter Mayer, Adam J. Aviv, & Florian Schaub. 2022. Motivating Password Changes After Data Breaches Using the Protection Motivation Theory. Manuscript in preparation.

Yahoo!, LinkedIn, and Dropbox [233]. These stolen credentials often end up on the dark web and become trading assets to cybercriminals [514]. Attackers can then try the stolen credentials on different online accounts at scale through automated login requests (i.e., credential stuffing), meaning that a breach of one service provider’s password database could put other accounts at risk when they use the same or similar passwords [557]. Once accounts have been compromised, attackers may use them for spam, fraud, identity theft, or distributing malware [380], leaving users—especially those who reuse their passwords—vulnerable to these security incidents.

With nudging consumers toward changing breached passwords as the end goal, we saw the promise of applying the Protection Motivation Theory (PMT) to motivate behavioral changes after data breaches based on evidence from previous chapters. Specifically, PMT posits that individuals form protection motivation via assessing the threat (threat appraisal) and evaluating coping strategies (coping appraisal). In Chapter III, most participants did not think they would be personally affected by the Equifax data breach [602], suggesting low perceived threat vulnerability. In Chapter IV, most participants considered the breach’s impact on their lives to be “none” or “very little” [329], suggesting low perceived threat severity. Moreover, the presentation of recommended actions in breach notifications might fail to address the action’s response efficacy (since they often recommend what to do but not why doing so is necessary), and the overwhelming amount of information might hamper consumers’ self-efficacy [600].

To assess the effectiveness of nudges based on PMT at encouraging password changes after breaches, we conducted a longitudinal experiment, measuring participants’ intention ( $n=1,386$ ) and action ( $n=1,175$ ) in changing their breached passwords. We randomly assigned participants to a control condition or one of three treatment conditions (a threat appeal, a coping appeal, or both) and situated our nudges in real-world breaches known to have exposed their information using data

from Have I Been Pwned (HIBP). With this method, we managed to achieve an ecologically valid evaluation for our nudges, which has been a key challenge in testing breach notifications due to the difficulty of ethically simulating a data breach or legal implications when partnering with companies to send out experimental notifications.

We found that the threat appeal alone made participants 1.48x more likely to form password change intention; when presented together, the threat and coping appeals made participants 1.54x more likely to change their passwords eventually. However, we did not find any statistically significant difference between the three treatment conditions. Our manipulation check results suggest that the threat appeal was more powerful than the coping appeal, and our nudges affected the factors we targeted to some extent but in directions we did not expect: participants who received the threat appeal also had higher perceptions of the response efficacy, which was a component of coping assessment. Our regression models to identify additional covariates associated with password changes highlight the significant influence of user characteristics (e.g., security attitudes, password reuse, and demographics) and contextual factors (e.g., time passed since the breach occurred). In addition, our analysis of participants' open-ended responses identifies the numerous challenges they encountered when trying to change their passwords. Taken together, these findings illuminate that while PMT-based nudges can be useful, we need to rethink whether password changes should be the nudging goal when consumers take alternative actions that improve their security outcomes or when usability issues in the password ecosystem make the act of changing passwords itself impossible.

## 9.1 Related Work

We review the Protection Motivation Theory, which serves as the theoretical foundation of this work. We then review prior work of password-related behaviors (including but not limited to those after data breaches) since they are closely related to

the context in which we deployed our nudges.

### 9.1.1 Protection Motivation Theory

The Protection Motivation Theory (referred to as “PMT” onward) seeks to explain individuals’ cognitive responses in the face of fear [422, 423]: people conduct two appraisal processes, one focusing on the threat itself (threat appraisal) and the other focusing on their ability to act against the threat (coping appraisal). In threat appraisal, people consider the negative consequences of the threat if it occurs (threat severity) and how susceptible they are to the threat personally (threat vulnerability) [422]. In coping appraisal, people evaluate whether taking a recommended action will remove the threat (response efficacy) [422], their confidence in carrying out the action (self-efficacy) [574], and barriers to taking action such as financial costs, time, and effort (response costs) [160]. Higher perceptions of threat severity, threat vulnerability, response efficacy, and self-efficacy increase protection motivation, whereas higher perceived response costs decrease protection motivation [423].

**PMT in security and privacy research.** While PMT was originally proposed to study health behaviors, the theory has been applied to a wide range of domains, such as environmental concerns and political issues. Within the security and privacy domain, a large body of work has examined PMT in the context of employees’ computer security behaviors and adherence to organizational security policies [63, 102, 216, 237, 363, 463, 588]. Other studies involving average users focus on PMT’s role in data backups [51], security behaviors on home computers [24, 202], smartphone locking [17, 18], password creation [544], and adoption of tools such as secure mobile payment apps [484] and Tor browser [485]. Most of these studies sought to achieve one or more of the following goals: (1) validating that PMT can be applied to explain security behaviors in specific contexts; (2) integrating PMT with other

social and behavioral theories, and (3) applying PMT to develop and evaluate nudges that promote security behaviors.

Among the studies that evaluated PMT-based nudges, the majority of them followed the experimental design of a PMT condition (including threat and coping appeals) versus a control condition [18, 51, 246, 484, 485]. Some studies termed the PMT condition “fear appeal”—which makes it sound like the condition exclusively focuses on threats—but in essence, an ideal fear appeal should address both the threat and the individual’s ability to deal with it [51]. The results have shown that the PMT condition promoted more secure behaviors consistently across contexts. Yet the format of the fear appeal matters; e.g., in Vance et al.’s study, an interactive fear appeal treatment that changed based on users’ password input resulted in significantly stronger passwords, but a static fear appeal did not [544].

The effectiveness of using PMT in developing nudges may indicate that increasing protection motivation requires both threat and coping appeals. But to prove this argument, we need to examine threat and coping appraisals in isolation, and prior work has provided relatively limited insights on this topic. Our study took inspiration from van Bavel et al.’s study, in which the authors separated threat and coping appeals in examining their effects on online security behavior including choosing a secure connection, selecting a trusted vendor, choosing a strong password, and logging out); the authors found that the coping appeal was as effective as both appeals combined, but not so the threat appeal alone [542]. Other survey-based studies also validated constructs within the coping appraisal as significant predictors of security intention or behavior, usually through structural equation modeling [49, 63, 102, 202, 216, 237, 252, 280, 289, 363, 463].

The salience of coping appraisal does not undermine the importance of threat appraisal, however, as many of these studies also found significant effects for threat-related constructs, either threat severity/vulnerability alone [63, 216, 237, 252, 289]

or both [24, 463]. Other times, threat perceptions did not directly impact intention or behavior but moderated the effect of coping-related constructs [363] or had interaction effects with the subject's occupational background [102]. Through a literature review, Mayer et al. summarized that threat vulnerability and rewards of maladaptive response are not reliable predictors of security intention; other PMT constructs have reliable albeit weak to medium effect sizes [326]. Nonetheless, the authors cautioned that the non-reliable findings were based on a limited set of studies, and future research is needed to confirm this assessment [326].

The inconclusive findings of threat vs. coping comparisons might also result from the limitations of PMT itself. As Boss et al. summarized, PMT best explains behaviors in situations involving a highly personally relevant threat along with strong efficacy for the coping response [51]. Furthermore, like any other behavioral change theory, PMT hardly captures all nuances in people's behaviors and behavioral changes. Prior work has integrated PMT with other theories and highlighted the importance of other factors such as personal responsibility [49, 280, 455], attitudes [463], social norms/influences [216, 252, 463], psychological capital [63], prior negative experience [289], and more.

**A possible intention-behavior gap.** Behavioral intention is often used as a proxy for protection motivation and is considered a reasonable predictor of behavior according to the Theory of Planned Behavior [15]. However, prior work has revealed a possible intention-behavior gap, calling for measuring both variables when possible. For example, a meta-analysis of experimental evidence showed that a medium-to-large-sized change in intentions led to only a small-to-medium-sized change in behavior [570]. Key challenges that people may encounter as they strive to enact their intentions [453] include failing to get started, failing to keep goal pursuit on track, and failing to bring goal pursuit to a successful close [453]. Correspondingly, there



are self-regulatory mechanisms that target these challenges to help people realize their intentions, such as forming implementation intentions [181] and monitoring goal progress [208].

Among the limited number of studies that examined the intention-behavior gap in security contexts, Crossler et al. found that the costs of implementation (e.g., time and inconvenience) could be a strong deterrent to full compliance for employees to follow Bring Your Own Device policies [102]. Similarly, Jenkins et al. found that high levels of required effort negatively moderated users' intentions to follow security policies [245]. In the follow-up work of the study presented in Chapter IV [328], we found that the presence of an intention-behavior gap for actions after data breaches largely depended on the specific action: participants followed through with their intentions across all levels for actions like changing the password and signing up for credit monitoring services, but for actions that require more efforts and sound more serious (e.g., freezing credit reports and filing a complaint), intention did not predict action reliably.

In the privacy literature, Norberg et al. were the first to demonstrate the intention-behavior gap by showing that people are more likely to share personal information than they intend to during marketing exchanges [365]; the authors coined the term “privacy paradox” to describe this phenomenon. The privacy paradox remains a topic of debate since some studies found a positive correlation between intention and behavior with large effect sizes [120, 262], and others found a reversed intention-behavior gap where participants disclosed less information in the behavior condition than in the intention condition [487]. Research on the underlying mechanisms of the concern-behavior gap [41] (which might also explain the intention-behavior gap) argues that the gap could occur after an explicit risk-benefit calculation (which could be biased or unbiased), or individuals may engage in little or no risk assessment due to a lack of knowledge [4] and learned helplessness [458].

Altogether, our study contributes to the PMT literature by (1) applying PMT to encourage protection motivation in a new context (changing passwords after data breaches); (2) separating threat and coping appeals to understand their relative effectiveness; and (3) measuring both intention and behavior to get a complete picture of protection motivation. We found that compared to the control condition, threat appeal alone significantly raised password change intentions but threat and coping appeals combined significantly motivated password change behaviors.

### 9.1.2 Password Policies, Behaviors, and Nudges

**Password Policies.** Service providers use password-composition policies to prevent users from creating easily guessed passwords [451], as strong passwords (measured by length and complexity) are harder to crack than shorter and simpler passwords [185]. The U.S. National Institute of Standards and Technology (NIST) password guidelines call for a strict eight-character minimum length requirement [185]. Many service providers go above and beyond to require different character classes to be included or periodic password changes [99], although this requirement becomes ineffective as users tend to bypass it in predictable ways, such as by simply capitalizing the first letter of their password or adding a “!” to the end [538]. There are also usability costs when users struggle to keep track of updated passwords [241]. Shay et al. suggested that requiring longer passwords (e.g., 12 or 16 characters) with fewer composition requirements provides both security and usability benefits over requiring 8-character passwords with many character classes [449].

Research of password corpora has also identified characteristics of common passwords, which inform the creation of password policies about certain phrases to avoid. According to an analysis of leaked passwords from RockYou, the most popular passwords were “123456,” “password” and “iloveyou” [545]. Other prevalent semantic themes in passwords include names, locations, dates, animals, and money [338, 549].

In summary, a good password policy should require a minimal length (eight characters according to the NIST [185] or longer if the organization prefers), some but not too many classes of special characters, highlight common phrases to avoid, and remove arbitrary password expiration periods. We drew on these guidelines from prior work as we developed our coping appeal about what to choose for the new password.

**User behaviors.** Despite good-faith efforts in protecting the security of their personal information, most users struggle to comply with password policies [449] and fail to create strong passwords [196, 388, 538]. Password reuse is prevalent [105, 158, 388], which enables attackers to use the same compromised password in other protected accounts or website login. Das et al. estimated that 43-51% of users reuse the same password across multiple sites [105], whereas Pearman et al. found that password reuse is more rampant than previously believed when partial reuse is also taken into account [388]. The likelihood of password reuse increases as more passwords are created or when the password is short and simple [174]. Other research has revealed that users match password strength to the account's relative importance [538] and rarely change their passwords unless in the case of a breach or forgotten password [483].

Even with external stimuli, users do not always comply with password change advice. Thomas et al. developed and evaluated a protocol for detecting breached credentials and found that only 26% of the warnings resulted in users migrating to a new password [515]. In Bhagavatula et al.'s longitudinal study based on real-world password data, only 33% of all participants who had accounts on the breached site changed their passwords, and only 13% did so within three months of the breach announcement [46]. Similarly, in a case study of LinkedIn, Huh et al. found that less than half of participants changed their LinkedIn password upon receiving the password reset notification from the company [231]. Even for users who change the breached passwords, they rarely change the same or similar passwords on other sites [46]; their

new passwords tend to be similar to their other passwords and could still leave them vulnerable to future attacks [180]. Furthermore, users may have trust issues with third-party advice providers when there is a lack of explanation regarding how password leaks are detected and why password change is needed. In studying users' perceptions of Chrome's compromised credential notifications, Huang et al. found that some participants falsely assumed Chrome learns about users' plain-text credentials or expressed privacy concerns about Google's management of users' data [229].

Methodologically, prior research has measured password behaviors by tracking participants' log data [46, 158, 163], analyzing passwords from public or private datasets [105, 330], observing password creation in-situ [538], and relying on participants' self-reported data [180, 231, 566]. We evaluate our nudges based on participants' self-reported data, which could be prone to social desirability and recall biases, but is also grounded in high ecological validity as we present participants with real-world breaches that affect their email addresses and potentially other personal information.

**Helping users with passwords.** Password managers are tools that combine secure password storage and retrieval with random password generation to help users deploy strong, unique passwords without memorability issues [389]. However, users are often uncertain about what password managers are, how to use them, and whether they are trustworthy [19, 483]. Users may believe there is little to protect, worry about having a single point of failure, or have prior negative experiences with password managers [389, 408]. Alkaldi et al. studied how to encourage the adoption of password managers by satisfying users' self-determination needs in terms of autonomy, relatedness, and competence [20]. Other research has sought to develop nudges that help users create stronger passwords, such as visual/text indicators that provide feedback on the password's strength [70, 536], estimating how long it would take to

crack the password [544, 577], suggesting that users create passwords by concatenating a series of unrelated words together [450], and comparing the strength of the user’s password with other users of the system [135]. Peer et al. showed that password nudges could be more effective when they are also personalized to individuals’ decision-making style [392]. Our study expands prior literature by examining the applicability of PMT in encouraging better password behaviors and grounding our nudges in an environment with high ecological validity, in which participants act on their own passwords and accounts based on notifications of real-world breaches.

## 9.2 Method

Between July and August 2022, we conducted an online experiment to evaluate the effectiveness of PMT-based nudges in encouraging people to change passwords after data breaches. Our experiment was approved by our university’s Institutional Review Boards (IRB).

To increase the ecological validity of our evaluation, we adapted the survey infrastructure in Chapter IV to similarly pull breach records from Have I Been Pwned (HIBP) using its public API, thereby situating our nudges in real-world breaches known to affect individual participants. HIBP is a database maintained by security expert Troy Hunt who routinely analyzes password dumps and text storage sites on the Internet to collect information about leaked account credentials. As of August 2022, the site listed 622 pwned websites with over 11 billion pwned accounts.

### 9.2.1 Study Design

The purpose of our study is to evaluate to what extent PMT-based nudges are effective at encouraging consumers to change their passwords after being affected by a data breach. The key dependent variables are participants’ self-reported intention and behavior to change their breached passwords. Because prior work has not

provided conclusive evidence regarding the relative importance between threat and coping appraisals, our experiment follows a  $2 \times 2$  between-subjects design to examine the effect of threat and coping appeals together and in isolation. In total, our experiment has four conditions:

- Threat only: Participants were presented with information about the risks after passwords are leaked in a data breach (threat severity) and how the risks could affect them personally (threat vulnerability), cf. Figure 9.2.
- Coping only: Participants were presented with information about how to change the password (self-efficacy), how changing the password reduces the threat (response efficacy), and estimated time (response costs), cf. Figure 9.3.
- Threat and coping combined: Participants were presented with information about both threat and coping described above.
- Control: Participants were not presented with any threat or coping-related information. As a baseline, they would still see a prompt “We recommend that you change the password for your [site name] account” as well as information about a data breach that affected them. This prompt appeared in all conditions.

We carried out our experiment through three online surveys following a longitudinal design to cover the measurement of both intention and behavior, and to leave participants enough time to follow through with their intention if they wanted to. In the screening survey, we asked participants to provide an email address to be queried in the Have I Been Pwned database for the email’s associated breach records. We then determined each participant’s eligibility based on the query results and invited those eligible back for the main survey. In the main survey, we showed participants one of the four nudges (randomly assigned) and elicited their password change intentions. We then invited participants who had not changed their password before our

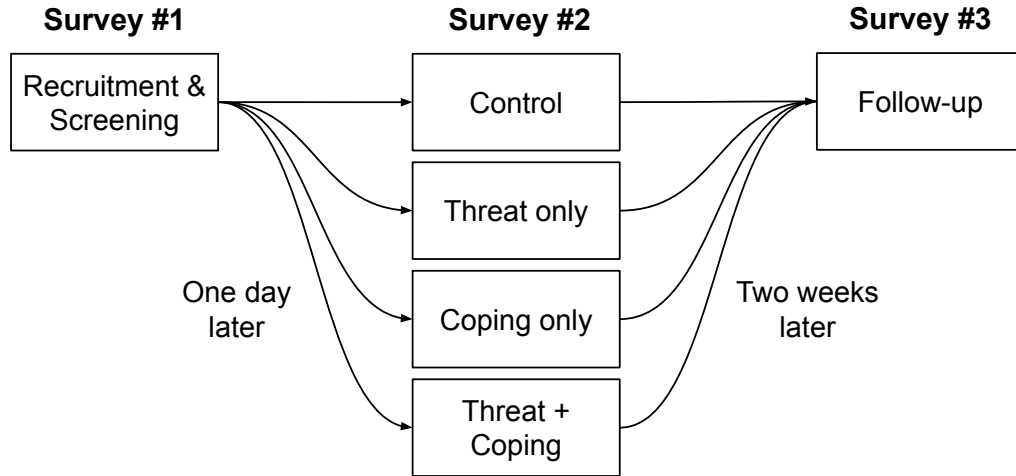


Figure 9.1: An overview of our experiment’s procedure.

study to a follow-up survey two weeks later, in which we elicited their actual password change behaviors. Figure 9.1 shows an overview of the experiment’s procedure, and complete survey materials are included in Appendix G.1.

### 9.2.2 Protocol

**Pilot testing.** To ensure the nudges we developed work as intended (e.g., the threat appeal indeed raises people’s attention to threat severity and vulnerability regarding breached passwords) and the questionnaire is understandable, the first author conducted four cognitive walkthrough sessions [286] with participants recruited via their social networks. Each session lasted about 60 minutes, in which the participant was asked to take all three surveys while verbalizing their thoughts and questions as they moved through the pages (see Appendix G.2 for the cognitive walkthrough protocol). All four participants confirmed that the text was intuitive and aligned with their understanding of the threats regarding breached passwords and the challenges regarding password changes. Based on the feedback received, we made a few changes to improve the questionnaire design, such as including social login [159] as a way of account creation.

To iterate on the questionnaire design and troubleshoot potential technical issues

**What are the risks**

- Criminals may access your account to steal your personal information, impersonate you, or make fraudulent purchases in your name.
- If you used the same password elsewhere, criminals may take over your other accounts too.
- Criminals use automated programs to test compromised passwords on hundreds of accounts in just a few seconds. You're at risk regardless of whether you are a promising target or not.
- Once your password is out there, criminals may try to take over your account anytime after a breach, no matter how long ago the breach happened.

Figure 9.2: The threat appeal (1) features negative things that could happen to affected users when a data breach involves login credentials [180, 514] (matching threat severity) and (2) addresses optimism bias and hyperbolic discounting that people may experience after data breaches [260, 602] (matching threat vulnerability).

**How to change your password** Changing your [site name] account password would prevent criminals from using the breached password to access your account. It only takes a few minutes. Just follow these easy steps:

1. **Go to [site URL] and log into your account.**  
Unsure if you have a [site name] account or can't log into it? Contact [site name] to recover the account or have your account deleted. You can usually find contact information in the privacy policy.
2. **Create a unique and strong password in account settings.**  
Longer passwords are best. Do not reuse the same password for other accounts. Check out [this guideline](#) for more do's and don'ts about passwords.
3. **You're all set!**  
If you used your old password for other accounts, make sure to change your password for those accounts too.

Figure 9.3: The coping appeal (1) builds the connection between changing the password and reducing the chance of account compromise (matching response efficacy), (2) gives a list of concrete steps to take, including URLs when applicable (matching self-efficacy), and (3) clarifies that the effort “only takes a few minutes” and is “easy” (matching response costs).



at a larger scale, we conducted another round of pilot testing with participants recruited via Prolific, a crowdsourcing platform commonly used in social science and behavioral research. We collected 107, 76, and 69 complete responses for the screening, main, and follow-up surveys, respectively. Based on pilot data, we made further changes to the questionnaire design (e.g., adding “already changed password” as an answer option for the question about password change intention). We also used the pilot data to determine participants’ compensation and effect sizes to be considered in power analysis (more in Section 9.2.3).

**Recruitment & Data Collection.** We recruited participants for the screening survey via Prolific. We made our screening survey open to prospective participants who speak English, are at least 18 years old, and are currently living in the United States. We also asked participants to take our survey on a desktop device and in the Chrome or Firefox web browsers, since we found through pilot testing that some survey questions might cause technical issues when displayed on a tablet or mobile device or when opened in other less mainstream browsers.

We began collecting data for the screening survey in July 2022 and concluded data collection for the follow-up survey in August 2022. We obtained 2,412 complete responses to the screening survey. In creating our sample, we sought a diverse representation of gender, age, and educational attainment. Specifically, we used the “balanced sample” feature on Prolific, which allows us to distribute the study evenly to male and female participants.<sup>2</sup> Furthermore, we released slots for the screening survey in small batches ( $n \approx 100$ ), monitored the demographic distributions of incoming data, and used Prolific’s pre-screeners to balance age and education as needed. Since we found that random sampling on Prolific tends to attract younger and more educated participants, we targeted the screening survey to middle-aged ( $\geq 35$  years

---

<sup>2</sup>While Prolific uses sex as one of its pre-screener, we used gender in our questionnaire, also including “non-binary” “prefer to self-describe” “prefer not to say” options.

old) and older ( $\geq 45$  years old) participants in nine out of 24 batches and people who do not have a college degree in two out of 24 batches. We only did such stratification for the screening survey since the main and follow-up survey participants came from the screening survey.

Our experiment began in the main survey and continued in the follow-up survey. Among participants who completed the screening survey, 1,824 (75.6%) were deemed eligible. We invited these participants back for the main survey one day after they completed the screening survey. Among the 1,654 (90.7%) participants who returned and completed the main survey, 1,388 (83.9%) selected “yes” or “no” for the password change intention question, and the remaining 266 (16.1%) selected “already changed.” We excluded participants who selected “already changed” for the follow-up survey and in our data analysis since the nudge does not apply to them, and there is no need for them to change their password again. Among participants who received the follow-up survey invitations, 1,176 (84.7%) returned and completed the follow-up survey.

In terms of compensation, our pilot data suggested that participants on average took three, nine, and four minutes to complete the screening, main, and follow-up surveys respectively. Aiming to compensate participants at least \$15/hour, we set the compensation accordingly: \$0.80 for the screening survey, \$2.40 for the main survey, and \$1.00 for the follow-up survey. The actual median completion times were 1.84 minutes (screening), 6.64 minutes (main), and 3.71 minutes (follow-up), corresponding to a rate of \$26.08/hour, \$21.69/hour, and \$18.76/hour<sup>3</sup> respectively.

**Screening Survey.** After informed consent, we asked participants to provide an email address, which we used to query the HIBP database and obtain information about data breaches in which the participant’s email address had been exposed. We clearly noted how we do the querying in a privacy-preserving manner—we kept par-

---

<sup>3</sup>The rate takes into account bonus payment for participants who received a \$1.00 bonus payment for uploading the password reset confirmation email ( $n=77$ ).

participants' email addresses in ephemeral storage and deleted them after the query had been sent; we would never have direct exposure to participants' email addresses, nor did we store them anywhere or include them in our analysis. Participants who were uncomfortable with providing their email addresses could opt out.

For participants who provided an email address, we did the query and informed them of their eligibility for our experiment based on the query results. The participant is eligible if they provide an email address for their own or shared email account, rather than an email address that belongs to someone else entirely or is made up for the study. Furthermore, the participant needs to have at least one valid breach for their provided email address. We considered that a breach is valid (1) if passwords were among the types of breached information, which makes changing passwords a potentially reasonable option, and (2) changing the password is indeed a viable option on the breached site. Specifically, the site should still be functioning, the site should have an account creation feature, and the site should conduct business with average consumers rather than other businesses. To determine criterion (2), we manually inspected each breached site recorded by HIBP in June 2022, right before launching the experiment. Our manual filtering led us to exclude 230 out of the 598 breaches (38.5%) in HIBP's database at the time of our study. For participants who were ineligible for our experiment, we redirected them to the final page showing the breach records associated with the email they provided to ensure they were aware of the situation.

To better characterize our sample, we further asked eligible participants to indicate their usage of the email account using the same questions in [329], including how often they checked the email, what they used the email for (e.g., professional/personal correspondence or account creation), and how long it has been used. We also collected eligible participants' demographic and occupational background information (CS/IT and law) at the end of the screening survey.



### Overview

In April 2019, the PDF management service [Lumin PDF](#) suffered a [data breach](#). The breach wasn't publicly disclosed until September when 15.5M records of user data appeared for download on a popular hacking forum. The data had been left publicly exposed in a MongoDB instance after which Lumin PDF was allegedly been "contacted multiple times, but ignored all the queries". The exposed data included names, email addresses, genders, spoken language and either a bcrypt password hash or Google auth token. The data was provided to HIBP by a source who requested it be attributed to "JimScott.Sec@protonmail.com".

### What data was compromised

- Passwords
- + **Additional Information:** Auth tokens, Email addresses, Genders, Names, Spoken languages, Usernames

Figure 9.4: Example breach information shown to participants.

**Main Survey.** Participants who passed screening received invitations to the main survey one day after they completed the screening survey. The main survey started by showing a valid breach involving passwords randomly chosen from their breach records. The breach information included a short description, the site's logo and name, and types of compromised data as provided by HIBP, highlighting passwords and listing other data types under "additional information" (Figure 9.4). We then asked participants about their awareness in three dimensions: if they had heard of the site, if they had known they were affected by the breach, and if they had an account with the site. For those with an account, we further asked them to specify their account usage, including the account's age, frequency of use, and perceived importance.

Next, we showed participants one of the four nudges randomly assigned. All nudges included the password change prompt, but those in the treatment conditions

received additional information about threat, coping, or both. In line with prior work’s recommendations [100, 139, 229], we followed a layered approach in displaying the nudging text by displaying it in multiple interactive, expandable boxes with a short delay between boxes, hoping that by doing this, we could avoid overwhelming participants while nudging them to pay close attention and read through the text. After showing the nudge, we asked participants to indicate their password change intention (“yes” “no” “already changed”) and explained their choice open-endedly. We also asked participants whether they had used the breached password for their other online accounts.

Following the password reuse question, we asked participants to rate their perception of the five PMT constructs (threat severity, threat vulnerability, response efficacy, self-efficacy, and response costs) using scales adapted from prior work [49, 484]. We used these responses as manipulation checks to gauge how effective our nudges were at addressing our participants’ threat and coping appraisals. We also included two attention check questions in this part, asking participants to select a fixed answer option (“very unlikely” for threat vulnerability and “extremely serious” for threat severity). The main survey ended with questions about the participant’s security attitudes [147] and prior negative experience [603] (namely with account compromises, data breaches, and identity theft) as additional possible covariates of password change intention and behavior.

**Follow-up Survey.** Participants eligible for the follow-up survey received an invitation two weeks after completing the main survey. We decided the follow-up time to be two weeks later after trying various options in pilot testing, in which all participants who ended up changing their password did it within two weeks. We reasoned that this timeframe gave participants enough time to change their password if they wanted to do it on certain less busy days (e.g., weekends) but was not too long to

cause recall issues.

The follow-up survey began by reminding participants of the previous surveys they took and the breach featured. Next, we asked participants to describe what they had done in response to learning about the breach in the main survey open-endedly, followed by an attention check asking participants to select the name of the featured breach from a correct option and three decoy options. We then asked participants to specify whether they had changed their password since taking the main survey, why, and what they did to passwords for other accounts. Those who indicated they had changed their password since taking the main survey were asked to provide more details, such as when they changed the password and what mechanisms they used to remember the new password. They could optionally upload a screenshot of the password reset confirmation email. We adopt this approach from [231] as a measure to validate participants' self-reported behavior and set a \$1.00 bonus payment for those who uploaded a valid screenshot.

### 9.2.3 Data Analysis

We pre-registered our study protocol, hypotheses, and data analysis plan on Open Science Framework prior to data collection.

**Hypotheses.** Our key independent variable is the condition (control, threat only, coping only, threat and coping), and our key dependent variables are participants' intention and action to change their breached passwords. Because the four conditions had different amounts of threat- and coping-related information, and because the control condition does not have any threat or coping information, we made the following hypotheses for password change intention:

**H1:** There will be different levels of password change intention across the four experimental conditions.

- **H1a:** Participants in the control condition will exhibit lower password change intentions than those in the threat only condition. (control < threat only for intention)
- **H1b:** Participants in the control condition will exhibit lower password change intentions than those in the coping only condition. (control < coping only for intention)
- **H1c:** Participants in the control condition will exhibit lower password change intentions than those in the threat and coping combined condition. (control < threat + coping for intention)

Assuming that the nudge’s effect on intention in the main survey carries over into the follow-up survey, we made the following hypotheses for password change behavior:

**H2:** There will be different levels of password change behavior across the four experimental conditions.

- **H2a:** Fewer participants in the control condition will end up changing their breached passwords than those in the threat only condition. (control < threat only for behavior)
- **H2b:** Fewer participants in the control condition will end up changing their breached passwords than those in the coping only condition. (control < coping only for behavior)
- **H2c:** Fewer participants in the control condition will end up changing their breached passwords than those in the threat and coping combined condition. (control < threat + coping for behavior)

We did not include directional hypotheses regarding differences among the threat only, coping only, and threat and coping combined conditions because prior work has not provided conclusive evidence on which of the threat vs. coping nudges is more effective or whether the two nudges combined would be more effective than each presented in isolation [326, 542]. However, we were still interested in understanding

their differences, and we ran additional pairwise comparisons as exploratory analyses to obtain insights.

**Data Cleaning.** We retained all completed responses for the screening survey. We used the attention checks in the main and follow-up surveys to flag responses that require further examination but did not exclude these responses immediately. Among the 1,388 participants who completed the main survey (already excluding those who already changed passwords), 32 failed one of the attention checks, and two failed both attention checks. Among the 1,116 participants who completed the follow-up survey, one failed the attention check. We excluded responses from the two participants who failed both attention checks in the main survey and the one who failed the attention check in the follow-up survey. We retained the remaining responses since the data for other parts of the survey was detailed and insightful. We also retained responses from the 211 participants who completed the main survey but did not complete the follow-up survey since they still contributed insights for half of our key hypotheses.

After data cleaning, the final sample size is 1,386 for the main survey (control: 349/25.2%; threat only: 339/24.5%; coping only: 348/25.1%; threat and coping combined: 350/25.2%) and 1,175 for the follow-up survey (control: 291/24.8%; threat only: 304/25.9%; coping only: 297/25.2%; threat and coping combined: 283/24.1%).

**Statistical Analyses.** To confirm H1 and H2, we conducted two omnibus  $\chi^2$  tests to detect significant differences between the four conditions, with intention (yes/no, excluding “already”) and action (yes/no) as the dependent variable, respectively. To confirm H1a-H1c and H2a-H2c, we conducted pairwise  $\chi^2$  comparisons to detect the differences between the control and any treatment conditions.

We also planned and conducted a series of exploratory analyses. To gather insights into the relative performance of threat vs. coping appeals, we conducted pairwise comparisons between the three treatment conditions. To understand to what extent



the threat and coping appeals address participants' threat and coping appraisals respectively, we conducted Kruskal-Wallis tests to investigate the relationship between the condition (IV) and each of the five PMT constructs (DV) as manipulation checks. We were also interested in knowing whether the effect of the nudge would remain robust after controlling for other covariates, such as individuals' account usage and demographics. To this end, we built logistic regression models with the condition and other covariates as IVs and intention/action as the DV. To avoid model fit problems caused by too few observations in a category, we binned the demographic data in Table 9.1 into fewer categories for the regression analyses: gender (binary, men or women), age (three levels: 18-34, 35-54, 55+), educational attainment (three levels: high school or less, some college, Bachelor's degree or above), and annual household income (three levels: <\$50k, \$50-100k, >\$100k). We report odds ratios, confidence intervals, and p-values for regression results: an odds ratio below one would indicate a decrease of likelihood in intention/action for changing the password, and an odds ratio above one would indicate an increase.

**Qualitative Analysis.** We analyzed participants' free text responses using thematic coding [434] to gather deeper insights into their reasoning behind password change intention (Q26; main survey), what they did in general after learning about the breach (Q38; follow-up survey), and reasoning behind password change action (Q41). The first author developed a draft codebook, iteratively improved it with input from other co-authors, re-applied new codes as they came up to earlier batches, and coded all responses eventually. Following established guidelines [336], we did not calculate inter-rater reliability since the coding was done by a single researcher who is an expert in the field, and the data was relatively straightforward.

**Power Calculation.** We conducted a priori power analysis to determine the sample size required for the study based on effect sizes observed in the pilot data. Our

pilot data suggested a small-to-medium effect ( $w=0.15$ ) for the omnibus  $\chi^2$  test on intention and a small effect ( $w=0.10$ ) for pairwise comparisons on intention. We based our power analysis on pairwise comparisons since we wanted to ensure we would have enough data to detect between-condition differences. For 80% power at  $\alpha=0.05$  and  $df=1$ , G\*Power suggested we need 1,570 participants in total for the main survey. We did not achieve this goal due to budgetary constraints and an underestimation of participants whose data got excluded in the analysis because they had changed the password already (16% of all completed main survey responses, whereas our original estimate was 5%). We conducted a post hoc power analysis and confirmed that the sample sizes we had still enabled us to detect a small effect for the omnibus  $\chi^2$  test ( $w=0.09$  for intention;  $w=0.10$  for action) and a small-to-medium effect for the pairwise comparisons ( $w=0.11$  for intention;  $w=0.12$  for action).

## 9.3 Findings

Our analysis indicates that threat appeal alone performed significantly better than the control condition in raising password change intention. The threat and coping appeals combined performed significantly better than the control condition in encouraging actual password change behavior. While being small, the significant difference for both comparisons remains robust after controlling for other covariates. Our qualitative results further illuminate the hurdles in participants' attempts to change their passwords and the alternative actions they took when they did not change their passwords. We describe our participants' demographics and the breaches in our sample before diving into the result details.

### 9.3.1 Sample

**Participant profile.** Table 9.1 summarizes the demographics of our main survey participants compared to the US census bureau's data. Our sample has a quite bal-

<b>Metric</b>	<b>Sample</b>	<b>Census</b>
Women, Men, Non-binary	50.9%, 46.5%, 1.9%	51%, 49%, n/a
18-24, 25-34 years	16.0%, 22.9%	7%, 14%
35-44, 45-54 years	21.5%, 19.5%	13%, 13%
55-64, 65 years or older	13.3%, 5.4%	13%, 15%
High school or less, Some college	11.6%, 23.7%	37%, 15%
Associate’s degree (aca./voc.)	11.8%	11%
Bachelor’s degree	36.9%	24%
Advanced degree (Master’s/professional/doctoral)	15.7%	14%
<\$25k, \$25k-\$50k	16.0%, 22.3%	19%, 20%
\$50k-\$75k, \$75k-\$100k	20.1%, 14.9%	16%, 12%
\$100k-\$150k, >\$150k	13.5%, 10.0%	14%, 19%
Asian, Black	7.8%, 6.8%	6%, 14%
White, Two or more races	76.9%, 5.0%	76%, 3%
Hispanic/Latino	9.8%	19%
Other (e.g., American Indian, Pacific Islander)	0.9%	2%

Table 9.1: Gender, age, education, income, race/ethnicity compositions among participants of the main survey ( $n=1,386$ ). Census statistics from [531, 532, 533, 534] as of 2022. Some percentages do not add up to 100% due to non-reporting.

anced distribution regarding gender, income, and race/ethnicity, but it is slightly more educated and younger than the US population. For participants’ occupational background, 314 (22.7%) reported having studied or worked in computer science/information technology, and 49 (3.5%) reported having studied or practiced law or other legal services.

We further asked participants how they used the email account corresponding to the email addresses they provided for our study. Most participants used the email for an extended period (mean: 12.81 years, median: 12). Most participants checked the email daily (85.4%); the rest checked the email weekly (11.7%), monthly or less frequently (2.9%). In addition, participants could choose multiple options for what they used the email for. Most participants selected using the email account for personal correspondence (86.4%), followed by signing up for medium-value accounts (75.0%), signing up for sensitive accounts (57.4%), signing up for low-value accounts (53.1%), and professional correspondence (40.8%). These results indicate that participants were checking breach records for the emails they used often and for important pur-

Name	Freq.	Year	Acc. No.	Breached Data
Zynga	132 (9.5%)	2019	172M	Email addresses, Passwords, Phone numbers, Usernames
MyFitnessPal	97 (7.0%)	2018	144M	Email addresses, IP addresses, Passwords, Usernames
Chegg	55 (4.0%)	2018	39M	Email addresses, Names, Passwords, Usernames
MySpace	53 (3.8%)	2008	359M	Email addresses, Passwords, Usernames
Canva	51 (3.7%)	2019	137M	Email addresses, Geographic locations, Names, Passwords, Usernames
Wattpad	44 (3.2%)	2020	268M	Bios, Dates of birth, Email addresses, Genders, Geographic locations, IP addresses, Names, Passwords, Social media profiles, User website URLs, Usernames
LinkedIn	39 (2.8%)	2012	164M	Email addresses, Passwords
ClearVoice Surveys	38 (2.7%)	2015	15M	Dates of birth, Email addresses, Genders, IP addresses, Names, Passwords, Phone numbers, Physical addresses
Evite	38 (2.7%)	2013	101M	Dates of birth, Email addresses, Genders, Names, Passwords, Phone numbers, Physical addresses
Mathway	36 (2.6%)	2020	25M	Device information, Email addresses, Names, Passwords, Social media profiles
ParkMobile	36 (2.6%)	2021	20M	Email addresses, Licence plates, Names, Passwords, Phone numbers

Table 9.2: Top breaches in our sample. Freq. indicates how many times the breach was featured in the main survey.

Conditions	% w/ intention	OR	95% CI	p-value
Threat only vs. Control	67.3% vs. 58.2%	1.48	[1.07, 2.04]	.02
Coping only vs. Control	62.9% vs. 58.2%	1.22	[0.89, 1.68]	.23
Combined vs. Control	62.3% vs. 58.2%	1.19	[0.88, 1.63]	.30

Table 9.3: Pairwise comparisons between control vs. treatment for the percent of participants who reported intending to change the password in the main survey.

poses, which adds to our findings’ ecological validity.

**Overview of breaches.** For each participant, we displayed the nudge in the context of a breach that involves passwords randomly selected from the participant’s breach records. Our sample consists of 127 unique breaches. Table 9.2 lists the top 11 breaches featured in the main survey and more details of each breach to give a snapshot. We also calculated breach age, defined as the time between the breach’s occurrence date and the survey’s completion date. The average age of breaches in our sample was 5.16 years (median: 4.2, sd: 3.13).

Conditions	% w/ action	OR	95% CI	p-value
Threat only vs. Control	28.0% vs. 22.7%	1.32	[0.90, 1.95]	.14
Coping only vs. Control	27.0% vs. 22.7%	1.26	[0.85, 1.86]	.23
Combined vs. Control	31.1% vs. 22.7%	1.54	[1.04, 2.27]	.02

Table 9.4: Pairwise comparisons between control vs. treatment for the percent of participants who reported having changed the password in the follow-up survey.

### 9.3.2 Results of Confirmatory Analyses

Overall, a majority of participants (868, 62.6%) in the main survey stated they intended to change the breached password, but only a minority (319, 27.1%) ended up changing the password in the follow-up survey.

Looking at the descriptive statistics, we observed the following trends: threat appeal alone > coping appeal  $\geq$  threat and coping appeals combined > control for motivating password intention; threat and coping appeals combined > threat appeal alone  $\geq$  coping appeal alone > control for motivating actual password changes. However, the differences between groups were minor. For the two omnibus  $\chi^2$  tests of independence we conducted on intention and action respectively, both tests revealed non-significant results ( $\chi^2(3)=6.10$ ,  $p=.11$  for intention;  $\chi^2(3)=5.27$ ,  $p=.15$  for action). As such, we rejected both H1 and H2: there were no significant differences between the four conditions regarding their impact on participants' password change intention and action.

While the overall difference was insignificant, the pairwise  $\chi^2$  tests of independence we ran between the control vs. treatment conditions revealed significant differences for specific pairs. Tables 9.3 and 9.4 show the results; we report odds ratios as the  $\chi^2$  test's effect size and the associated 95% confidence intervals, with the control condition as the baseline. Participants who saw the threat appeal alone were 1.48x more likely to intend to change their breached password than the control condition ( $p=.02$ ), confirming H1a. Participants who saw the threat and coping appeals com-

bined were 1.54x more likely to change their breached password than the control condition ( $p=.02$ ). The odds ratio for both comparisons only corresponds to a small effect size [78]. We did not find a statistically significant difference between coping only vs. control for intention ( $p=.23$ ), between threat and coping combined vs. control for intention ( $p=.30$ ), between threat only vs. control for action ( $p=.14$ ), or between coping only vs. control for action ( $p=.23$ ). H1b, H1c, H2a, H2b are rejected. These results show that compared to the control condition, threat appeal alone could effectively motivate password change intention. Still, participants were much more likely to act on the intention only when both threat and coping appeals were present.

### 9.3.3 Results of Exploratory Analyses

In our exploratory analyses, we compared differences in intention and action between the three treatment conditions, conducted manipulation checks on the threat and coping appeals, and provided more details of the covariates we measured. Our regression results indicate that participants' password reuse behavior, security attitudes, demographics, and the breach's age help explain additional variances in password change intention or action in addition to the type of nudge they received.

**Pairwise comparisons between treatment conditions.** To compare the threat and coping appeals' effects as well as understand whether threat and coping appeals combined would be more effective, we ran pairwise  $\chi^2$  tests between the threat only, coping only, and threat and coping combined conditions, with intention and action as the dependent variable respectively. For intention, we did not observe any significant differences between threat only vs. threat and coping combined ( $p=.20$ ), coping only vs. threat and coping combined ( $p=.92$ ), or threat only vs. coping only ( $p=.27$ ). The same pattern of non-significant pairwise difference also applies to action:  $p=.46$  for threat only vs. threat and coping combined,  $p=.31$  for coping only vs. threat and

coping combined, and  $p=.85$  for threat only vs. coping only. These results confirm the pattern in Tables 9.3 and 9.4: the differences in motivating password changes between the three treatment conditions were minor (1% to 4% based on descriptive statistics) and were not significant even with a large enough sample to detect a small-to-medium effect.

**Manipulation checks.** Table 9.5 shows the descriptive statistics of participants' ratings of the five PMT constructs: threat severity, threat vulnerability, response efficacy, self-efficacy, and response costs. To understand to what extent our nudges produced the intended effect on participants' threat and coping perceptions, we conducted Kruskal-Wallis tests to detect whether participants' ratings of each construct differed between conditions, followed by post hoc Dunn tests to detect significant differences between any pairs. We applied Holm-Bonferroni correction to the post hoc Dunn tests to control for Type I error due to the exploratory nature of the analysis.

Results of the Kruskal-Wallis tests showed that our nudges had different effects on participants' perceived threat vulnerability ( $H(3)=14.11$ ,  $p=.002$ ) and perceived response efficacy ( $H(3)=21.52$ ,  $p<.001$ ). The post hoc Dunn tests further revealed that participants who received the threat appeal had significantly higher ratings of perceived threat vulnerability, evidenced by the significant differences between threat only vs. control ( $p=.01$ ) and threat and coping combined vs. control ( $p=.004$ ). Interestingly, the threat appeal also seemed to increase participants' coping perceptions. Participants who received the threat appeal had significantly higher ratings of response efficacy, evidenced by the significant differences between threat only vs. control ( $p=.009$ ), threat and coping combined vs. control ( $p<.001$ ), and threat and coping combined vs. coping only ( $p=.009$ ). There was no significant difference between conditions for participants' ratings of threat severity ( $H(3)=0.97$ ,  $p=.81$ ), self-efficacy ( $H(3)=1.30$ ,  $p=.73$ ), or response costs ( $H(3)=4.06$ ,  $p=.26$ ); the post hoc Dunn tests

<b>Variable</b>	<b>Condition</b>	<b>Mean</b>	<b>Median</b>	<b>SD</b>
Perceived threat severity	Control	4.28	5.00	0.90
	Threat only	4.37	5.00	0.79
	Coping only	4.31	4.50	0.88
	Combined	4.32	4.00	0.83
Perceived threat vulnerability	Control	2.73	3.00	0.06
	Threat only	2.99	3.00	0.06
	Coping only	2.89	3.00	0.06
	Combined	3.01	3.00	0.06
Perceived self-efficacy	Control	3.99	4.00	1.14
	Threat only	3.96	4.00	1.15
	Coping only	4.04	4.00	1.13
	Combined	4.00	4.00	1.07
Perceived response efficacy	Control	3.20	3.00	1.13
	Threat only	3.44	4.00	1.15
	Coping only	3.28	4.00	1.17
	Combined	3.55	4.00	1.13
Perceived response costs	Control	1.51	1.00	0.74
	Threat only	1.52	1.00	0.79
	Coping only	1.51	1.00	0.79
	Combined	1.60	1.50	0.80

Table 9.5: Mean, median, and standard deviation of participants' rating of the PMT constructs, divided by conditions. We adapted scales from prior work [49, 484] in measuring these variables: threat severity (Q28), threat vulnerability (Q29), response costs (Q32) were measured using 5-point Likert scales, taking the median. Response efficacy (Q30) and self-efficacy (Q31) were measured using a single 5-point Likert-type item.



did not reveal any significant pairwise differences for these three variables.

These results suggest that our nudges' manipulations of PMT constructs were somewhat successful. The threat appeal primarily influenced participants' threat perception by influencing perceived threat vulnerability. In addition, the threat appeal's presence strengthened participants' perception of response efficacy, even though response efficacy is part of coping appraisal according to PMT, whereas the presence of coping appeal alone did not help much. This could mean that our coping appeal did not perform as intended. Perhaps a more plausible explanation is that participants might not need the coping appeal for the action featured (changing passwords after data breaches). As shown in Table 9.5, participants across conditions had high ratings of perceived self-efficacy and low ratings of response costs. Participants might have already felt confident about their ability to change the password without help from others, and they did not see many costs associated with this action. As such, there was little room for the coping appeal to have further effects.

**Awareness, security posture, and account usage.** In the main survey, we measured participants' prior awareness of the company/breach, security posture, and account usage as covariates, i.e., additional factors we wanted to understand and control for in examining our nudges' effect on participants' password change intention and action.

For prior awareness (Q16, Q17), the majority (75.8%) of participants had heard of the breached site before our study, 17.1% had not, and 7.1% were unsure. By contrast, the majority (82.3%) of participants were unaware that they were affected by the breach we showed them before our study, 7.0% had known they were affected, and 10.8% were unsure.

We examined participants' security posture by measuring whether they used the breached password elsewhere (Q27), security attitudes using SA-6 [147] (Q33), and

prior experience with account compromise (Q34), data breach (Q35), and identity theft (Q36). Excluding 11.0% of participants who claimed they did not have an account with the breached site, 12.3% knew for sure they used the breached password elsewhere, 32.5% did not reuse the breached password, and 44.2% were unsure. Most participants' SA-6 ratings were between the medium to high end (mean: 3.44; median: 3.75; sd: 0.95), indicating that most participants were fairly motivated to learn about security and follow expert-recommended advice. More than half of the participants knew that their information was exposed in other data breaches before our study (57.2% yes; 36.1% no; 6.6% unsure). Fewer participants had experienced account compromises (29.1% yes; 54.6% no; 16.3% unsure) or identity theft (12.0% yes; 81.4% no; 6.6% unsure).

We further asked participants whether they had an account with the breached site (Q18), hypothesizing that password change would only be actionable, and thus more likely, for those with an account.<sup>4</sup> More than half of participants indicated they had created an account with the breached site, either by providing an email address and a password (48.5%) or by using the social login feature such as “sign in with Google” or “sign in with Facebook” (6.9%). The rest did not think they had an account (16.6%) or were unsure about the account's existence (28.1%). For participants who reported having an account with the breached site, we followed up with questions about the account's age (Q20), frequency of use (Q21), and perceived importance (Q24). Most accounts had existed for an extended period (mean: 5.98 years; median: 5; sd: 4.36). Most participants (87.2%) logged into the account only yearly or even less frequently; the rest checked the account monthly (8.0%), weekly (3.5%), or daily (1.3%). When asked about the account's perceived importance, most participants indicated that the

---

<sup>4</sup>In Q22 and Q23 we also asked participants to select what types of information the account might have about themselves. However, only in the middle of data collection did we find out that Q22 did not have a “none of the above” option, and participants could not skip this question. As such, we refrain from reporting the results of this question since the responses were likely skewed due to the survey question's design.

account was “very unimportant” (51.1%) or “unimportant” (23.1%) to them. Average ratings of the perceived importance level were low (mean: 2.04; median: 1; sd: 1.44 for a 7-point scale; the higher, the more important). Altogether, these results shed light on the nature of breaches curated by HIBP and the accounts on these breached sites among our sample: most participants created the account a while ago, only used it very infrequently, and did not attach much value to the account.

**Covariates associated with password change intention and action.** Having found that our nudges only had a minor effect on participants’ intention and action of changing the password, we ran logistic regressions to identify whether the effect would remain after controlling for the covariates we measured, as well as whether any of these covariates were associated with password change intention or action.

First, we built a logistic regression model on factors associated with password change intention measured in the main survey. This model (results in Table 9.6) contains the following 14 variables: condition, prior knowledge of the breached site, prior knowledge of the breach, account existence, reuse of the breached password, security attitudes (SA-6), prior negative experience (with account compromise, data breach, and identity theft), demographics (age, gender, education, income), and the breach’s age. We found that the significant differences in intention between the threat and control conditions remained after controlling for other covariates: participants who saw the threat appeal only were 1.84x more likely to have password change intention compared to those in the control condition.

For other covariates, the model suggested password reuse and security attitudes, as part of one’s security posture, as influential predictors of password change intention. Participants who reused the breached password for other online accounts, were unsure about the password reuse situation, or held more proactive security attitudes were significantly more likely to form intentions. The model further revealed demographic

	<b>B (SE)</b>	<b>OR</b>	<b>95% CI</b>	<b>p-value</b>
(Intercept)	-2.88(0.59)	0.06	[0.02, 0.18]	< .001
Condition: coping (vs. control)	0.28(0.24)	1.32	[0.83, 2.11]	.24
Condition: threat (vs. control)	0.61(0.24)	1.84	[1.15, 2.95]	.01
Condition: combined (vs. control)	0.28(0.29)	1.32	[0.82, 2.17]	.25
Account exist: yes (vs. no)	-0.15(0.32)	0.86	[0.45, 1.60]	.63
Account exist: yes (vs. no)	-0.28(0.30)	0.75	[0.42, 1.37]	.35
Aware account: yes (vs. no)	0.50(0.35)	1.64	[0.82, 3.29]	.16
Aware account: unsure (vs. no)	0.15(0.34)	1.17	[0.59, 2.28]	.65
Password reuse: yes (vs. no)	1.10(0.31)	3.01	[1.66, 5.69]	< .001
Password reuse: unsure (vs. no)	0.60(0.19)	1.82	[1.25, 2.67]	.002
Security attitudes (5-point scale)	0.65(0.10)	1.92	[1.58, 2.34]	< .001
Acc. Compromise: yes (vs. no)	0.15(0.20)	1.16	[0.79, 1.71]	.44
Prior breach: yes (vs. no)	-0.24(0.19)	0.78	[0.54, 1.13]	.20
Identity theft: yes (vs. no)	0.12(0.29)	1.13	[0.64, 2.04]	.68
Age: 35-54 (vs. 18-34)	0.29(0.20)	1.33	[0.90, 1.96]	.24
Age: 55+ (vs. 18-34)	0.62(0.26)	1.87	[1.12, 3.15]	.15
Gender: women (vs. men)	0.68(0.18)	1.97	[1.38, 2.82]	< .001
Edu.: ≥Bach. (vs. ≤high school)	0.28(0.30)	1.32	[0.72, 2.30]	.36
Edu.: Some College (vs. ≤high school)	0.47(0.30)	1.60	[0.88, 2.90]	.12
Income: 100+K (vs. <50K)	0.52(0.25)	1.69	[1.05, 2.76]	.03
Income: 50-100K (vs. <50K)	0.03(0.20)	1.03	[0.69, 1.52]	.89
Breach age (years)	-0.09(0.03)	0.92	[0.87, 0.97]	.001

Table 9.6: Logistic regression for predicting password change intention in the main survey.  $n=725$  after excluding incomplete or not applicable responses. Cox and Snell  $R^2=0.14$ . Model  $\chi^2(21)=111.05$ ,  $p<.001$

differences across gender and income. Participants who identified as women (vs. men) and participants with more annual household income (>100K vs. <50K) were more likely to form intentions. How much time has passed since the breach also influenced password change intention: participants were less likely to form intentions for older breaches.

Next, we built a logistic regression model on factors associated with password change action measured in the follow-up survey. We used the same covariates in Table 9.6 but included intention as an additional covariate to examine whether there was an intention-behavior gap in participants' act of changing the breached password. Results in Table 9.7 confirmed that the significant difference between the threat and coping combined and control conditions remained after controlling for other covariates: participants who saw both the threat and coping appeals were 2.37x more likely to end up changing their breached password compared to those in the control condition.

Compared to results for the intention model, breach age and security attitudes still served as significant predictors for action. Password reuse, however, no longer predicted action reliably, and the significant gender and education differences for intention also disappeared. Instead, the action model surfaced age differences: middle-aged participants (35-54 years old) were less likely to change the password than younger participants (18-34 years old). The addition of intention led us to find that the intention-behavior gap was not particularly pronounced in our study. Participants' password change behavior largely aligned with their intention: those who stated intention in the main survey were 11.24x more likely to report having changed their password in the follow-up survey than those without intention.

**More details of password changes.** For all participants in the follow-up survey ( $n=1,175$ ), we asked whether they changed the password for other online accounts

	<b>B (SE)</b>	<b>OR</b>	<b>95% CI</b>	<b>p-value</b>
(Intercept)	-3.46(0.77)	0.03	[0.01, 0.14]	< .001
Condition: coping (vs. control)	0.42(0.30)	1.52	[0.85, 2.75]	.16
Condition: threat (vs. control)	0.50(0.28)	1.66	[0.96, 2.89]	.07
Condition: combined (vs. control)	0.86(0.31)	2.37	[1.30, 4.36]	.005
Aware site: yes (vs. no)	0.29(0.39)	1.34	[0.63, 2.91]	.46
Aware breach: yes (vs. no)	0.47(0.35)	1.60	[0.79, 3.20]	.19
Account exist: yes (vs. no)	0.16(0.43)	1.17	[0.50, 2.78]	.72
Account exist: unsure (vs. no)	0.03(0.42)	1.03	[0.45, 2.39]	.94
Password reuse: yes (vs. no)	-0.37(0.32)	0.69	[0.37, 1.28]	.24
Password reuse: unsure (vs. no)	-0.22(0.23)	0.81	[0.51, 1.26]	.35
Security attitudes (5-point scale)	0.27(0.12)	1.31	[1.03, 1.67]	.03
Acc. Compromise: yes (vs. no)	-0.22(0.23)	0.80	[0.51, 1.26]	.34
Prior breach: yes (vs. no)	0.02(0.22)	1.02	[0.66, 1.56]	.93
Identity theft: yes (vs. no)	0.34(0.30)	1.40	[0.77, 2.54]	.27
Age: 35-54 (vs. 18-34)	-0.62(0.24)	0.54	[0.34, 0.86]	.01
Age: 55+ (vs. 18-34)	0.05(0.29)	1.05	[0.60, 1.84]	.87
Gender: women (vs. men)	-0.34(0.21)	0.71	[0.47, 1.07]	.10
Edu.: ≥Bach. (vs. ≤high school)	0.14(0.36)	1.15	[0.57, 2.36]	.69
Edu.: Some College (vs. ≤high school)	0.02(0.36)	1.02	[0.51, 2.08]	.96
Income: 100+K (vs. <50K)	-0.13(0.27)	0.88	[0.52, 1.48]	.63
Income: 50-100K (vs. <50K)	-0.06(0.24)	0.94	[0.58, 1.51]	.79
Breach age (years)	-0.10(0.04)	0.90	[0.84, 0.96]	.003
Intention: yes (vs. no)	2.42(0.32)	11.24	[6.25, 21.74]	< .001

Table 9.7: Logistic regression for predicting password change action in the follow-up survey.  $n=615$  after excluding incomplete or not applicable responses. Cox and Snell  $R^2=0.21$ . Model  $\chi^2(22)=141.59$ ,  $p<.001$

(Q42). For those who indicated they changed the breached password ( $n=319$ ), we further asked how soon they changed the password since taking the main survey (Q43), what they used for the new password (Q44), and what techniques they used for remembering the new password (Q45) to characterize their password change behaviors, using survey questions from prior work [180, 327].

For password changes regarding other accounts, almost half of participants (48.9%) left passwords for other accounts the same. Fewer participants prioritized changing the password for other accounts using the same or similar passwords (18.5%) or for accounts they thought were important, such as bank accounts (16.3%). Only 6.1% participants said they changed the password for every online account. The remaining 10.2% described their action for other accounts via free text, mostly mentioning that they had the habit of changing passwords regularly (e.g., *“I often change all my online passwords at least once or twice yearly”*) or they changed the password for other accounts affected by breaches after checking HIBP (e.g., *“I already use different passwords for most accounts, but I did check the list of hacked sites and I think I changed the password on one”*). Our findings are similar to results by Golla et al. [180] in that participants developed their priorities in changing passwords for other accounts and rarely changed passwords across all accounts, but even more participants in our study exclusively focused on changing the password for the account affected by the breach.

Regarding password change timing, most of our participants who changed their password did it very promptly after receiving our nudges in the main survey (mean: 1.46 days; median: 1; sd: 2.24). Compared to prior work, our participants reacted to the password change prompt much faster—e.g., the mean time taken to reset the password was 26.3 days in Huh et al.’s case study of the LinkedIn breach [231], and only 13% of participants in Bhagavatula et al.’s study changed their passwords within three months of the breach announcement [46]. The differences could come

from the study’s methodology: we followed up with our participants two weeks since they received the breach notification, whereas in Huh et al. the time gap between the breach date and data collection was several months [231]; we proactively reached out to our participants to ask about their password change, whereas Bhagavatula et al. derived their findings based on naturalistic observational data [46]. The finding implies that our nudges effectively prompt participants to follow through with their intention, and participants who forgot or did not have intention were likely to leave their password as it is without additional reminders.

Regarding participants’ strategies in creating the new password, about half of our participants (50.2%) used a password completely unrelated to the old one they created. Some participants (32.9%) used a unique, random password generated by a password manager. Fewer participants were exposed to risks of password reuse attacks in creating the new password, either by changing a few characters in the old password (8.8%) or by using a password that they already used for other accounts (4.7%). The remaining 3.4% self-described their new password, such as following their own password creation heuristics (e.g., *“the same password schema I use for other sites, which is known to me but generates a different password for different sites”*) and using a random password created by themselves without trying to remember it (e.g., *“I won’t use that account again so I just typed in a lengthy string of numbers. If I want it again I will just recover the password”*). Compared to findings in Golla et al. [180], our participants had stronger new passwords and much fewer reused their old passwords. This is likely because our participants were creating passwords for their own accounts in real life rather than in a hypothetical scenario. The message in our coping appeal (Figure 9.3) could have also helped as it specifically discouraged password reuse and linked a guideline for making strong passwords.

Participants could choose multiple options to indicate their strategies for remembering the new password. Participants’ strategies were diverse; the most popular



options were saving the new password in the browser (27.3%), remembering the new password without writing it down or storing it digitally (25.4%), using a third-party password manager such as 1Password and LastPass (21.0%), and writing the new password down on paper (17.9%). Less commonly, participants stored the new password in a digital file (12.2%), used a system-provided password manager such as iOS Keychain (10.3%), or planned to reset the password every time they logged in rather than remember it (2.5%). Our findings align with Mayer et al. [327] and Pearman et al. [389]: participants' strategies for password management were primarily using a password manager and trying to remember it mentally.

For participants who indicated they had changed the password, we asked them to optionally upload a screenshot of the password reset confirmation email to verify the validity of their responses. We adapted the method from Huh et al. [231] by making the question optional but having a \$1.00 bonus payment to incentivize participants and honor their time. We also explicitly reminded participants to double-check the screenshot did not include sensitive or personal information in our instructions (Appendix G.1.3). Only 77 (24.1%) participants uploaded a screenshot; the rest either could not find the email (127, 39.8%) or chose not to upload a screenshot (115, 36.1%). In open-ended responses, participants mainly explained that they had deleted the email permanently (*"I deleted the email after receiving it, and it has probably been cleared from my Trash folder too since I try not to make my inbox too cluttered"*), they did not think it was worth the effort (*"I don't feel like searching through my email right now. Even though I would like the extra payment"*), or they were uncomfortable with the question (*"I don't feel comfortable sending my information to someone I don't personally know"*). Due to the low response rate, we did not use this question as a validation mechanism for participants' self-reported password change behaviors. While the question being optional could contribute to the low-response rate, we believed that it was a more ethical approach so that participants were not put in a position

<b>Intend to Change</b> ( $n=868$ )	<b>Count</b>	<b>No Intention</b> ( $n=518$ )	<b>Count</b>
bad things	226 (26.0%)	inactive use	189 (36.5%)
to be safe	195 (22.5%)	no account	152 (29.3%)
other actions	179 (20.6%)	no sensitive info	81 (15.6%)
inactive use	103 (11.9%)	forget password	45 (8.7%)
triggered by breach	95 (10.9%)	other actions	45 (8.7%)
action good idea	60 (6.9%)	unimportant password	43 (8.3%)

Table 9.8: Top reasons for intending to change or not change the breached password in the main survey.

to sacrifice their data for monetary gains unwillingly. The different findings between our study and Huh et al. [231] could also imply shifting norms among crowdworkers as they become more attentive and protective of their privacy [435].

### 9.3.4 Qualitative Insights

We analyzed participants’ open-ended responses regarding the rationale behind password change intention and action and their general reactions after learning about the breach. We found that intention and action occurred when participants had concerns over negative consequences or held a proactive attitude toward security. On the contrary, participants refrained from changing the password for an inactive account or resigned to inaction when they believed there was no account for which to change the password.

**Reasoning behind password change intention.** To gather insights into participants’ motivations and hurdles in forming password change intentions, we asked all main survey participants to explain why they intended or did not intend to change the password for the breached account (Q26). We summarize the top reasons in Table 9.8, differentiating between those with vs. without intention, and unpack the findings below.

Among participants who intended to change their password, 226 (26.0%) mentioned various negative consequences that they feared could arise from being affected

by the breach. The top potential consequences mentioned were personal information getting exposed, stolen, or misused (114, 13.1%) and account compromises (84, 9.7%). Some participants detailed how their personal information could be misused (*“My LiveJournal account contains fiction I wrote that I never want to lose. The knowledge that someone could delete all that work is quite perturbing”*). Other participants highlighted possible compromises of their other accounts when credential stuffing attacks happen. This concern was particularly salient among participants who reflected that they might use the password elsewhere or had the habit of reusing passwords (47, 5.4%); as one participant described, *“I want to change it because I use a lot of recycled passwords. Therefore, criminals could access other, more important accounts like e-mail accounts, bank accounts, my Amazon account, and other online retailer accounts”*). The finding also supports our previous quantitative results that participants who reused (or were unsure about reusing) the leaked password elsewhere were much more likely to form password change intention (Table 9.6).

A common theme that exclusively applied to participants with intention was a proactive attitude toward staying safe and secure (195, 22.5%), usually in the expression of *“just to be safe,” “to have a peace of mind,”* and *“better safe than sorry.”* Such a proactive attitude could play a dominant role in participants’ reasoning even when they thought the risk was limited or the account was unimportant; as one participant explained, *“I’m not too concerned about my MySpace info being hacked as it was such a long time ago and any of the information I entered then would not be relevant at all now. I will however change my password if that is an option just to be on the safe side.”* In addition, 95 (10.9%) participants mentioned that they intended to change the password as a result of learning about the breach and being reminded that they had an account on the breached site; some explicitly highlighted that our nudges convinced them to do so even though they had received breach notifications before (*“I changed my password based off of the previous screen in this survey. I had known*

*about the breach prior to taking this survey but didn't feel prompted to change my password as I haven't used Chegg in years").*

In addition to, or even on top of, changing the account password, participants showed interest in taking other actions (179, 20.6%), mostly deleting the account (119, 13.7%). Participants might expect to change the account password in the process of deleting the account, but account deletion is usually the ultimate goal (*"I don't remember making an account so I plan to do a password reset to gain access and then delete the account, if there is in fact an account as it says"*). The intention of account deletion largely depends on participants' evaluation of the account's utility and possible use; as one participant wrote, *"I will probably return to the site to determine if I am still interested in maintaining an active account. If so, I will attempt to log in and change my password. If not, I will try to contact the site to get my account and associated data disabled and/or removed."*

Switching to the rationales behind participants who did not intend to change the password, the most popular reason was inactive use of the account (189, 36.5%), i.e., the participant stopped using the account or still used it but very rarely. Comments on inactive use often came together with comments that the account did not have sensitive, important, or relevant information about themselves anymore (81, 15.6%) or that they did not care about the account even if it was compromised (42, 8.1%). As one participant reasoned, *"This account isn't important, and I'm pretty sure I haven't logged into it in almost a decade. The information contained in it would be minimal since I never really shared personal details or provided accurate information to websites for certain questions."* Similarly, 43 (8.3%) participants noted that the password used for the account was an old, unimportant, or "thrown-away" password and had limited impact even when leaked (*"The password I used for it is an old and way too simple password that I don't use for my important accounts. I have used that password for some accounts like free online games but that is it"*).

Participants also commented on the practical constraints that made password change impossible or difficult. 152 (29.3%) participants believed that they did not have an account with the breached site, as they did not even recognize the site's name, they recognized the site but did not remember making an account, or they believed they had already deleted the account (*"I don't have an account with MyFitnessPal – never have, so this is kind of confusing to me"*). 45 (8.7%) participants mentioned they forgot the password, which creates another barrier for them to change the password even when they knew password reset was an option (*"I don't really use this account much anymore and it's likely I don't remember my password, so I don't really want to reactivate the account just to change the password"*).

Looking at the overlapping reasons in Table 9.8, we found that inactive use of the account was a common theme across the board. Among participants with password change intention, 103 (11.9%) mentioned inactive use, but were still motivated to change the password due to fears of negative consequences (*"I do not want hackers to have access to an account that bears my name"*) or prior negative experience (*"I recently became aware of a similar hack with Chewy.com and changed my password and removed all payment data"*). Considerations of other actions were another common theme mentioned by those with (179, 20.6%) or without (45, 8.7%) intention, and account deletion was the top considered action for both groups. By contrast, whether the account has sensitive information could be a differentiating factor: 81 (15.6%) participants without intention referred to the account's lack of sensitive information, whereas 20 (2.3%) participants with intention were alerted by the types of information the account had about themselves (*"My data is very important to me. Even if it were to only be my date of birth, name and password, I would change it. I gave my address and other important information"*).

<b>Pw. Changed (<math>n=319</math>)</b>	<b>Count</b>	<b>Pw. Unchanged (<math>n=856</math>)</b>	<b>Count</b>
to be safe	91 (28.5%)	inactive use	283 (33.0%)
bad things	76 (23.8%)	no account	169 (19.7%)
other actions	41 (12.9%)	no sensitive info	86 (10.0%)
triggered by breach	41 (12.9%)	try and fail	86 (10.0%)
inactive use	20 (6.3%)	forget to do	79 (9.2%)
to access account	18 (5.6%)	other actions	74 (8.6%)

Table 9.9: Top reasons for changing or not changing the breached password in the follow-up survey.

**Reasoning behind password change action.** In the follow-up survey, we similarly asked participants to explain their reasons for changing or not changing the breached password (Q41). Table 9.9 summarizes the top reasons, divided by those who changed vs. did not change the password. Participants’ reasoning for action mostly aligns with the reasoning they gave for intention but with a few notable differences, as we discuss below.

Among participants who ended up changing the password, a proactive attitude toward staying safe and secure (91, 28.5%), concerns over negative consequences (76, 23.8%), and being alerted by the breach (41, 12.9%) continued to be the primary motivators. Participants continued to describe other actions they took (41, 12.9%), mostly deleting the account (22, 6.9%) and changing the password for other accounts that were subject to password reuse or for important accounts that they cared about (15, 4.7%). In pursuing other actions, 18 (5.6%) noted that they changed the password through the password reset feature to gain access to the account. However, changing the password was not their primary goal (“*I reset the password for my account, then went into settings and began the process of deleting the account*”).

Looking at the reasons given by participants who did not change their password, participants continued to refer to the account’s lack of use (283, 33.0%) and lack of sensitive information (86, 10.0%) to demonstrate the account’s irrelevance or insignificance. The absence of an account continued to be a practical constraint that made password change unrealistic (169, 19.7%). Consider a quote like “*I haven’t used the*

*site in many years so I'm sure that my account has been deleted due to inactivity,"* it is possible that the account still exists and the participant just forgot about it. Nonetheless, 86 (10.0%) participants confirmed that the account no longer existed or tried to log into the account but failed. Sometimes the site did not recognize the email address they provided (*"I went to Adobe site to see if I could sign in with my email address and it got [an] error message that there was no account associated with my email"*). Some participants never received the password reset email (*"I requested to reset my password but they never emailed me and it wasn't in my spam or trash email folders"*). Other participants found out the site was down (*"Heroes of Newerth no longer exists in any form. No website, no working game client, nothing. They basically have keys to a lock that is long gone"*).

Failing to recover account access assured participants that they might not have an account in the first place or had deleted the account. The proliferation of data brokers might also contribute to the problem; as one participant speculated, *"I don't think my Bonobos account ever existed, for example, and when I tried to reset my password, the email never came. This happened with other sites too. Maybe these sites had bought my information somehow from some other place, and therefore were able to obtain a username/email and password from me despite me never opening an account with them directly."* Even when account recovery might be an option, some participants could not find it due to problematic interface design (*"No, I couldn't even verify that I had a Zynga account. I think it told me to go to the specific game in question, but they have hundreds of them, and I didn't recognize them. So I have no idea what's going on"*). These friction points add to prior work on users' struggles in account management [201], generate frustrations, and could make participants abandon their plan (*"I tried to log in to change the password/delete the account, but couldn't get into it. I didn't try very hard, though. When I couldn't get past the forgotten password piece, I just gave up and left it abandoned"*). As much as we sought to address the costs

of password change in our coping appeal (Figure 9.3, “It only takes a few minutes. Just follow these easy steps”), our findings show that the time and effort required in password change are highly individual-dependent, and claiming the ease of password change does not apply to everyone.

Another notable theme among participants who did not change the password was that they had the intention but simply forgot about it or got distracted (79, 9.2%). Participants might delay the action because they did not take password change as a priority (“*I intended to change it but I procrastinated and eventually forgot all about it*”), but the forgetfulness could also be due to competing needs and valid challenges in their daily lives (“*My best friend passed away on July 29th and then I flew to Arizona for her funeral. I completely forgot about this*”). Along this line, 38 (4.4%) participants mentioned they did remember that they needed to change the password but were too busy or occupied with other duties (“*I was not able to follow up, [I] have had an on-going family situation taking up my time*”). The finding about participants forgetting about or putting aside password change is not surprising—prior work has well-documented the secondary nature of security and privacy mitigations, along with possible countermeasures such as commitment nudges [170] and implementation planning nudges [484, 485]. However, it could be challenging to find a balance between sending useful reminders and annoying users: participants who are motivated to follow through with their intention might appreciate such reminders, but a reminder could be meaningless to those who do not see the value of password change.

**(In)actions in response to the breach.** In the follow-up survey, we asked participants if they did anything in response to learning about the breach (Q38) before probing about password change specifically. Almost half of our participants (538, 45.8%) said they did not do anything; the rest (637, 54.2%) reported various actions,



<b>Action Taken</b> ( $n=637$ )	<b>Count</b>
change other passwords	174 (27.3%)
delete this account	90 (14.1%)
try and fail	86 (13.5%)
check breach records	65 (10.2%)
check password reuse	50 (7.8%)
check/delete info in account	39 (6.1%)
two-factor authentication	19 (3.0%)

Table 9.10: Additional actions participants took in response to learning about the breach.

and Table 9.10 summarizes the most popular actions as alternatives to changing the account password on the breached site.

As participants recognized threats of password reuse, 174 (27.3%) participants went above and beyond to change the password for other online accounts. Some participants made the change to all accounts (*“I ended up changing all my passwords because I used that password for Canvas for everything else”*) while others prioritized important accounts (*“I changed my Gmail password since that is the only account I’m particularly worried about”*). In determining which accounts were at higher risk, 65 (10.2%) participants obtained a more comprehensive list of breach records to guide their password changes (*“I changed my password on Houzz, but I also checked the website, ran a few checks via Google and Firefox’s password/breach checker, and found other breaches as well”*). 50 (7.8%) participants reviewed their existing passwords to identify those that were reused and should be changed first (*“I did nothing with respect to Zynga, but did spend a day updating all my passwords, including eliminating all duplicates, with my password manager”*).

In line with participants’ stated intentions, deleting the breached site’s account instead was another popular action (90, 14.1%), especially for accounts that were rarely used, as one participant recounted, *“After the previous survey, I did change my password. Later that day, I decided to just delete my account. I never really did use that site anyway, and it just got me thinking that I really didn’t need to have any info*

*at all on it.*” Many participants who chose to delete the account followed the logic that since the account no longer exists, they no longer need to worry about it being compromised, and the service provider now has less information about themselves (“*I was going to change my password but I decide on deleting my account instead. I don’t use this account anymore and if I delete it, no one can get a hold of it*”). Nonetheless, deleting the account does not mitigate the risk of credential stuffing attacks on other accounts using the same or similar passwords. When participants mentioned deleting the account but nothing else, the key issue was whether they reused the same password for other accounts. If yes, focusing on the account with the breached site exclusively could leave out additional risks. Otherwise, deleting the account could be a reasonable and sufficient action.

Some participants expected they might still need to use the breached account in the future. In this case, they did not delete the account, instead checked or modified the content in the account to ensure there was no sensitive information (39, 6.1%). As one participant detailed their experience, “*I immediately went to my LiveJournal account, where I changed my password and the email address associated with the account. I also confirmed that all posts which contained potentially identifying information were either deleted or set to private.*” Other less commonly taken actions include checking or deleting other inactive accounts (28, 4.4%), enabling two-factor authentication for important accounts (19, 3.0%), informing friends and family of the breach (10, 1.6%), researching cybersecurity topics and tools (10, 1.6%), and starting to use a password manager (9, 1.4%). Even when participants did not change the password for the breached site, these alternative actions could still increase their knowledge of online self-defense and lead to better security outcomes.

Usability issues remain a major hurdle in participants’ attempts to take action. Confirming our prior findings about reasons for not changing the password, 85 (13%) participants documented how they tried to log into the breached account but failed.

Changing the password across different accounts remains a burdensome process, even with the help of tools like HIBP and password managers. In addition, participants were still apprehensive of the marginal benefits compared to the efforts required (*“I only did a few of them. The whole process becomes such a pain in the ass given the various devices I have. And all of the sites in question are non-financial”*). The task became even more daunting for someone who had to do it manually (*“I tried to make a list of accounts I might have out there with that email and password. I didn’t change any passwords because, honestly, it was taking a long time, and I bailed on the task”*).

## 9.4 Discussion

Through a longitudinal experiment, we compared the effectiveness of nudges that incorporate a threat appeal, a coping appeal, or both with a control condition, focusing on their impacts on participants’ intended and actual behaviors around password changes after data breaches. We found that a threat appeal alone was most effective at motivating password change intention, and threat and coping appeals combined were most effective at motivating password change behavior; both generated a statistically significant yet small difference compared to the control condition. Participants further described challenges they experienced in attempting to change their passwords and alternative strategies they developed when they did not see the benefit of changing the password for the particular breach site. We next discuss our study’s limitations before summarizing how our findings contribute new knowledge to PMT literature and implications for designing password change notifications.

### 9.4.1 Limitations

Our work has multiple limitations. First, we situated our nudges in breach records curated by HIBP. While HIBP is reputable and has generated large impacts (e.g., being integrated into Firefox Monitor and 1Password), the database is built on scans

of account credential dumps and pastes, thereby consisting of mostly username and password breaches. Other types of personally identifiable information, such as social security numbers and medical records, rarely appear in HIBP’s database, but they do appear in other databases of breach records [236]. Some of our findings could be an artifact of us using HIBP as the source of breaches and our random sampling of breaches. For example, many participants mentioned their accounts with the breached site were old, rarely used, or contained little sensitive information; therefore, they did not feel motivated to change the breached password to protect the account. We showed participants a password breach selected randomly from HIBP rather than breaches that exposed more intimate details of themselves exclusively, which we expect might trigger different reactions. Future research could consider building a tool to curate breach records from multiple sources or partner with companies or non-profit organizations to achieve this goal.

Second, we conducted our experiment only with participants recruited in the US. We made this decision because the messages in our threat and coping appeals were highly language- and locale-dependent—e.g., identity theft in the US is largely fueled by the problematic way of social security numbers being used as identifiers [238]. Similarly, we relied on Prolific for recruitment, which led to the exclusion of less tech-savvy participants such as those without Internet access. These recruitment criteria limit our findings’ generalizability beyond the US and introduce opportunities for future research to replicate our study in different locales with different legal and consumer protection frameworks around data breaches.

Third, we recognize that some of our survey questions were not perfect measures for their intended constructs. As discussed in Section 9.3.3, we had to throw away responses about the account’s information type due to the absence of a “none of the above” option. Along this line, our questions about the account’s age, the purpose of use, and perceived importance might not capture all nuances as participants’ account

usage evolves.<sup>5</sup> This might also explain why none of the account-related factors were significant predictors in the regression models. Still, participants described a lot of benefit-cost analyses around the account’s utility in open-ended responses. As lessons learned, we suggest that future work could consider specifying the timing (e.g., ask “When was the last time you checked the account?” rather than “How often do you check the account?”) and using Likert scales rather than single Likert items for concepts like perceived account importance to more accurately characterize users’ diverse account usage behaviors.

#### 9.4.2 Theoretical Contributions

##### **The need for both threat and coping appeals in driving password changes.**

Our findings contribute insights to the ongoing discussion about whether the threat or coping appraisal plays a more prominent role in individuals’ formation of protection motivation. Specifically, we observed an interesting pattern that compared to the control condition, the threat appeal alone performed significantly better at raising higher password change intention, but the threat and coping appeals combined performed significantly better at encouraging actual password changes. The finding that threat and coping appeals combined performed better than the control condition at encouraging action also appeared in prior work in other security and privacy contexts, such as the adoption of secure mobile payments [484] and Tor browser [485]. The wider psychology and health literature on PMT has similarly suggested that strong appeals work only when accompanied by equally strong coping appeals [584].

A possible interpretation of this finding is that threat appeal alone is enough to raise intention, and coping appeal might even weaken intention by reminding partic-

---

<sup>5</sup>One participant wrote, “I had a hard time answering one question: ‘To the best of your memory, how long have you been using your LiveJournal account?’ I used it all through college, and then forgot about it for 13+ years, until just now? So I used it for a few years, but it’s not true that I’ve ‘used it for a few years,’ but it’s definitely also not true that I’ve used it for 13+ years?” Unfortunately, we got this response in the middle of data collection, so we could not make further changes to the question.

ipants of the effort required for changing the password. However, once the intention is formed, participants benefited from the coping appeal to overcome challenges in the intention-behavior gap [453]. Notably, our interpretation is speculative and could benefit from future replications to validate this proposed cognitive model, such as by testing the relationship between threat appeal, coping appeals, and various PMT constructs via structural equation modeling. We also envision experimental research that validates the role of coping appeal in bridging the intention-behavior gap. For instance, future work could take inspiration from prior literature on action and coping nudges [485] and test their effectiveness in motivating password changes. An action nudge could be a reminder to users about changing their passwords, and a coping nudge could help users persevere and troubleshoot when they cannot find their accounts or when the website does not respond to their password reset requests.

**Comparing the effectiveness of threat vs. coping appeals.** We provide two speculative explanations for why the threat appeal rather than coping appeal was significantly more effective than the control condition at the intention stage. First, our findings suggest that the threat appeal might have more significantly impacted the PMT constructs than the coping appeal. Looking at the manipulation check results, participants who saw the threat appeal had not only higher perceptions of threat vulnerability but also higher perceptions of response efficacy (which is a component of the coping assessment). By contrast, participants who received the coping appeal did not have significantly higher perceptions of response efficacy or self-efficacy or significantly lower perceptions of response costs. Our qualitative results indicate a similar pattern: participants who intended to change their passwords or ended up changing their passwords often commented on concerns over negative consequences and a feeling of “wanting to be safe.” whereas much fewer participants noted that changing the password was an easy act or commented on other aspects about coping.

Second, our findings suggest that a coping appeal (or our coping appeal in particular) might not apply well to password changes as the nudging goal or action of interest. Most participants already had high self-efficacy in their ability to change the password (Table 9.5), leaving little room for further increase. The response efficacy of changing the password largely depends on the threat they perceived, which further relates to whether a threat appeal is present and the participant’s password reuse habit. More fundamentally, no matter how much we emphasize the ease (i.e., low response cost) of changing the password, it would not apply to someone who has to go through all the hassles of recovering the account or someone who does not have an account to start with—both cases had a high presence in our qualitative findings. The focus on password changes might also contribute to the differences between our findings and prior work on PMT. For instance, van Bavel et al. found that the coping appeal was more effective than the threat appeal in nudging participants toward secure online purchases [542]; yet the experiment was conducted in a simulated setting, which means participants would not be able to experience the potential errors and usability issues that could occur to their real-world online accounts and impact their coping behaviors subsequently.

**The limitations of applying PMT to motivate password changes.** Our work illuminates the limitations of PMT when we look into factors that explain the additional variances in participants’ password change behavior and the hurdles participants experience in changing their passwords. Prior work that integrated PMT with the Theory of Planned Behavior has showcased how attitude influenced security intention and behavior [237, 463, 542], yet “attitude” is a broad term—the subject that the attitude applies to has many possibilities, from the recommended action to risks in general. Our work more precisely highlights the importance of security attitudes both quantitatively (evidenced by SA-6 being a significant predictor in both the in-

tention and action regression models) and qualitatively (evidenced by the prevalence of “to be safe” in participants’ reasoning about motivators).

An additional context-specific factor that impacts password changes was the breach’s age: the regression results suggest that participants were less likely to change the account password for breaches that happened a long time ago. Our qualitative analysis further provides insights into why, as participants reasoned that so much time has passed that the breached information was no longer relevant (*“The breach happened so long ago that I’m sure that password is not being used anywhere else for my accounts online. Along with the other compromised information, it’s no longer the same, except my name which is common knowledge and published on the internet anyways”*) or the lack of negative consequences so far assures them nothing will happen in the future (*“I feel if something malicious were to happen as a result of this breach it either already happened, or will never happen”*).

Moreover, while PMT exclusively focuses on appraisals and hurdles in human cognitive processing, our qualitative results demonstrate the larger systemic issues in the password ecosystem and digital platforms more broadly. Forgetting or not knowing the old password was a common theme for inaction, yet when an average American user has over 150 online accounts that require a password [111], keeping track of different passwords for every single account is fundamentally challenging, especially for those who have not developed a password management system [158, 388]. In attempting to change their passwords, participants encountered many issues that were technically out of their control, e.g., the site told them there was no account that matched the provided email address, they struggled to find the password reset button, or they struggled to identify which account to change when the breached company owned many products that the account could belong to. Some of these issues might be unintended consequences of the site’s attempt to comply with privacy regulations, e.g., when the site deletes inactive accounts to meet GDPR’s data minimization re-



quirement [144]. But others were akin to the usability issues of privacy controls in prior work [197, 198], and could even be viewed as dark patterns that sites deploy to coerce users into staying in business with them [356].

### 9.4.3 Practical Implications

At the time we conducted the research, compromised credential checking has been adopted in browsers and browser extensions (e.g., Google Chrome and Microsoft Edge) [284, 406] as well as password managers (e.g., 1Password and LastPass) [169, 456] across desktop, tablet, and mobile devices. Some participants even reported using these tools to check for potential password reuse, indicating that these are promising venues to deploy our nudges in more naturalistic and higher-impact settings. Our findings confirm and add to prior work on guidelines and best practices for designing password notifications, as we discuss below.

**Highlight threats, but with caution.** Prior work has surfaced misunderstandings users may have in assessing the risks of data breaches and password reuse [46, 180, 329, 602] as well as ideas to address them. For example, Golla et al. advocated that the notification should encourage changing similar passwords on other accounts and thoroughly explain why doing so mitigates password reuse attacks [180]. Huang et al. drew specific recommendations for Chrome’s compromised credential notification, suggesting that there should be more explanations about why it is necessary to change the breached password and what risks the user may face if they do not change their passwords [229]. As we incorporated these suggestions in developing our threat appeal, our findings about the effectiveness of the threat appeal confirm that it is important to help users recognize the risks of breached credentials, and doing so can motivate them to take concrete actions. The action might not apply to the specific breached site (Table 9.10) but could still lead to better security outcomes as

participants start to adopt two-factor authentication as an additional layer of defense or delete other inactive accounts to remove their digital footprints.

Nevertheless, product developers and service providers should consider the returns and potential unintended consequences in implementing the threat appeal. Our findings demonstrate that the threat appeal alone only performed marginally better than the control condition for intention, and the difference no longer existed regarding password change behaviors when taking away the coping appeal. The small effect size raises the question of whether iterating, implementing, and evaluating the threat appeal is worth the effort if a plain notification (which also means less text for users to read) can achieve the same purpose most of the time. To avoid overwhelming users, prior work has recommended providing information in a layered form [100, 139, 229], and we indeed incorporated this in our visual displays of the nudging text. Unfortunately, we cannot know if the layered approach truly results in an improvement in our study since we did not conduct A/B testing with this feature, but this could be an opportunity for future research, especially considering that the specific layer designs will differ a lot between different products and interfaces.

Furthermore, there might be inevitable tensions between using a threat appeal and making the notification trauma-informed. Prior work has highlighted that a too strong threat appeal could be counter-productive, especially for those already in a vulnerable state: they may engage in risk denial or simply refuse to engage with the fearful message as a self-protective mechanism because the threat makes them feel uncomfortable [433]. Our nudges mostly received appreciation from participants for informing them of the breach, and no one explicitly commented on the message being too triggering. Yet some participants recounted how busy they were or how they got distracted by other things, and this finding suggests how an inappropriately designed threat appeal could generate unnecessary burden or anxiety. One participant brought up other duties they were struggling to deal with in life (*I have saved the information,*

*but I was ill and I had to work three days 12 hours shift in a row, so I haven't had a chance to do anything, yet"); another mentioned the breached site reminded them of family members who had passed away ("I don't do any business with Dave. If this is true then it would have been my son who used it on my phone when his was broken. My son is now dead (suicide), so I cannot ask him"). These findings reinforce our previous recommendation that threat appeals should ideally be used together with coping appeals so that the coping appeal could ease potential stress reactions triggered by the threat appeal (e.g., by telling the user that they could pick a later time to deal with the issue and asking them if they want a reminder for this).*

**Communicate data flows to mitigate distrust.** A major factor that prevents users from adopting password managers or adhering to security advice is trust issues. In Huang et al.'s study, many concerns participants had were related to Google as they assumed that Google checks their non-saved credentials or worried about Google taking control over their own data [229]. By contrast, none of our participants expressed concerns about HIBP (which we explicitly highlighted as the source of the breach records we showed), perhaps because of the service's non-profit nature. Interestingly, our participants' concerns mostly centered around the breached site, e.g., when they were unfamiliar with the site's name (*"I never heard of them and am not going to ask for problems. This could be a scam"*) or when they were concerned about additional data leaks that could happen as they sought to recover their account in order to change the password (*"I don't really want to make the account 'active' again and get new spam from Chegg or risk it getting hacked again"*).

These findings highlight the necessity of providing more transparency to mitigate potential questions, distrust, and misunderstandings among users, but how to do so well remains the key challenge. Prior work has recommended that browsers and password managers that send compromised credential notifications should explain

how the product detects leaked or reused passwords [229] and, on a higher level, what measures the company takes to protect users' data [409]. Service providers are already doing this to some extent in their messaging,<sup>6</sup>. Still, the deeper issue is users' lack of awareness of and trust in encryption [115, 589], which will likely cause comprehension problems for such messaging. Helping users build trust with the breached site is more tricky, as users have reasonable doubts about the site for leaking their data in the first place. Proactively reaching out to affected users might help, as our participants expressed that they would rather learn about and deal with the breach compared to stay uninformed (*"I do not remember receiving any notice from Poshmark. They can't even spend the resources to notify us. There is no responsibility for protecting the data of customers"*). Suppose users trust the site enough to check it out. In that case, it might help to provide additional explanations of when and how their account was created (if it still exists) or why their account no longer exists (e.g., explaining that the account was deleted due to inactivity if that is the case, rather than leave it up to the user to guess).

**Standardizing and semi-automating the password change experience.** The various challenges our participants experienced in attempting to change their passwords illuminate larger problems with the password ecosystem. As much as we tried to predict edge cases in developing our coping appeal (e.g., "Unsure if you have a [site name] account or can't log into it? Contact [site name] to recover the account or have your account deleted"), our findings reflect that the real situation is far more complex. Participants could encounter all kinds of scenarios in which things could go wrong. It then becomes difficult to predict all scenarios and enumerate corresponding solutions in the coping appeal.

Our findings reflect that the experience of changing a password could be drastically

---

<sup>6</sup>E.g., see "Privacy is at the heart of our design" in Google's blog post about its password checkup feature [406].

different from site to site. However, it should not be this way, and we see the need for standardizing the password change processes by providing more industry guidelines or pushing for stronger regulations. For example, the California Consumer Privacy Act (CCPA) requires businesses to include a “Do Not Sell My Personal Information” link in their privacy policy that allows consumers to submit an opt-out request (if the business sells personal information) [375]. We imagine similar efforts could be made to standardize where websites should provide the account login fields as well as information regarding how to change the account password. The CAN-SPAM Act in the US requires that for commercial messages, businesses should honor consumers’ opt-out requests within ten business days [497]. Similarly, we see opportunities for standardizing the turnaround time for sending password reset confirmations, given that one of the hurdles our participants experienced was unable to find the password reset email.

Beyond standardization, we see opportunities for partially automating the password change experience for consumers as our participants complained about the burden and effort required. For instance, for someone who already adopts a password manager, the password manager could provide a feature that allows the user to scan through every saved login credential and determines whether it still works for the corresponding site.

**Opportunities and challenges for providing personalized advice.** Our findings reveal that participants had diverse ways of using the account and strategies for managing their passwords. Changing the account password for the breached site was reasonable for participants who used the account extensively and felt motivated to protect it from potential compromises. However, participants who did not care about the account tended to delete the account instead, and participants who were already using strong and unique passwords for different accounts had valid reasons for not

changing the password since their risk levels were minimal. While standardized mechanisms can help consumers more easily change their passwords when they want or need to, our findings call for a deeper reflection on whether password changes should be the nudging goal for everyone given the diverse risk levels and preferences among individuals.

We see the promises of personalized advice to help consumers better assess whether password changes are needed in their particular situations. For instance, existing password managers such as 1Password and LastPass are already providing dashboards and scores that help users evaluate overall password strength and the extent of password reuse in their stored credentials [119, 282]. Going beyond these features, password manager could provide personalized and data-driven recommendations for which accounts to prioritize in cleaning up old or weak passwords (e.g., prioritizing banking and shopping sites that are more likely to contain financial information) since otherwise, users might not be able to recall all precisely on their own. In addition, advice on which actions to take should ideally tailor to the specific types of information compromised. As our study focused on password breaches, changing the password for the breached site and other accounts using similar passwords were the most intuitive and likely applicable actions. However, breaches that involve more sensitive or unique data types, such as social security numbers and medical records, might require more serious measures. Existing advice for what to do after data breaches mostly enumerates different actions depending on the different breached data types [250]. Nevertheless, a more efficient way could be presenting personalized advice about actions to take upfront rather than burden the user to figure out. The interactive resource provided by the US Federal Trade Commission [506] is a good step toward this direction, although the extent of personalization there is still quite limited.

As much as we see the exciting opportunities offered by personalizing post-breach coping advice, we should not ignore the challenges and open questions raised by per-

sonalization. Too much automation, while reducing users' burden, could reduce their sense of agency [88]. The personalized advice might be inaccurate or irrelevant, which may further cause trust issues or annoy users. When the advice is about password changes in particular, the personalization will likely need to be offered via password managers, which have an increasing yet still small user base. The dependence on password managers might also exclude specific populations like older adults, who are known to have more trust issues with cloud storage of passwords [408]. The questions then become: how do we find a balance between making the personalized advice useful versus learning too much about users' preferences to the extent of causing privacy concerns? How do we make personalized advice accessible to everyone and truly reflect users' diverse needs and preferences?

## CHAPTER X

# Conclusion and Future Work

In this dissertation, I bring together several threads of research to demonstrate how we can carefully, effectively, and ethically understand and support consumers' adoption of online privacy-protective behaviors across multiple contexts. I first summarize how my work contributes to prior literature, before reflecting on the lessons I have learned through this journey and potential directions for future work.

### 10.1 Summary of Key Contributions

Through in-depth interviews, surveys, and content analyses, the first part of my dissertation highlights the challenges consumers face in both reacting to data breaches and adhering to expert advice about security, privacy, and identity theft more broadly.

- In Chapter III, I document an interview study on consumers' reactions to the 2017 Equifax data breach, providing novel insights into people's reasons for inaction after a data breach. The findings illuminate how prior work on consumers' bounded rationality, heuristics, and biases in privacy decision-making applies to the data breach context, which was relatively understudied when we conducted the study. As consumers refrained from taking action, they exhibited potential optimism bias in discounting their personal possibility of suffering harm and



tended to delay action until harm had occurred. Furthermore, financial costs and usability issues could hamper consumers' adoption of certain protective measures.

- In Chapter IV, I present a study that similarly investigated consumer reactions to data breaches but using survey as a complementary method. One of the study's key contributions is the high ecological validity setting. Rather than asking participants to recall breaches they were affected by, we presented them with breaches known to have exposed their information using custom survey software and records from the Have I Been Pwned database. The study confirmed issues identified in the previous qualitative work about consumers' reasoning of breaches' limited impact on themselves but also contributed new knowledge of consumers' low awareness of data breaches: for 73% of the breaches, participants were not aware that their data had been exposed before taking our study.
- In Chapter V, I approach consumers' struggles with reacting to data breaches from a different angle via a content analysis of 161 breach notifications sent to consumers by breached companies. In line with prior work on privacy policies, we found that similar issues exist in breach notifications as companies frequently used hedging terms in describing whether the recipient was affected by the breach and the associated risk. The recommended actions were often presented in large chunks of text without highlighting their priority or efficacy, raising the question of whether these breach notifications could effectively help consumers cope with the aftermath.
- In Chapter VI, I extend the inquiry of protective behaviors' adoption from data breaches to a broader set of contexts related to security, privacy, and identity theft via an online survey with 902 US Internet users. The study surfaced demographic differences (e.g., in terms of gender, age, income, and

education) in behavior adoption. We further looked into behavior abandonment to get a complete picture of the adoption cycle, and our findings suggest that abandonment occurred when participants found the practice inconvenient (e.g., due to the requirement of intermittent engagement) or low-value (e.g., when it does not provide tangible rewards and better alternatives arise).

The insights in the first part of my dissertation have laid the foundation for developing approaches that encourage consumers to adopt privacy-protective behaviors, but any approach would need to consider context-specific consumer needs and constraints. In the second part, I demonstrate the development of these approaches in three contexts that address the needs of different populations and the challenges they face. I also discuss their connections with one another.

- In Chapter VII, I document a series of studies focusing on the design and evaluation of icons and accompanying link texts to convey the presence of privacy controls. This work represents a user-centered approach to designing privacy interfaces—by basing the icon and link text designs on participants’ needs and preferences—to help consumers overcome hurdles in exercising privacy controls. The work yielded icon and link text combinations that clearly conveyed the presence of privacy controls without generating substantial misconceptions. Our recommendations have directly influenced the California Consumer Privacy Act’s rulemaking process.
- In Chapter VIII, I show how encouraging the adoption of privacy-protective behaviors should be sensitive to the needs of at-risk populations by focusing on survivors of intimate partner violence (IPV). This work demonstrates that rather than develop a completely new approach, we could repurpose existing sociotechnical infrastructure toward the population’s needs—in this case, customer support at computer security companies. We engaged with IPV and cus-

customer support professionals to elicit their insights on incorporating IPV-related resources into customer support's existing practices. Key insights include using trauma-informed language, avoiding promises about solving problems, and making referrals to external resources for holistic safety planning. Based on the insights, we have developed guidelines and training materials for improving customer support and shared them with partnering companies.

- In Chapter IX, I draw on insights from my prior work to develop and evaluate approaches that encourage consumers to take action after data breaches—specifically, nudges toward changing the breached password using threat and coping appeals from the Protection Motivation Theory. By evaluating the threat and coping appeals together and in isolation compared to a control condition in a longitudinal field experiment ( $n=1,386$ ), we confirmed the promise of PMT-based nudges as the threat and coping appeals combined most effectively increased password change behaviors, but only by a small margin. Our study further suggests the limitations and additional considerations in deploying PMT-based nudges and providing password-related advice broadly, as the advice needs to be attentive to consumers' diverse needs and the practical constraints they face.

## 10.2 Reflection on Lessons Learned and Future Work

**Evolving interpretations of inaction.** Prior work has extensively documented the various cognitive and behavioral biases consumers face in privacy decision-making [4, 7]. The different studies across my dissertation contribute nuanced findings about why inaction occurs and the need for interpreting inaction with a critical lens rather than assuming that inaction is a bad outcome and consumers are biased or lazy.

For instance, the study in Chapter III referred to heuristics and biases and specif-

ically highlighted optimism bias and a tendency to delay action until harm has occurred as two examples of biases that prevent consumers from reacting to data breaches. However, the limitation of this study is that we suggested biases without contextualizing them in the ground truth of participants' objective privacy risks, i.e., whether they were truly affected by the Equifax data breach in this case. The Equifax breach indeed impacted almost half of the US population, and statistically speaking, some of our participants would be affected. Yet we did not confirm whether individual participants were affected by the Equifax breach. Consumers' inaction could be reasonable when their personal information was not exposed in this breach. In the studies presented in Chapters IV and IX, we managed to address this limitation by eliciting consumers' reactions to breaches that are known to have exposed their personal information. Inaction was still a common theme, as consumers rationalized their inaction or developed workarounds when the original recommended action did not suit their needs. These findings beg more questions about how we should view inaction: Is inaction necessarily bad? When is inaction justified, and when do we need to encourage action?

In answering these questions, the blurry line between risk and harm adds more complexity but also ideas for future research. Following the epistemology of risk analysis, risk is the probability of an adverse event happening, and harm is the consequences associated with the adverse event [586]. Risk and harm are closely connected, but not all risks will eventually lead to concrete harms. Yet what counts as harm is another topic up for debate. Legal scholars have argued that in the context of data breaches, the risk of possible future injury itself could be sufficient to establish harm because of the anxiety victims experience about the increased risk of future harm [475]. Arguments along this line are helpful for plaintiffs to win data breach lawsuits and establish stronger legal protections for consumers. Findings from my research, however, contribute complementary insights by showing that most consumers

are not concerned about the risks or harms of data breaches when the breach does not involve important accounts or sensitive information about themselves. In addition, most consumers focus on more concrete forms of privacy harm (e.g., take the breach seriously when the breach concerns financial information or implies risks of financial loss) but do not think of the abstract risk of future harms when the breach does not trigger anxiety.

Findings of consumers' reasons for inaction and the uncertainty in risk converting to harm have suggested that we need to avoid taking inaction as a reflection of consumers' incorrect risk assessments or that they are hopelessly lazy and unmotivated. Consumers' reactions to breaches should be contextualized in relation to the objective risk level; using breach records from the HIBP database is a good first step, and future work can seek to expand the methodology to other sources. Moreover, there is some extent of rationality in consumers' inaction—a low-income consumer struggling to make ends meet will likely deprioritize privacy-protective behaviors to ensure job security and housing first; a consumer whose account with the breached site no longer exists might as well focus on changing similar passwords for other accounts, or do nothing if they are already using strong and unique passwords.

Recognizing the rationality in inaction does not mean we should let inaction happen. Rather, a possible way forward is to give consumers the help and resources to take action when they are highly motivated to do so or when there is a high risk that they will suffer concrete harm from the breach. My prior work has contributed concrete recommendations for making data breach notifications useful, both for industry practitioners (e.g., design implications for compromised credential notifications in Chapter IX) and for policymakers (e.g., policy implications for strengthening regulations about businesses' post-breach responses and consumer protections they should be providing in Chapters III, IV and V). Going beyond the design space of breach notifications, future work could seek to improve a broader spectrum of tools that help

consumers monitor and recover from breaches by both examining the effectiveness and improving the usability of existing tools (e.g., credit or identity monitoring services, credit freezes), and by developing and evaluating new tools (e.g., integrating account checkups with password checkups in password managers). Another fruitful but more challenging line of work could be measuring the longitudinal risks and harms for consumers affected by breaches and using such insights to develop personalized advice for consumers to react to breaches (e.g., basing recommendations on factors such as objective risk level and consumers' own attitudes and preferences).

**Individual differences in adopting privacy-protective behaviors.** Prior work has identified demographic differences (or assumed differences) in consumers' privacy behaviors. However, findings on this topic are largely context-dependent and remain inconclusive. For instance, some studies have suggested older adults are more vulnerable to protecting their privacy due to difficulties in processing information or a lack of appropriate social networks [53, 142, 459]. Other studies have found no significant age-based differences [542] or demonstrated older adults' more rationally calculated privacy decisions [178], possibly because the older adult population itself is quite diverse. When it comes to gender, prior work has found gender stereotypes about security and privacy fueled by sexism (e.g., men are more skilled at protecting themselves than women) [572] and additional challenges faced by LGBTQ+ communities [175]. Some studies have found gender-based differences for individual protective behaviors [228, 381], but others did not [91, 386]. As for socioeconomic status, some studies have illuminated the unique challenges low-income consumers face in navigating privacy and security risks in online services [551], often compounded by limited technical skills. Other studies have surprisingly revealed that consumers' experiences with negative incidents were more correlated with their advice sources rather than educational attainment or resources [412].

My dissertation contributes to this line of discussion by examining demographic differences in consumers' adoption of privacy-protective behaviors (Chapters III, IV, and V) and considering demographic differences in evaluating the effectiveness of approaches to encouraging adoption (Chapters VII and IX). My findings reveal several recurring themes among low-income consumers in their reasoning of action after data breaches, such as "I've got nothing to lose" (as they felt demotivated or resigned to protect their already limited assets) and "I'm a small fish in a big pond" (as they expected that attackers would go after higher-income consumers for the higher return). Admittedly, there are misunderstandings in these reasonings, as we look at how credential stuffing attacks happen en masse to millions of accounts or how low-income consumers compromise at least thirty percent of all identity theft victims in the US [188]. Nevertheless, we should connect low-income consumers' inaction to their unique challenges, such as competing needs for financial security and limited access to costly privacy-enhancing strategies. More work is needed to protect low-income consumers by default (e.g., stopping the pervasive surveillance that targets low-income consumers and makes them more vulnerable) [310] and to make the recovery process more equitable after negative incidents [188].

Furthermore, findings across different chapters in my dissertation have underscored that the individual differences are highly dependent on the study's scope, sample, and method among many other factors, calling for caution in generalizing the results. For instance, the study in Chapter VI found that men had higher adoptions of practices for security, privacy, and identity theft protection broadly. Looking at studies focusing on individual contexts, we then found that women were more likely than men to form password change intentions after data breaches (Chapter IX), or there were no significant gender differences in consumers' preferences for privacy icon designs (Chapter VII). Most of the studies in my dissertation were conducted with US Internet users recruited from crowdsourcing platforms such as Prolific and MTurk.

While prior work has found that security and privacy survey responses collected via these crowdsourcing platforms are somewhat generalizable [413, 489], it is important to note that some protective behaviors we examined and corresponding findings are strictly US-specific (e.g., freezing credit reports being a recommended action for mitigating risks associated with a data breach). There is ample space for future work to validate to what extent the findings shift based on insights from non-WEIRD (White, Educated, Industrialized, Rich, and Democratic) samples [214].

Chapter VIII reflects a shift to examining vulnerability as a particular lens of individual differences in adopting privacy-protective behaviors. The work highlights how a caring approach is needed in aiding survivors of intimate partner violence as survivors face unique and persistent threats from their abusers [167, 325]. Additionally, even the mere perception of a threat can lead survivors to distrust or have traumatic stress reactions toward technology and further isolate from support [166, 319, 521]. We saw the promise of activating computer security customer support as a novel avenue for helping IPV survivors due to their expert technical knowledge. Meanwhile, we realized traditional ways of operating commercial customer support could be ill-suited for serving IPV survivors by potentially amplifying trauma and providing advice that could escalate the abuse. As such, we worked with IPV and customer support professionals to elicit recommendations on how to cater customer support to survivors' needs. This work has provided broader implications of how the possible effects of trauma and traumatic stress reactions should factor into user experience design, as my co-authors and I articulated in our recent framework of trauma-informed computing [80]. In particular, future research should consider making security and privacy tools, advice, and the entire support ecosystem more trauma-informed, especially given that security incidents and privacy violations could be highly traumatic.



**Interdisciplinary thinking improves the development of approaches.** In related work (Chapter II), I summarized that existing approaches that encourage consumers to adopt privacy-protective behaviors include providing privacy notices and choices and implementing privacy nudges. My dissertation contributes new knowledge to both areas by identifying user-tested icons to make privacy controls easier to find (Chapter VII), and by developing and evaluating nudges to encourage password changes after data breaches (Chapter IX) informed by knowledge of consumer reactions to data breaches (Chapters III and IV) and issues in breach notifications (Chapter V). Both lines of work have generated broader impacts. Our icon recommendations were adopted into the official California Consumer Privacy Act regulations [375]. Our line of work on data breaches has informed the presentation of advice in Firefox Monitor [343], and I have been invited to present my findings to the US Federal Trade Commission to share our research’s implications for future regulatory and enforcement practices around data breaches.

That being said, privacy notices, choices, and nudges are minor interface-level changes that influence individuals’ outcomes. Operating under the larger environment of surveillance capitalism [606], they are necessary but insufficient to protect individual consumers’ privacy [100, 555]. In parallel to seeking improvements in approaches within the existing notice and choice framework, I have sought to expand the epistemology of developing approaches that encourage adoption, such as those that go beyond individual and interface levels. Chapter VIII has demonstrated how computer security customer support, as an example of existing sociotechnical infrastructure, can adapt to IPV survivors’ needs and challenges informed by insights from relevant stakeholders. In work outside of this dissertation, I have examined the privacy needs of older adults and turned the work into online self-defense workshops tailored to older adults, delivered via partnering with local senior centers. As prior work highlighted the significant correlation between advice source and privacy expe-

rience [412], educational efforts could serve as an effective approach in encouraging adoption, especially for populations like older adults who may have relatively limited online experience but are similarly, if not disproportionately, impacted by risk factors.

Lastly, my dissertation has demonstrated the value of interdisciplinary thinking and context-specific approaches in addressing specific challenges that consumers face in adopting online privacy-protective behaviors. As shown in Chapter VII, eliciting input directly from consumers helps inform the design of intuitive privacy icons and effective privacy user experiences. However, in Chapter VIII, we considered that direct engagement with IPV survivors might cause unnecessary retraumatization by requiring them to remember and recount traumatic events. This is not to say that IPV survivors' insights are not important or they should be excluded from research processes—in fact, prior work has provided excellent examples of how to conduct participatory research with IPV survivors [291] and other marginalized or vulnerable populations [86, 121, 436] in careful, ethical, and trauma-informed ways. In our case, we chose not to directly interact with IPV survivors because there were other means for us to achieve the research goal: we had access to customer support chat logs that could help us better understand existing problems, and we were able to recruit IPV and support professionals who could contribute meaningful qualitative insights on improving customer support.

The study in Chapter IX further demonstrates how approaches informed by a well-established theory could encounter challenges in adapting to a new context (i.e., changing the password after data breaches). Evaluating the approaches with consumers in a rigorous, controlled experiment provides insights on whether and how to implement the approach in the real world, as our findings highlight the statistically significant yet marginal improvement of using threat and coping appeals to highlight the risks of passwords getting breached and the practical constraints participants faced in changing their passwords. Altogether, my dissertation highlights the value of

getting consumers involved in developing and evaluating approaches that encourage privacy-protective behaviors. Still, researchers need to make decisions about whether (e.g., weighing potential benefits and harms to the specific population), when (e.g., during or after the approach development stage), and how (e.g., qualitative explorations for understudied topics versus quantitative validations when there is a sufficient amount of prior work to establish hypotheses) about getting consumers involved on a case-by-case basis.

## APPENDICES

## APPENDIX A

### Supplemental Materials: “Consumer Reactions to the 2017 Equifax Data Breach”

#### A.1 Interview Protocol

1. Could you tell me how you manage your personal finance, such as income and credit cards? Has it changed over time?
  - (a) If yes, could you tell me any particular points that the change occurred?
  - (b) If no, could you explain why?
  
2. What’s the first thing that comes into your mind when you hear the term “credit bureau”?
  - (a) From your point of view, what do credit bureaus do?
  - (b) You just said credit bureaus do... How do they do this? Could you draw or sketch on the paper to make it clear? (A few prompts listed as below if necessary)
    - i. What information do they collect?

- ii. What parties do they share information with?
    - iii. What do they know about you?
    - iv. What information you can get from them?
    - v. What is their purpose?
  - (c) Could you name some credit bureaus?
- 3. Could you tell me your personal experience with credit bureaus?
  - (a) Have you ever interacted with credit bureaus directly?
    - i. If yes, when was the last time, and how was the experience?
    - ii. If no, could you explain why?
  - (b) What do you know about your credit history and credit scores?
  - (c) Have you ever checked your credit report?
    - i. If yes, when was the last time? What prompted you to check it? How did you do it? With which credit bureau? Only one or multiple?
    - ii. If no, why not?
  - (d) Have you ever checked your credit score? (if they have checked report, ask if credit report included credit score)
    - i. If yes, when was the last time? What prompted you to check it? How did you do it? With which credit bureau? Only one or multiple?
    - ii. If no, why not?
  - (e) Do you feel that credit bureaus have an impact on your life?
    - i. If yes, what is the impact?
    - ii. If no, could you explain why not?
- 4. Have you ever heard of Equifax?

- (a) If yes, what do you know about it?
  - (b) If no, "Equifax is one of the big three consumer-focused credit bureaus in the United States".
5. Equifax experienced a data breach in 2017. How much do you know about the data breach of Equifax?
- (a) Have you ever heard of this Equifax data breach before this interview?
    - i. If yes, could you describe what happened based on your understanding?
    - ii. If no, "It happened between May and July in 2017 and compromised the personal information (i.e. names, addresses, birth dates and Social Security Numbers) of over 145 million Americans".
  - (b) In your view, what are the potential consequences of this breach?
  - (c) What was your reaction when you heard about the Equifax data breach?
  - (d) How do you feel about your data at Equifax now? Did it change after the breach?
  - (e) Do you know if you were personally affected by this breach?
    - i. Do you know if data about you was exposed in the data breach?
      - A. If yes, how do you know?
    - ii. Did you check if you were affected?
      - A. If yes, how did you do it?
    - iii. Did you check your credit reports at any point since you learned about the breach?
      - A. When did you do it? How often?
      - B. How did you do it?

- C. Only at Equifax or also at other credit bureaus?
- (f) Do you know what you could do to protect your credit data in general?
- (g) Did you do anything to protect yourself in response to the breach?
  - i. Have you heard of fraud alerts?
    - A. Can you describe what it is?
    - B. Have you placed a fraud alert before or after the Equifax data breach?
    - C. Can you describe how?
    - D. With Equifax? With other credit bureaus? Which ones?
    - E. Did you pay money for it?
    - F. How long has the fraud alert been active for?
  - ii. Have you heard of a credit freeze?
    - A. Can you describe what it is?
    - B. Have you placed a credit freeze before or after the Equifax data breach?
    - C. Can you describe how?
    - D. With Equifax? With other credit bureaus? Which ones?
    - E. Did you pay money for it?
    - F. How would you unfreeze your credit?
  - iii. Did you start monitoring your credit and bank accounts more often since then?
    - A. Can you describe how?
    - B. With Equifax? With other credit bureaus? Which ones?
    - C. Do you pay money for it?
  - iv. Have you heard of identity theft protection?



- A. Can you describe what it is?
  - B. Have you signed up for any identity theft protection services?
  - C. Can you describe how?
  - D. With what company/entity?
  - E. Do you pay money for it?
- v. Did you do any other things not mentioned previously?
6. Before this breach occurred...
- (a) Have you ever experienced any data security problem, such as someone secretly changed your password?
  - (b) Have you ever experienced identity theft, such as someone applying for credit cards under your name?
- (For each question, if yes, follow up with “Could you tell me more about the experience? Do you feel it has any impact on you?”)

## A.2 Screening Survey for Recruitment

Thank you for your interest in our study! Please answer a few questions about your demographics and availability for the interview.

1. In which year were you born?
2. What is your current gender identity?  Male  Female  Non-binary/third-gender  Not listed (please specify)  Prefer not to answer
3. What is the highest level of education you have completed?  Less than high school  High school degree of equivalent  Some college but no degree  Trade, technical, or vocational degree  Associate’s degree  Bachelor’s degree  Master’s degree  Doctoral degree  Professional degree (JD, MD, etc.)  Other (please specify)  Prefer not to answer

4. Which of the following categories best describes your occupation?  Administrative support (e.g., secretary, assistant)  Art, Writing, or Journalism (e.g., author, reporter, sculptor)  Business, Management, or Financial (e.g., manager, accountant, banker)  Education or Science (e.g., teacher, professor, scientist)  Homemaker  Legal (e.g., lawyer, law consultant, or law professor)  Medical (e.g., doctor, nurse, dentist)  Engineering or IT Professional (e.g., programmer, IT consultant)  Service (e.g., retail clerk, server)  Skilled Labor (e.g., electrician, plumber, carpenter)  Unemployed  Retired  College student  Graduate student  Not listed (please specify)  Prefer not to answer
5. What was your total household income before taxes during the past 12 months?  Less than \$25,000  \$25,000 to \$49,999  \$50,000 to \$74,999  \$75,000 to \$99,999  \$100,000 to \$124,999  \$125,000 to \$149,999  \$150,000 or more  Prefer not to answer
6. What is your citizen status?  I am a citizen of the United States.  I am a permanent resident of the United States.  I am neither a citizen nor a permanent resident of the United States.  Other (please specify)  Prefer not to answer
7. (If answer to above question was “citizen of the United States” or “permanent resident”) How many years have you been living in the United States?  Less than 1 year  1-2 years  2-3 years  3-4 years  4-5 years  > 5 years  Prefer not to say

### A.3 Qualitative Analysis Codebook

Below is the codebook used for interview transcript analysis, grouped into four big categories.

## **Category: Financial Management**

**Financial status:** Description of general financial situation, e.g., income, number of checking/saving accounts, number of credit cards currently held, late payment, as well as the mentioning of occupation, big purchases (e.g., cars and mortgages).

**Financial tracking:** The way to keep track of earnings and spendings, manage different credit cards, use checks or do everything online, the way of paying bills (e.g., set up automatic withdrawals or pay bills whenever it comes).

**Financial behavior change:** Any particular change in the ways of managing one's finance, how and why it occurred, may also include behavioral change resulting from attitudinal change (e.g., I tried to spend less because I wanted to save money).

## **Credit Bureau Related**

**Understanding of credit status:** (1) The knowledge of the meaning and components of credit scores in general, how credit score is generated, whether it costs money to check credit scores, the mentioning that different bureaus may have different scores etc. (2) The impression of whether the participant's own credit score is good or bad, the description of when's the last time checking it and how to check it, where does the credit score come from (e.g., one of the three big bureaus or banks) (3) The impression of one's credit history, things included in the credit report, whether or not they have things like late payments and debts.

**Awareness of credit bureaus:** The number of credit bureaus, specific names of credit bureaus, also use this when they say they can't remember it or can't give the full name, also include the participant's knowledge or guess about whether there are bureaus other than the big three.

**Impact of credit bureaus:** "What impacts do credit bureaus have on you": how credit bureaus may impact consumer lives by giving credit ratings/scores or in other ways. Also include cases where participants say credit bureaus have little or no

impact on them personally because of various reasons.

**Check credit status at credit bureaus:** Directly contact credit bureaus to access credit reports or sign up for other credit-related products and services, description of the process (e.g., schedule times to make use of the free opportunity to check credit reports annually).

**Check credit status at other places:** Usually through banks and third-party financial aggregation app (such as NerdWallet, Credit Karma, and Mint) to check credit history, credit score, or credit status in general, and the reason for doing it (e.g., it's free and more convenient), the frequency of the received updates, whether or not it might be helpful.

**Reasons for no interactions:** Description of having little or no interactions with credit bureaus, didn't check credit status through either credit bureaus or other places, and the reasons for doing it, e.g., I don't need to make big purchases or I don't want to know my credit status because it's poor.

**Dispute process:** Anything related to the dispute system within the credit reporting system, can be (1) the general telling that consumers have the right to dispute incorrect information; or (2) the complaint that the current dispute system doesn't work to solve consumers' problems (e.g., they have to spend a lot of time filing the dispute and it's hard to get the error eventually corrected).

**Information providers of credit bureaus:** Companies and organizations that provide information to credit bureaus, e.g., government, IRS, lending companies.

**Customers of credit bureaus:** Entities to which credit bureaus share or sell individual consumer's information, who may have the access to consumer credit files at credit bureaus. Also include cases where participants may not explicitly mention it but rather say it's an information exchange process, e.g., "I think that banks quarry them but they would also ask banks about".

**Types of information collected:** The types of information credit bureaus col-

lect from their providers (e.g., checking accounts, savings, credit history, loans) about individual consumers, usually the answer following “what types of information do credit bureaus collect?” and “what do credit bureaus know about you?”

**Offerings of credit bureaus:** What information consumers can receive from credit bureaus, such as the annual free credit reports, credit reports that cost money to see credit scores, credit monitoring services.

**Purpose of credit bureaus:** This will refer to how credit bureaus use the collected information for, what their purposes are, e.g., assessing one’s creditworthiness, generating credit scores. Answers following the question “what’s the first word that you associate with credit bureaus” and “what are their purpose” might fall under this category.

**Errors in mental models:** This code encompass any obvious errors that we capture in participants’ describing of credit bureaus.

**Inaccurate credit files:** Specific instance of negative perception - the experience that credit bureaus get errors on consumer credit files or retrieve the file of the wrong person, and hence leading to bad or unpleasant experience for consumers.

**Opaque data aggregation process:** Specific instance of negative perception - mentioning of the process how credit bureaus collect and aggregate all different types of information as opaque, unclear, not idea about what’s going on behind the curtain.

**Abusive use of power:** Specific instance of negative perception - the mentioning that credit bureaus (and other related institutions such as governments and banks) are in the position of holding great power/have little interest in protecting consumer rights; consumers are in a relatively weak position.

**Insidious data collection:** Specific instance of negative perception - describing the data sharing between credit bureaus and data furnishers as passive, creepy or scary, without obtaining consent from consumers. As for consumers, they have limited control and choice over this kind of data collection.

**Positive perception of credit bureaus:** Positive description of credit bureaus in general, the statement that credit bureaus have a positive image in the participant's mind.

**Negative perception of credit bureaus:** Negative description of credit bureaus, the statement that credit bureaus have a negative image in the participant's mind, note that if they just say "credit bureaus steal money from people" it doesn't count, there should be specific negative adjectives to describe it being bad or their negative feelings about it.

### **Risk Perception**

**Emotional feelings of the breach:** The emotional feelings that participants experienced after heard of the breach (e.g., angry, disgusting, indifferent, not surprised), the emotional/attitude change towards Equifax (or other bureaus) after the breach compared to the time before.

**Change of trust:** Mentioning that after this breach, Equifax (or other credit bureaus) will have a less reputable image in the mind of consumers, or the participant personally will have less trust in the company.

**Expectation of credit bureaus:** Expectations towards Equifax, or other companies that have experienced data breaches about what they should do as the countermeasure of the breach, whether they have meet or failed the expectations in the past, as well as their expectations to these companies' future actions.

**The class action lawsuit:** The specific mentioning of the class action lawsuit against Equifax following the breach, whether participants might have heard of it or joined it, how they feel about it.

**Prevalence of data breaches:** The mentioning that there are too many previous data breaches in recent years that the occurrence of the Equifax breach doesn't make the participant too surprised, and that there is too much data available online.

**Mentioning identity theft:** Direct mentioning of identity theft or indirect conceptualization through examples as a consequence of the Equifax breach, or just identity theft in general.

**Victims of the breach:** Talking of targets that are more likely to be affected by the breach, e.g., people who have good credit.

**Likelihood of being personally affected:** The knowledge, assumption or assessment of whether participants themselves are personally affected, and if yes, to what extent, can be either an assured response or a guess.

**Negative consequences of the breach:** Mentioning consequences that's not about identity theft but can still happen after the Equifax data breach, such as invasion of personal privacy when so much personal and financial information was exposed.

**Knowledge of Equifax:** Impression of Equifax as a company, e.g., it's one of the big three credit bureaus, it's the one that got hacked, also include cases where participants say they've never heard of it.

**Cause of the breach:** The description that this breach was conducted by people other than hackers, such as governments, and/or it was profit-driven, e.g., some participants assumed that hackers will sell the stolen data to someone else, others believed that it's an internal breach and someone's disclosing the information intentionally.

**Types of exposed data:** Description of the general impression of some data being exposed in the Equifax data breach (e.g., a lot of personal information released) or specific types of data (e.g., SSN, credit card numbers). Also include cases where participants say they don't know.

**Awareness of the breach:** Memory of whether or not this participant has heard of the breach, what happened in general in the breach.

**Previous data security experience:** Previous experience of data security prob-

lem, such as being involved in a data breach and having password compromised somewhere.

**Previous identity theft experience:** Previous experience of being an identity theft victim, such as someone else applying for credit-related products under the participant's name, the effort in solving the related problems, or the reason for not conducting any kind of follow-up investigations.

### **Protective Actions**

**Check Equifax's website:** The mentioning of someone (either the participant or other related people) check the Equifax website for his or her own breaching status. Also include this code when participants say they didn't check it.

**Credit freeze:** The action of placing a credit freeze, the interpretation of what credit freeze means/what's the expected outcome, the cost of credit freeze, why someone may want to initiate a credit freeze, their assumptions of what a credit freeze may do.

**Check credit report after the breach:** The mentioning of checking credit report following the data breach as a safeguard measure.

**Fraud alert:** The mentioning of placing a fraud alert on file, either for this breach or previous ones, their assumptions of what a fraud alert might do, the process of how to place a fraud alert.

**Credit monitoring service:** Enroll in credit monitoring services provided by credit bureaus, governments, or other entities.

**Self-monitoring:** The action of checking accounts more frequently, keeping an closer eye on them, and the related outcomes.

**Identity theft protection:** Conceptualization of what this type of service does, why someone may want it.

**General security practices:** Strategies to protect one's credit data/online pri-



vacy in general, e.g., don't disclose personal information such as SSN and passwords to others, avoiding suspicious emails, not using PayPal.

**Self-initiated actions after the breach:** Things that the participant has done in reaction to the breach or knows that they could have done, also include cases where they say they don't know.

**Reasons for taking actions:** Any reasons why the participant chose to take any one of the suggested actions above.

**Reasons for not taking actions:** Any reasons why the participant chose to not taking any one of the suggested actions above.

**Triggering new actions:** Any places where participants say they will or might consider doing some actions after the interview, the conversation inspires them to do something, and the reasons behind.

**Suggestion from participants:** The suggestion or proposal made by participants throughout the interview, e.g., credit bureaus shouldn't charge money for their certain offerings such as credit freeze, and there should be a consistent way to calculate credit scores.

**Sources of recommendation:** Protective actions recommended by anyone who's considered as reputable, trustworthy or expert by the participant, e.g., family member, financial advisor. Also include cases where participants said they provided recommendations for other people and hence became the source of knowledge.

**Usability issues:** Reporting about problems and hurdles participants encountered (or other people they know) when trying to initiate any one of the suggested actions.

**Compensations after data breaches:** Description of products and services offered by companies following previous data breaches that the participant or someone he/she knows was involved in (e.g., some companies may offer free or paid credit monitoring services and fraud alerts for victims).

## APPENDIX B

# Supplemental Materials: “Consumer Reactions to Data Breaches that Affected Them”

### B.1 Survey Materials

#### Email address-related questions

We are going to ask you to enter your most commonly used email address at the bottom of this page. We will use your email address to look up whether your email address has been disclosed in any data breaches (also called “security breaches”), using the public lookup service for data breaches [haveibeenpwned.com](https://haveibeenpwned.com). If your email address was involved in any data breaches, we will ask you some questions about those breaches.

**Privacy Notice:** We do not track or store your email address as part of this study, and we will not be able tie your email address to any results or analysis. All records of your email address will reside only in temporary storage to facilitate the lookup of data breaches your email address was involved in and will be deleted following the completion of this task. The researchers will never see your email address.

To access information about breaches, your email address will be communicated to haveibeenpwned.com, a public service not operated by us, which maintains a database of data breaches involving email addresses. Communication with haveibeenpwned.com will occur on secure and encrypted channels, and haveibeenpwned.com also does not permanently store email addresses used in queries. As described in their privacy policy: “Searching for an email address only ever retrieves the address from storage then returns it in the response, the searched address is never explicitly stored anywhere.”

If you have any further concerns about providing your email address, you may opt-out of the survey at this time. We will remove any record of your participation. Note that if you choose to opt out, you will not be compensated.

1. Please enter your most commonly used email address. After the task, you may search for another email address, but for now, we are primarily interested in breaches that may have involved your most commonly used email address. [free text]
2. Thank you for providing your email address. Please tell us more about this email address. Whose email address is it?  It is my own account / I have sole ownership of this account  It is my shared account / I share the account with someone else (e.g., a partner or family member)  It is someone else’s account / someone else has sole ownership of this account  I made up an email address just for this study
3. How often do you check emails in this account?  Every day  A few times a week  A few times a month  A few times a year
4. What do you use this email account for? Choose all that apply.  For professional correspondence (e.g., with colleagues, business partners)  For personal correspondence (e.g., friends and family members)  Account creation / signup for sensitive accounts (e.g., banking, taxes, etc.)  Account creation / signup

- of medium sensitive accounts (e.g., social media, online shopping)    ◦ Account creation / signup for low value accounts (I used it when I'm prompted to sign up but don't really care)    ◦ Other [free-text]
5. Approximately for how long have you been using this email account? [number entry] ◦ year(s)    ◦ month(s)    ◦ week(s)    ◦ day(s)
6. How many other email addresses/accounts do you regularly use? (Not counting the one you entered) [number entry]

### **Breach-related questions**

*(if email not involved in a data breach)* **Your email address has not been part of any of the data breaches recorded by haveibeenpwned.com.** That is great news for you, but we still would like to ask you some further questions.

7. In your opinion, what might be reasons that your email address has not been part of any data breach? [free text]
8. Do you believe another email address that you regularly use is more likely to have breaches? [yes/no]
9. Would you like to take this survey with that email address instead? [yes/no] *(if yes return participant to questions in Section B.1, if no continue to demographic questions in Section B.1)*

*(if email involved in a data breach)* **Your email address was part of a data breach:** According to haveibeenpwned.com your email address was part of one or more data breaches.

10. In your opinion, what might be reasons that your email address has been part of data breaches? [free text]

We will now ask you questions about three of these breaches. We will show you the full data breach history for your email address at the end of the survey.

*(for up to three data breach, the following ...)*

**Your email address was part of the following breach**

[Image and description of breach (see Figure 4.1)]

Please make sure you read the description of this breach, since we will now ask you a few questions with respect to this breach (the description of the breach will be available to you while answering the questions).

11. In your opinion, what might be reasons that your email address has been part of data breaches? [free text]
12. Prior to this study, were you aware that you are affected by this breach?  yes  
 no  unsure
13. *(if yes aware)* How did you first become aware that you are affected by this breach?  I was notified by the breached company.  I was notified by my bank or credit card company.  I was notified by a third-party breach notification service (e.g., Have I Been Pwned, Firefox Monitor, Breach Clarity).  
 I was notified by my credit monitoring or identity theft monitoring service (e.g., LifeLock, Credit Karma).  Someone else (e.g., a romantic partner or a family member) told me about it.  I found out myself through negative events in real life (e.g., suspicious activity on my credit card, locked out of online accounts.)  I learned about the breach through news media.  I do not remember.  Other [free text]
14. *(if yes aware)* Please describe how you felt when you learned that your information was part of this breach  
*(if no/unsure aware)* Please describe how you feel after now learning that your information was part of this breach. [free text]
15. *(if yes aware)* How concerned were you when you learned that your information

was part of this breach?

*(if no/unsure aware)* How concerned are you after now learning that your information was part of this breach?  Not at all concerned  Slightly concerned  Somewhat concerned  Very concerned  Extremely concerned

16. *(if yes aware)* Please describe how you think this breach has or will impact your life. If you suspect or have experienced impacts resulting from this breach, please describe them.

*(if no/unsure aware)* Please describe how you think this breach will impact your life. If you suspect or have experienced impacts resulting from this breach, please describe them as well. [free text]

17. How concerned are you about the following data being compromised in this breach? [for each data type in the breach as provided by HIBP]  Not at all concerned  Slightly concerned  Somewhat concerned  Very concerned  Extremely concerned  I don't know  Does not apply to me (the company does not have my real information)

18. What did you do, if anything, after learning that your information was part of this breach? Please explain why. [free text]

19. Regarding this specific breach, please select how likely you are to initiate each the of the following actions within the next 30 days, or whether you have taken the action already.  Not likely  Somewhat likely  Very likely  I did/do this already  This does not apply to me / I don't understand

*(For each of the following actions:)*

- Change the password of my account for the breached company, if it exists
- Change the password of other accounts that used the same password
- Delete or deactivate my account for the breached company, if it exists
- Enable two-factor authentication on my account for the breached company, if it is available

- Use a credit or identity monitoring service (e.g., LifeLock, Identity Guard, IdentityForce, Credit Karma, Credit Sesame)
  - Use a breach notification service (e.g., Firefox Monitor, Breach Clarity, Have I Been Pwned)
  - Take legal action against the breached company
  - Review my credit reports and/or, bank/credit card statements for suspicious activity
  - File a complaint against the breached company with a consumer protection agency (e.g., FTC, CFPB, State Attorney General)
  - Place a credit freeze on my credit reports
20. Are there any other actions you would like to initiate within the next 30 days or other actions you have already taken? [free text]

### **Demographics & attention checks**

21. Which of the following breaches were you asked about in this study? [multiple choice of the correct answer and four decoys]
22. What is your age?  18-24  25-29  30-34  35-39  40-44  45-49  
 50-54  54-59  60-64  65+  Prefer not to say
23. What is your gender?  Man  Woman  Non-Binary  Prefer not to answer  Other [free text]
24. What is the highest level of education you have completed?  Less than high school  High school or equivalent  Some college, no degree  Associate's degree, occupational  Associate's degree, academic  Bachelor's Degree  Master's Degree  Professional degree  Doctoral degree  Prefer not to say

25. What is the shape of a red ball?  Red  Blue  Square  Round   
Prefer not to answer
26. Which of the following best describes your educational background or job field?  
 I have an education in, or work in, the field of computer science, computer engineering, or IT.  I do not have an education in, or work in, the field of computer science, computer engineering, or IT.  Prefer not to answer
27. Which of the following best describes your educational background or job field?  
 I have an education in or work-in/practice law or other legal services.   
I do not have an education in or work-in/practice law or other legal services.  
 Prefer not to answer
28. What was your total household income before taxes during the past 12 months?  
 Under \$15,000  \$15,000 to \$24,999  \$25,000 to \$34,999  \$35,000 to \$49,999  \$50,000 to \$74,999  \$75,000 to \$99,999  \$100,000 to \$149,999  
 \$150,000 or above  Prefer not to say

## Debrief

**Information on breaches your email address was part of:** Thank you for completing our study. Please note that the information about data breaches we showed to you is real. Your email address, and potentially other personal information has been part of these breaches and could be used by criminals to steal your identity or access your accounts.

**List of breaches your email address was part of:** Below is the full list of breaches in which the email address you entered was involved according to haveibeenpwned.com. Please note that you can always obtain the same results by checking your email address on haveibeenpwned.com, which, in addition, also provides records with sensitive breaches upon the verification of your email account. Please keep in mind



that this list only reflects breaches that are registered in the haveibeenpwned.com database, your information may have been exposed in other breaches.

**Resources for breach recovery and further reading** Here is a list of resources to help you prevent or recover from harm due your information being exposed in data breaches, as well as help you better protect yourself from data breaches in the future.

- *Resources about recovering from a data breach:*
  - Federal Trade Commission: Identity theft recovery steps
  - Federal Trade Commission: Credit Freeze FAQs
  - Firefox Monitor: What to do after a data breach
  - Norton: What to do after 5 types of data breaches
- *Resources about protecting yourself against future breaches:*
  - Firefox Monitor: How to create strong passwords
  - Firefox Monitor: Steps to protect your online identity

## B.2 Qualitative Analysis Codebook

In the following we provide our unified codebook with the primary codes, their respective counts, and their first-level sub-codes.

- **bad actors (17)**: company sell data, hackers, department stores
- **behaviour (94)**: continue use as before, insecure, keep using email, secure practice, email practice, insecure practice
- **cannot recall (17)**: confused, unconcerned, surprised, concerned
- **consequence experienced (97)**: compromised accounts, information disclosure, spam, data on the dark web, scam, attempted login, other account with same pwd, email disclosure, identity theft, social media account hacked, physical, financial disadvantage, unrecognized new account, past event, reputation,

job offer missed, upset, site breached

- **consequence potentially (92)**: spam, identity theft, compromised accounts, information misuse, financial disadvantage, scam, physical, financial account hacked, information disclosure, stalking, other account with same password, unrecognized new account
- **data not relevant (84)**: outdated, fake data, not sensitive, unique password, not primary email, little data, will be caught by spam filter, so much data out there, account not used, unimportant password, unique username
- **data relevant (3)**: sensitive
- **defense intended to be put into place as reaction to breach (180)**: change password, monitor email, use secure passwords, monitor suspicious activity, monitor financial information, do not use facebook login for shopping sites, increase protective measures, change email, be more cautious, 2FA enabled, limit online disclosure, review accounts, stop using, reduced use email, check suspicious emails, signing up to websites less often, new email account, learn more about breach, reduced use site, close account, scan computer frequently, re-link security accounts, change financial information, change employer, monitor accounts, use vpn, review financial information, unique password, change username, use password manager, go after companies, learn about safeguarding, solve issues as they appear, security checkup, check financial information, protective measures, stop using email, protect email, stop using service, tor, investigate, strong password, location setting, no reuse password, be more careful, legal action
- **defense put into place as reaction to breach (226)**: use password manager, change password, reduced use site, change emails, protective measures, 2FA enabled, change password creation strategy, unique password, no cc info in unused apps, actions caused by other breach, close account, change username,

remove email from accounts, use secure passwords, review financial information, use breach monitors, be more cautious, review account information, update browser, check suspicious emails, change email, stop using site, nothing, changed info, check account, 2fa enabled, limit data disclosure, unsubscribed from mailer, change info, reviewed prior steps, monitoring, check financial, email practices, contacted company, changed email, unsubscribed, changed password, delete account, learn about breach, antivirus, called credit card company, recover hacked account, careful disclosure, no reuse password, strong password

- **defense put into place pro-actively before breach (40)**: use secure passwords, 2FA enabled, be cautious, change password, don't answer phone calls, review financial information, unique password, use password manager, monitor accounts, monitor emails, unique email, protective measures, monitor credit reports, spam filter, stop using site, change email, account not used, monitor financial information
- **do not know hibpwnd (2)**
- **feeling (929)**: unconcerned, concerned, violated, annoyed, negative, skeptical, uncomfortable, fatigued, paranoid, cautious, hopeful, upset, scared, unsurprised, would have been contacted, overwhelmed, disappointed, unsure, reassured, don't care, curious, not worried, relief, insecure, no fear, worried, unhappy, not important enough, confused, indifferent, surprised, unsafe, ashamed, regret, informed, used to breaches, no blame on company, upset
- **first breach (1)**
- **immediately informed (1)**
- **impact (525)** impact little, impact none, impact large, impact positive, impact unsure, impact negative, unconcerned
- **needs more info (1)**
- **not hacked into a lot (1)**

- **third party (11):** *bad security, good security at company*
- **unclear (2)**

## APPENDIX C

# Supplemental Materials: “An Empirical Analysis of Data Breach Notifications”

### C.1 Analyzed Breaches Notifications

Case Title	Case No.	Date Received	No. of MD Residents
intuit Inc.	itu-297341	4/26/2018	9
intuit Inc.	itu-295061	2/22/2018	149
SunTrust Bank	itu-298137	5/8/2018	1
Engle Martin & Associates	itu-295041	2/14/2018	5
COUNTRY Mutual Insurance	itu-294968	2/9/2018	4
OneMain Financial Group, LLC	itu-294973	2/6/2018	12
CNU Online Holdings, LLC	itu-295185 (2)	3/7/2018	24
Lincoln Financial Group	itu-295367 (2)	3/29/2018	316
Medical science & Computing, LLC	itu-295048	2/13/2018	104
TeenSafe	itu-298253	5/31/2018	4
Kinetics Systems, Inc.	itu-294972 (2)	2/6/2018	106
Marriott International Inc.	itu-295059 (2)	2/22/2018	3

Continued on next page

<b>Case Title</b>	<b>Case No.</b>	<b>Date Received</b>	<b>No. of MD Residents</b>
Under Armour, Inc.	itu-295369	3/29/2018	NA
Delta Air Lines, Inc.	itu-295677	4/12/2018	8,246
Farmgirl Flowers, Inc.	itu-298121	5/11/2018	24
Sears Holdings Corporation	itu-297330	4/24/2018	2,588
Thomas Edison State University	itu-295054	2/12/2018	12
Palo Alto Unified School District	itu-295046	2/14/2018	2
Nampa School District	itu-295387	3/16/2018	1
LendKey Technologies, Inc.	itu-295039	2/15/2018	256
Pershing, LLC	itu-296815	4/19/2018	10
UnitedHealthcare	itu-298215 (2)	5/22/2018	3
Capital Digestive Care	itu-297325	4/23/2018	8,713
Cambridge Dental Consulting Group	itu-298167	5/7/2018	4
The Childrens Mercy Hospital	itu-297380	4/30/2018	6
Coastal Cape Fear Eye Associates, P.A	itu-295071 (2)	2/23/2018	1
Eastern Shore Rural Health Inc.	itu-295051	2/13/2018	8
Flexible Benefit Service Corporation	itu-295047	2/14/2018	21
FastHealth Corporation	itu-295014	2/27/2018	NA
ATI Holdings, LLC	itu-295354	3/28/2018	1,823
Inogen, Inc.	itu-295778 (2)	4/13/2018	564
The Oregon Clinic	itu-298227 (2)	3/9/2018	2
Aflac	itu-294982 (2)	2/6/2018	1
W.W. Grainger, Inc.	itu-295831 (2)	4/17/2018	343
Heritage Land Bank, ACA	itu-298272	5/29/2018	NA
Cenlar FSB	itu-295287 (2)	3/20/2018	1
Draper and Kramer, Inc.	itu-295202	3/8/2018	69
Mercy Health Love County Hospital and Clinic	itu-295077	2/27/2018	3
Betterton, Tyler & Summonte, PL	itu-295520	4/10/2018	1
CPT Group Inc.	itu-297379 (2)	4/27/2018	67
Capital One, National Association	itu-298210	5/18/2018	1

Continued on next page

<b>Case Title</b>	<b>Case No.</b>	<b>Date Received</b>	<b>No. of MD Residents</b>
Jarrett & Luitjens, PLC	itu-298076	5/25/2018	7
Duquesne University	itu-294966	1/31/2018	6
The Arc Northern Chesapeak Region	itu-298199	5/11/2018	1,406
SA Stone Wealth Management Inc.	itu-297381 (2)	4/30/2018	2
Bennett Thrasher LLP	itu-298130	5/4/2018	4
Luxury Retreats	itu-298077	5/25/2018	9
Advanced Graphic Products, Inc.	itu-294999 (2)	2/1/2018	12
Ventiv Technology, Inc.	itu-294998	2/1/2018	3
Ullico Inc.	itu-298068	5/29/2018	8
Sallie Mae	itu-295218	3/13/2018	NA
Ameriprise Financial, Inc.	itu-295389	4/5/2018	1
M&E Tax and Financial Services	itu-295172	3/6/2018	300
Alabama Ballet	itu-295372	3/30/2018	2
UniCarriers Americas Corporation	itu-297398	4/23/2018	1
Rockville Eye Surgery Center, LLC	itu-295057	2/20/2018	2
Mattress Firm, Inc.	itu-294949	1/3/2018	1
Personal Care Products Council	itu-295201 (2)	3/7/2018	33
Waccamaw Management, LLC	itu-295169	3/5/2018	1
Northwestern Mutual Life Insurance Company	itu-296813 (2)	4/19/2018	2
Temp-Tations Home LLC	itu-295824	4/16/2018	114
Family Investment Administration at the Department of Human Services	itu-297340	4/26/2018	2
Allied Contractors, Inc.	itu-295037	2/16/2018	NA
K. Hovnanian American Mortgage, LLC	itu-295017	4/2/2018	10
HSBC Mortgage Services, Inc.	itu-294955 (2)	1/18/2018	1
ACTIVE Network	itu-295073 (2)	2/23/2018	6
Citizens Financial Group, Inc.	itu-295171	3/6/2018	1
MidCap Financial Services, LLC	itu-295096	2/5/2018	2
Cetera Advisors LLC	itu-297370	4/26/2018	60

Continued on next page

<b>Case Title</b>	<b>Case No.</b>	<b>Date Received</b>	<b>No. of MD Residents</b>
Rail Europe North America Inc.	itu-297372	4/27/2018	361
The Retirement Advantage, Inc.	itu-295004	3/28/2018	59
Purdy Insurance Agency	itu-298146	5/7/2018	8
Austin, Nichols & Co., Inc.	itu-295519 (2)	4/10/2018	1
Pendleton Square Trust Company	itu-298066	5/30/2018	1
NWAM, LLC	itu-295001	3/27/2018	2
talentReef, Inc.	itu-294943	1/10/2018	185
ABC Phones of North Carolina, Inc.	itu-295811 (2)	4/13/2018	11
Member First Mortgage, LLC	itu-295030	1/29/2018	10
Ellevest, Inc.	itu-298256	5/21/2018	1
JJ Haines	itu-295056	2/21/2018	NA
abod & caruso, llc	itu-298231	5/25/2018	1,080
Hampton Roads Shipping Association	itu-298061	5/2/2018	4
Mise En Place Restaurant Services, Inc.	itu-295804 (2)	4/11/2018	102
B.T.C.E., Inc. d/b/a HomeBrewIt.com	itu-295781	4/11/2018	11
Orbitz Worldwide, Inc.	itu-295291 (2)	3/21/2018	3,607
Hurst & Langlinais, Ltd.	itu-295292	3/21/2018	2
Teal Becker & Chiaramonte, CPAs, PC	itu-295002	3/27/2018	29
Santa Cruz Biotechnology, Inc.	itu-295064 (2)	2/22/2018	5
Best Buy Co., Inc.	itu-295777	4/13/2018	NA
Musicians On Call, Inc.	itu-295062	2/22/2018	33
Conservest Capital Advisors, Inc.	itu-298161	5/2/2018	3
Alaska Airlines	itu-298205 (3)	5/3/2018	10
Southwest Airlines Co.	itu-297324	4/20/2018	1,427
	(update)		
Inspire Home Loans Inc.	itu-294967	2/9/2018	10
Wealth Management, Inc.	itu-295355	3/28/2018	NA
Branton, de Jong and Associates	itu-295295 (2)	3/26/2018	4
Aberdeen Proving Ground	itu-295029	1/26/2018	36
Federal Credit Union			

Continued on next page



<b>Case Title</b>	<b>Case No.</b>	<b>Date Received</b>	<b>No. of MD Residents</b>
RBC Royal Bank	itu-294965 (2)	1/26/2018	NA
The Hartford	itu-297339 (2)	3/13/2018	16
Malley's Chocolates	itu-298150	5/7/2018	27
Varanko & Black CPA	itu-298163	5/4/2018	212
Boys & Girls Clubs	itu-295692	4/3/2018	3
Julian Sur, CPA	itu-297364	4/26/2018	4
University of Wisconsin- Superior Alumni Association	itu-295060	2/22/2018	NA
World Travel Holdings, Inc.	itu-295814	4/16/2018	94
Rail Europe SAS	itu-298252	5/30/2018	87
Bronson Nutritionals LLC	itu-295284	3/19/2018	375
Williams-Sonoma, Inc.	itu-295187 (2)	3/7/2018	1
Gibbs & Cox	itu-295843	4/9/2018	1
Franklin American Mortgage Company	itu-295340	3/28/2018	2
Clinical Pathology Laboratories Southeast, Inc.	itu-295290 (2)	3/21/2018	18
J.C. Penney Corporation, Inc.	itu-295339	3/27/2018	11
Thermo Fisher Scientific Inc.	itu-295368	3/29/2018	117
Bigfoot Gun Belts	itu-295684	4/4/2018	45
Vidant Medical Group LLC	itu-295000	3/27/2018	1
John Y. Trent & Associates LLC	itu-295278	4/4/2018	5
Employer Leasing Company	itu-294953	1/18/2018	3
Bell Partners Inc.	itu-298064	5/29/2018	7
Knape & Vogt Manufacturing Company	itu-298148	5/11/2018	2
TrueNet Communications, Inc.	itu-295094	2/8/2018	161
Global University	itu-298147	5/11/2018	1,320
Goldleaf Partners Services, Inc.	itu-294964	1/26/2018	11
Temecula Motorsports, Inc.	itu-298315 (2)	5/31/2018	11
Rea.deeming Beauty, Inc.	itu-294938	1/12/2018	364
Chopra Enterprises, LLC	itu-295184 (2)	3/6/2018	14

Continued on next page

<b>Case Title</b>	<b>Case No.</b>	<b>Date Received</b>	<b>No. of MD Residents</b>
Atrium Hospitality	itu-297329	4/24/2018	1
Whitmer & Company CPA's, LLP	itu-295697	4/10/2018	2
Guaranteed Rate, Inc.	itu-294937	1/12/2018	1,406
Broward College	itu-294944	1/10/2018	145
1st Mariner Bank	itu-294989	2/5/2018	929
Interstate Plastics, Inc.	itu-295177	3/6/2018	16
Manduka	itu-295333 (2)	3/27/2018	885
1st Mariner Bank	itu-295050 (2)	2/12/2018	558
North 40 Outfitters	itu-295065	2/23/2018	105
Corporation Service Company	itu-298062	5/17/2018	7,585
Strategic Analysis, Inc.	itu-296816	4/19/2018	356
The Meridian Group	itu-295063	2/22/2018	13
Lydia Security Monitoring	itu-295055	2/28/2018	186
ABC Bus Companies, Inc.	itu-295138	3/1/2018	NA
American National Insurance Company	itu-298172 (2)	5/7/2018	47
American Society of Anesthesiologists	itu-297322 (2)	4/20/2018	NA
Anchor Fund, LLC	itu-295206	3/9/2018	1
Bearing Distributors, Inc.	itu-295518	4/6/2018	1
Capital Farm Credit	itu-297337	4/25/2018	NA
Equias Alliance, LLC	itu-298152	5/9/2018	5
Invacare Corporation	itu-295269	3/14/2018	2
Kirby & Associates PC	itu-298200	5/11/2018	15
MCR Inestors LLC	itu-297327	4/23/2018	667
Mise En Place Restaurant Services, Inc.	itu-295804 (1)	4/11/2018	102
Parkway Corporation	itu-295144	3/2/2018	6
Premier Fixtures, LLC	itu-296814	4/19/2018	1
PrintingCenterUSA.com	itu-298254	5/24/2018	55
Ramy Brook	itu-297328	4/24/2018	13
Shutterfly, Inc.	itu-295693	4/2/2018	1
Sprint Corporation	itu-295363 (2)	3/28/2018	1

Continued on next page

<b>Case Title</b>	<b>Case No.</b>	<b>Date Received</b>	<b>No. of MD Residents</b>
Sprint Corporation	itu-298171 (2)	5/7/2018	1
Squire Patton Boggs (US) LLP	itu-295296	3/27/2018	9
St. Mary's Health, Inc.	itu-298247	5/29/2018	21
The Vanguard Group, Inc.	itu-295058	2/21/2018	1
US GreenFiber LLC	itu-295127	4/3/2018	4
WithumSmith+Brown, PC	itu-297336	4/25/2018	2

## C.2 Qualitative Analysis Codebook

<b>Category</b>	<b>Code</b>	<b>Values</b>
Other	Delivery methods	mailed letter, email, push notification, unclear
Formatting	Headings	yes (in separate lines), yes (in same line as the main text), yes (as tables), no
Formatting	Appendix	0, 1, 2, 3
Formatting	Format of presented actions	in numbered list, in bullet point list, in plain text, other
Formatting	Specifying the enrollment deadline	with text markups, in plain text, other
Formatting	Specifying the duration of benefits	with text markups, in plain text, other
Formatting	Including the recipient's breached information	with text markups, in plain text, other
Formatting	Explicit mentioning of not breached information	with text markups, in plain text, other
Formatting	Describing ways to keep consumers updated	with text markups, in plain text, other
Risk Comm.	Cause of the breach	specified, under investigation, other
Risk Comm.	When the breach occurred	specified, under investigation, other
Risk Comm.	When the breach was discovered	specified, under investigation, other
Risk Comm.	Types of breached info	specified, other

Continued on next page

Category	Code	Values
Risk Comm.	Info was breached	absolutely, maybe, no, no evidence, other
Risk Comm.	Breached info was misused	absolutely, maybe, no, no evidence, other
Risk Comm.	Consequences of breached info	specified, other
Actions	Location of described actions	compensations in main text, no other actions in main text; compensations and other actions in main text; compensations in appendix, no other actions in appendix; compensations and other actions in appendix; no compensations, other actions in appendix
Actions	Identity theft protection service	in main text, in appendix, other
Actions	Credit monitoring service	in main text, in appendix, other
Actions	Credit freeze	in main text, in appendix, other
Actions	Fraud alert	in main text, in appendix, other
Actions	Monitor accounts and credit reports	in main text, in appendix, other
Actions	Use two-factor authentication	in main text, in appendix, other
Actions	Change passwords	in main text, in appendix, other
Actions	Keep physical materials safe	in main text, in appendix, other
Actions	Other actions or compensations	in main text, in appendix, other
Actions	Compensation enrollment instructions	in main text, in appendix, other
Actions	Types of enrollment	automatic, link only, link and activation code, vague instructions, other

## APPENDIX D

# Supplemental Materials: “Adoption and Abandonment of Protection Practices”

### D.1 Survey Instruments

#### Introduction

Before you start, we would greatly appreciate it if you could switch off phone, email, music, or other distractions so that you can focus on this study. Thank you!

First, please enter your Prolific ID here. It should have 24 alphanumeric characters. Please make sure you complete this step so that we can verify your response and give you the compensation: \_\_\_\_\_

We are interested in critiquing and improving existing advice regarding how to stay safe online. We encourage you to be as HONEST as possible in your responses. Your individual responses will be anonymized.

Please proceed to the next page. Read each item carefully before answering.

## Show Practices

*[For the following 30 items, we selected ten of them (four security, four privacy, two identity theft protection) at random and displayed them to the participant.]*

*[For each item, the participant was asked to choose from the following answer options: ○ I am always doing this. ○ I am doing this but there are exceptions. Please describe it further: \_\_\_\_\_ ○ I am not doing this anymore, but I have done this before. Please describe it further: \_\_\_\_\_ ○ I have never done this before, but I expect to do this in the near future. ○ I have never done this before, and I do not expect to do this in the near future. ○ I have never heard of this/I do not understand. ○ Other (please specify): \_\_\_\_\_ ]*

- *[Update]* Have you ever installed operating system and software updates immediately after they become available?
- *[Automatic-update]* Have you ever kept automatic software updates turned on?
- *[Unique-password]* Have you ever used different passwords for each account?
- *[Strong-password]* A strong password should meet all of the following criteria:
  1. long – it should exceed the minimum password length (e.g., the US federal government requires employees to use passwords including at least 12 characters);
  2. complex – it should be a combination of letters, digits, and special characters;
  3. hard to guess – it should not contain common phrases (e.g., “iloveyou”) and personal information (e.g., your birth date or names of family members).

Have you ever used strong passwords for your online accounts?

- *[2FA]* Two-factor authentication (2FA) requires you to send another piece of information to verify your identity on top of username and password, such as a temporary code sent to your phone. Other forms of 2FA include USB ports like YubiKeys, and security apps like Duo. We want you to think about the time when you proactively enable 2FA. Examples of services that offer opt-in 2FA include shopping websites like Amazon, and email service providers such as Google and Yahoo.

Have you ever opted in to 2FA for online accounts that offer this option, outside of contexts where it is mandated (such as by your employers or banks)?

- *[Antivirus]* Anti-virus software is primarily designed to prevent, detect and remove software viruses and other malicious software such as worms, trojans, and adware.

Have you ever used anti-virus software?

- *[Password-manager]* A password manager is a piece of software that helps you generate strong passwords, stores your login information for all websites and apps you use, and helps you log into them automatically. It encrypts your password database with a master password – the master password is the only one you have to remember.

Have you ever used a password manager?

- *[HTTPS]* Hypertext Transfer Protocol Secure (HTTPS) extends the Hypertext Transfer Protocol (HTTP) to secure communications over the network. To check if a website uses HTTPS, you can look for a lock icon in the address bar. Below is an example of the HTTPS icon in Google Chrome:

*[Display a screenshot.]*

Have you ever checked if the website you are visiting uses HTTPS?

- *[Check-URL]* URL is the address of a website. To find the URL of the website you are visiting, you can look at the address bar of your web browser. The address bar is the long white bar at the top of your web browser, below the tabs at the top. Below is an example of a URL in the address bar in Google Chrome:

*[Display a screenshot.]*

Have you ever checked the URL to verify you are visiting the website you intended to?

- *[Install-software]* Have you ever only installed software coming from trusted and reputable sources on your computers or mobile devices?
- *[Link-clicking]* When checking your emails, have you ever avoided clicking links sent by people or companies you do not know?
- *[Attachment-clicking]* When checking your emails, have you ever avoided opening email attachments from people or companies you do not know?
- *[Temporary-credentials]* When it does not hamper the usability of the service, have you ever used fake identities for your online activities to prevent your real information from being tracked? E.g., use a fake name, phone number, or email address when creating a new account.
- *[Hide-info]* When it does not hamper the usability of the service, have you ever refused to provide information about yourself that is not essential to complete the transaction with a business? E.g., when filling out a final confirmation form for online shopping, you skip all optional fields.
- *[Real-name]* Have you ever avoided using a website because they ask for your real name?



- *[Extension]* There are browser extensions that can help protect your privacy in different ways. Some browser extensions remove advertisements from the webpage (e.g., pop-ups, banner ads, and other common forms of online advertisements). Some browser extensions block trackers used by websites to track your activity for site analytics and targeted ads delivery. Some others prevent unauthorized web sites from running JavaScript, Java, Flash, or other scripts.

Have you ever used a browser extension that blocks ads, trackers or scripts?

- *[Cookies-disable]* HTTP cookies (also called browser cookies) are small text files placed on your device after you visit a website. Cookies are designed for websites to record information about you, such as items in the shopping cart, browsing history, and login credentials that you previously entered.

First-party cookies are issued by a website that a user views directly – in most cases, they are essential to keep the website running smoothly. On the other hand, third-party cookies are not created by the website being visited, but rather by someone else – these cookies are used primarily for tracking and targeted advertising purposes.

Have you ever disabled or turned off third-party cookies in your browser?

- *[Cookies-clean]* HTTP cookies (also called browser cookies) are small text files placed on your device after you visit a website. Cookies are designed for websites to record information about you, such as items in the shopping cart, browsing history, and login credentials that you previously entered.

Cookies can be useful - they let sites remember you and what your preferences are, making it quick and easy to log into your favorite websites. On the negative side, cookies take up space on your device, can result in errors sometimes, and can pose a privacy risk when it stores sensitive information about you.

Have you ever cleared your web browser cookies and history?

- *[Anonymity-system]* Anonymity systems are infrastructure that provides anonymity on the Internet, through allowing information to be transmitted from one party to another without leaking their identity. Examples of anonymity systems include Tor, JonDonym, and VPN services.

Note: using private browsing mode (also called “incognito mode”) doesn’t make you anonymous on the Internet. Your Internet service provider, employer, or the sites themselves can still gather information about pages you visit.

Have you ever used anonymity systems, outside of contexts where they are mandated (such as by your employers)?

- *[Encryption]* Have you ever encrypted your phone calls, text messages or emails, outside of contexts where it is mandated (such as by your employers)?
- *[Public-comp]* Have you ever used a public computer to browse anonymously?
- *[Incognito]* Private browsing mode (sometimes called “incognito mode”) is a feature in popular web browsers. It automatically erases your browsing information, such as passwords, cookies and history.

Have you ever used the private browsing mode or incognito mode when browsing the web?

- *[Search-engine]* Have you ever used a search engine that does not keep track of your search history?
- *[Facial-recognition]* Facial recognition is a technology capable of identifying or verifying a person from a digital source such as a picture or video footage. It is widely used nowadays for ID verification, policing, airport screening, among many other purposes.

In some cases, you can choose to opt out of facial recognition. For instance, Facebook’s facial recognition feature, which recognizes you and suggests that

you be tagged when someone uploads a picture of you, can be turned off in the privacy settings. In other cases, opting out is much harder, such as law enforcement's use of facial recognition on public surveillance camera footage to detect criminals.

Have you ever opted out of facial recognition when possible?

- *[Credit-monitoring]* A credit monitoring service tracks activity on your credit reports at one or more major credit bureaus – Equifax, Experian, and TransUnion. Usually, credit monitoring services alert you when a company checks your credit history, when a new loan or credit card account is opened in your name, and so forth.

Have you ever used a credit monitoring service?

- *[Identity-monitoring]* An identity monitoring service tracks potential misuse of your personal information from a variety of databases. In addition to monitoring your credit reports at one or more major credit bureaus – Equifax, Experian, and TransUnion, identity monitoring also alerts you for potential risks that do not show up on your credit reports, such as change of address requests and personal information sold on the dark web.

Have you ever used an identity monitoring service?

- *[Credit-report]* A credit report is a statement that has information about your credit activity and current credit situation, such as loan paying history and the status of your credit accounts. Credit reports are provided by credit reporting companies, also known as credit bureaus or consumer reporting agencies, who collect and store financial data about you from creditors (e.g., lenders, credit card companies, and other financial companies).

Have you ever obtained free copies of your credit reports from **AnnualCreditReport**.

com?

- *[Statements]* Have you ever checked for fraudulent charges on your account statements (e.g., for credit card, bank, retirement, and brokerage)?
- *[Credit-freeze]* With a credit freeze, no one - including you - can access your credit report to open new accounts. You will get a PIN number to use each time you want to freeze and unfreeze your account to apply for new credit. To place a credit freeze, you should contact each of the three credit reporting agencies—Equifax, Experian, and TransUnion—individually at their credit freeze portals.

Have you ever placed a credit freeze on your credit reports at all big three credit reporting agencies?

- *[Fraud-alert]* With a fraud alert, businesses must try to verify your identity before extending new credit. Usually that means calling to check if you are at a particular store attempting to take out new credit. To place a fraud alert, you should contact any one of the three major credit reporting agencies—Equifax, Experian, and TransUnion—either by phone or online. The one you contact is required to notify the other two.

Have you ever placed a fraud alert on your credit reports at one of the three major credit reporting agencies?

- *[Attention-check]* Have you ever ignored any survey questions when taking online surveys?

This is an attention check question. Please select “I have never done this before, and I do not expect to do this in the near future” to let us know that you are paying attention.

## Demographics

1. Prior to this study, has anyone ever gained unauthorized access to one of your online accounts? E.g., someone secretly changed your password without you noticing it.  Yes  No  I am not sure
2. Prior to this study, have you ever been notified that your information was exposed in a data breach, by the breached company or other services?  Yes  No  I am not sure
3. Prior to this study, have you ever been a victim of identity theft? E.g., someone secretly applied to a new credit card under your name.  Yes  No  I am not sure
4. What is your age?  18-24  25-34  35-44  45-54  55-64  65-74  75 or older  Prefer not to answer
5. What is your gender?  Women  Men  Non-binary  Prefer to self-describe: \_\_\_\_\_  Prefer not to answer
6. What was your total household income before taxes during the past 12 months?  <\$20,000  \$20,000 to \$34,999  \$35,000 to \$49,999  \$50,000 to \$74,999  \$75,000 to \$99,999  \$100,000 or above  Prefer not to answer
7. What is the highest degree or level of school that you have completed?  Some high school  High school degree of equivalent (e.g., GED)  Some college, no degree  Trade, technical, or vocational training  Associate's degree  Bachelor's degree  Master's degree  Professional degree (JD, MD, etc.)  Doctoral degree  Other (please specify): \_\_\_\_\_  Prefer not to answer
8. Which describes your current employment status?  Employed full-time  Employed part-time  Self-employed  Care-provider  Homemaker

Retired    Student    Looking for work/unemployed    Other (please specify):  
 \_\_\_\_\_    Prefer not to answer

9. Which describes your educational background or job field?    I have an education in, or work in, the field of computer science, computer engineering or IT.    I do not have an education in, or work in, the field of computer science, computer engineering or IT.    Prefer not to answer

10. Have you ever worked or taken coursework related to security and privacy?    Yes    No    Prefer not to answer

11. *[If answer is “yes” to the question above]* Please describe your experience working or taking coursework related to privacy or security: \_\_\_\_\_

Do you have any feedback about this survey? If yes, please enter it here. Otherwise please skip this question and proceed to the next page. You will then be redirected back to Prolific.

## D.2 Qualitative Analysis Codebook

Code	Meaning	Example Quote
as-needed	Participants take, abandon or make exceptions for the practice simply when they needed to, without specifying any reasons or frequency with which they do so. Vague phrases like “I do this when needed” indicate this code.	“Yes when I needed to” (update)

Continued on next page

<b>Code</b>	<b>Meaning</b>	<b>Example Quote</b>
because-of-risk	Participants say they use (or used) this practice/service only when they are at risk, such as facing the risk of identity theft, when engaging in dangerous behavior other than visiting websites, or simply “when something is wrong.” Also includes cases in which the user does not trust the network, e.g., use VPN only when on public networks, such as when travelling.	“I have had to place a freeze on my credit last year when my information was stolen” (credit-freeze)
frequency	Participants simply specify the frequency of how often they take one practice/service, but do not explain why, e.g., I sometimes / often / rarely do this.	“I sometimes give a fake name” (real-name)
only-blocking	Participants adopt the practice to evade content blocking, censorship, geographic restrictions, access limitations (such as news site paywalls).	“I used vpn to bypass my school’s blocked sites” (anonymity-system)
only-email	The practice or service is only used when participants need to deal with email in general, such as when creating an account for an email address, when accessing an email account, when clicking on email links.	“Now I use a Gmail account with fake info” (temporary-credentials)
only-entertainment	Information related to the participant’s entertainment activities constitutes an exceptional case, such as accounts for social media and gaming. It may be the only context in which the practice is adopted, or not adopted.	“Sure, for sites like Reddit, where that is more the culture” (temporary-credential)
only-finance	The practice/service is only used when participants need to deal with financial information, such as banking sites and when shopping.	“I’ve encrypted financial documents I send to my accountant” (encryption)

Continued on next page

<b>Code</b>	<b>Meaning</b>	<b>Example Quote</b>
only-required	Participants only adopt the practice when it is mandated or expected behavior (e.g., by employer or institution), or when it is otherwise forced upon them, such as by a website or service that requires 2FA, or when using a device or service on which another person has made the decision to adopt the practice.	“Use what my work provides” (antivirus)
only-sensitive	The practice or service is only used when the site requires sensitive or private information, or the site / software is potentially incriminating e.g., porn sites, dark web, etc. However it is not specified that the sensitive info is necessarily financial info (which takes precedence) or other data types.	“I have used a public library’s computer and hotel lobby computers to look up sensitive information” (public-comp)
only-suspicious	The practice or service is only used when participants feel the site is shady / suspicious / odd.	“I only check on it if the website feels suspicious.” (check-URL)
platform-specific	Practice is or is not adopted depending on the platform (e.g., OS, browser, text messaging but not social media, etc.). For instance, a practice might not be adapted on platforms that are perceived to be secure, and on which the practice has low value or is irrelevant. Alternatively they may adopt the practice when it is offered by the platform.	“I currently have scripts blocked in Firefox, but not Chrome or IE.” (extension)

Continued on next page



<b>Code</b>	<b>Meaning</b>	<b>Example Quote</b>
site-specific	The practice is only used for specific sites, apps, or software, without any indicator of what kind of sites or apps they are (e.g., sensitive, suspicious, finance, entertainment). Use practice-unavailable for situations in which the practice is adopted whenever it is made available by the site. If they adopt the practice only when the site explicitly presents them with the option to adopt it, use only-when-offered. This code applies to not only sites, but also other situations.	“Some of the time I use anonymity systems, it just depends on the sites” (anonymity-system)
unrelated-reason	The practice is not used for the intended purpose, but rather, when used, it is used for a reason unrelated to security, privacy, or identity protection, such as to improve the performance of the computer or browser, to fix problems in the browser, convenience, etc.	“When I am having problems with my browser” (cookies-disable)
unwilling	Participants say they cannot follow the practice because they are simply “lazy” or “busy.” If they provide details that make it sound like the practice too impractical for them, use the code “impractical” instead.	“Sometimes I’m too busy to take the time to update right at that moment” (update)
when-offered-free	Participants will not proactively seek for this practice, but use it when it is offered to them, such as by being provided for free (e.g., like a free trial for antivirus). Similarly, they adopt the practice when it is conveniently offered as a simple choice, such as a pop-up option to disable facial recognition on social media, as compared to having to look through preferences to find the setting.	“Use what my work provides” (antivirus)

Continued on next page

<b>Code</b>	<b>Meaning</b>	<b>Example Quote</b>
impractical	The practice is hard, too inconvenient, or too unusable to be adopted or implemented in real life. This also includes case of failure, where participants want to follow the practice but fail to do so in real life.	“I usually use totally different ones for each account, but sometimes I just change a letter or a character in the same password” (unique-password)
not-needed	Participants simply say they do not need or want this practice anymore, perhaps indicating that it is of low value, without indicating specific reasons covered by other codes.	“I don’t currently need credit monitoring” (credit-monitoring)
performance-issues	Practice slows device or some component of the device down significantly, or causes other performance issues, such as disk space or memory hogging.	“Frequent, random updates are annoying and greatly slow down Internet speeds and are disruptive” (update)
practice-unavailable	Participants say they cannot always follow the practice because it is unavailable or non-functional in some (or in all) cases, or that they used the practice they used until it was discontinued or stopped working (in which case, if the company had continued to offer the product or service, they probably would have stayed enrolled).	“They discontinued the service” (identity-monitoring)
usage-interference	Adopting this practice will cause disruption of normal usage when it comes to browsing the internet or completing tasks on devices. These disruptions are more than mere inconveniences, they involve breaking existing functionality or impeding regular use of the device.	“I used it with Mozilla Firefox. For a while Java scripts and Flash were causing Firefox to crash. It seems to be fine now but the browser still asks me if I want to run them” (extension)

Continued on next page

<b>Code</b>	<b>Meaning</b>	<b>Example Quote</b>
using-substitute	The user has chosen to adopt a substitute practice that they feel can achieve the same purpose. The substitute need not necessarily be a sensible one, but the participant should at least implicitly think that it is.	“I use third party browser extensions such as privacy badger and ublock origin. So whichever cookies those allow by default are the ones not turned off” (cookies-disable)
for-sharing	Device, account, or password sharing is the primary factor in the participant’s decision to adopt, abandon or make an exception to the practice.	“If I am temporarily changing my password so a family member can access my account when away, I will temporarily use an easier password” (strong-password)
distrust-service	Participants express fear of negative outcomes that originate from using this practice, such as security or privacy problems, or system/software failures, explicitly mentioning that they do not trust whoever is behind the service.	“I dislike the feeling of someone being able to have access to my entire list of passwords which is extremely nerve-racking” (password-manager)
forgetting	Participants say they forget to use the practice without specifying why they do or do not forget. They may express that abiding by the practice does not occur to them on a consistent basis unless prompted by the software/tool they are using.	“I wanted to check my credit report a few years ago and read you could do it for free. I did it that year and the next, but it slipped my mind since then” (credit-report)
just-started	Participants say they are in the process of adopting this practice but have not adopted it fully.	“I just started doing this and I still have some passwords that are not synced yet but will be soon” (password-manager)

Continued on next page

<b>Code</b>	<b>Meaning</b>	<b>Example Quote</b>
own-judgment-sufficient	Participants disregard the practice when they trust more on the judgments and decisions made by themselves or other entities that make security decisions for them. Also used when the user abandons the practice because they have found it to be unreliable or inaccurate, implying they can do better on their own.	“I don’t click on obvious spam e-mails, but I am willing to open e-mails that seem legitimate even if I don’t know the senders” (attachment-clicking)
practice-by-default	Participants say they do not intentionally choose to use this practice because the practice is implemented by default.	“Since https is now default I rarely check unless it’s specifically not green or padlocked” (HTTPS)
other	Participants say something that we feel is still valuable to be coded, but none of the current codes applies.	n/a
inapplicable	The participant believes that the described practice never arises for them or does not apply to their situation.	“I’ve never been asked for non-essential information” (hide-info)
unclear	Participants say something that can hardly be interpreted into any reasons, user describes nature of partial retention without stating why, or responses are not related to the question at all.	“Attempted virus installs have happened in the past. The installs were blocked, but malware is getting smarter so I don’t trust my security program to always catch a virus attempt” (attachment-clicking)

## APPENDIX E

# Supplemental Materials: “Icons for Conveying Privacy Controls”

## E.1 Survey Instruments

### Icon Design Evaluation

Please answer the following questions with regards to the displayed symbol [and phrase]. Make sure not to reveal any private or personally identifiable information about yourself or others in your responses to any open-ended questions.

*[The symbol or symbol/phrase condition to which the participant was randomly assigned is displayed to the participant.]*

1. What, if anything, does this [symbol/symbol and phrase] communicate to you?  
Please be as complete as possible. (Open-ended response)
2. Imagine if you saw this [symbol/link/symbol and link] on a website. What do you think would happen if you clicked on this [symbol/symbol and phrase]?  
(Open-ended response)

*[The DAA’s blue AdChoices icon is displayed]*

3. Have you ever seen this symbol on a website before?  Yes  No  I'm not sure
4. Imagine if you saw this symbol on a website. What do you think would happen if you clicked on this symbol?  
  
*[Present the icon set in randomized order.]*
5. Which of these symbols do you think best conveys that there's an option to tell websites "do not sell my personal information?" *[The order of Q5/6 and Q7/8 was randomized for the icon refinement testing.]*
6. Please explain why you selected the icon above. (Open-ended response)  
  
*[Present the icon set in randomized order.]*
7. Which of these symbols do you think best conveys that there's an option to make choices about the use of your personal information?
8. Please explain why you selected the icon above. (Open-ended response)
9. Are you aware of any laws in the United States that require companies to provide a "do not sell my personal information" option?  No  Yes (please name or describe them): \_\_\_\_\_
10. What is your age?  18-24  25-34  35-44  45-54  55-64  65-74  
 75-84  85 or older  Prefer not to answer
11. What is your gender?  Women  Men  Non-binary  Prefer to self-describe: \_\_\_\_\_  Prefer not to answer
12. What is the highest level of education you have completed?  Less than high school  High school degree of equivalent  Some college, no degree  Associate's degree, occupational  Associate's degree, academic  Bachelor's

degree    Master's degree    Professional degree    Doctoral degree    Prefer not to answer

13. What was your total household income before taxes during the past 12 months?
- Under \$15,000    \$15,000 to \$24,999    \$25,000 to \$34,999    \$35,000 to \$49,999    \$50,000 to \$74,999    \$75,000 to \$99,999    \$100,000 to \$149,999    \$150,000 or above    Prefer not to answer
14. In which state do you currently reside? (Drop-down list)
15. Which of the following best describes your educational background or job field?
- I have an education in, or work in, the field of computer science, computer engineering or IT.    I do not have an education in, or work in, the field of computer science, computer engineering or IT.    Prefer not to answer
16. If you have any feedback on the survey, please leave it here. (Open-ended response)

### **Link Text Evaluation**

Please answer the following questions with regards to the web link. Make sure not to reveal any private or personally identifiable information about yourself or others in your responses to any open-ended questions.

Imagine if you saw this web link on a website.

*[The link text condition to which the participant was randomly assigned is displayed to the participant.]*

1. What types of ["selling" / "personal information" / "choices" / "options" / "opt-outs"] do you think this link refers to? (Open-ended response, displayed only if the participant saw a link text that includes the respective element)
2. What do you think would happen if you clicked on this link?

3. Which of the following do you think could happen if you clicked this symbol/link on a web page *[For each statement below, participants were asked to choose from a 5-point likert scale “Definitely” “Probably” “Not sure” “Probably not” and “Definitely not.” Statements were presented in randomized order.]*
- It will take me to the website’s Terms of Service statement.
  - It will take me to a page that verifies that the website does not sell my personal information.
  - It will take me to a page where I can pay to protect my personal information.
  - It will take me to a page with choices about the sale of my personal information.
  - It will immediately communicate to the website that I do not want my personal information to be sold.
  - It will take me to a page with choices about how my personal information is used and shared by the website.
  - It will give the website permission to sell my personal information.
  - It will take me to a warning not to share my personal information with websites.
4. Are you aware of any laws in the United States that require companies to provide a “do not sell my personal information” option? (See Appendix E.1 for answer options)
5. What is your age? (See Appendix E.1 for answer options)
6. What is your gender? (See Appendix E.1 for answer options)
7. What is the highest level of education you have completed? (See Appendix E.1 for answer options)

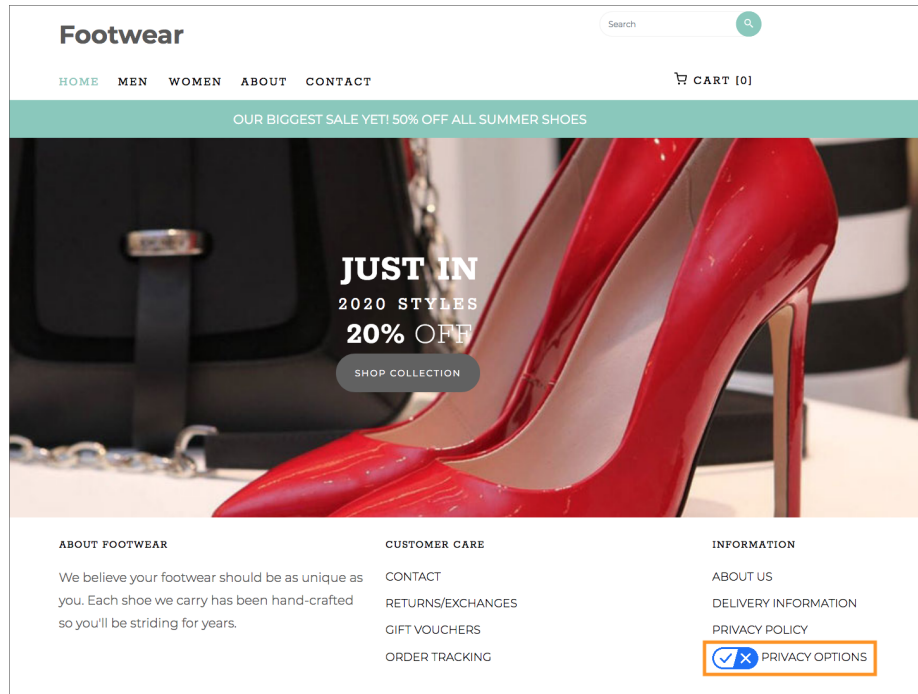


8. What was your total household income before taxes during the past 12 months?  
(See Appendix E.1 for answer options)
9. In which state do you currently reside? (Drop-down list)
10. Which of the following best describes your educational background or job field?  
(See Appendix E.1 for answer options)
11. Which of the following best describes your primary occupation?  Administrative support (e.g., secretary, assistant)  Art, Writing, or Journalism (e.g., author, reporter, sculptor)  Business, Management, or Financial (e.g., manager, accountant, banker)  Education or Science (e.g., teacher, professor, scientist)  Homemaker  Legal (e.g., lawyer, law consultant, or law professor)  Medical (e.g., doctor, nurse, dentist)  Engineering or IT Professional (e.g., programmer, IT consultant)  Service (e.g., retail clerk, server)  Skilled Labor (e.g., electrician, plumber, carpenter)  Unemployed  Retired  College student  Graduate student  Mechanical Turk worker  Other: \_\_\_\_\_  Prefer not to answer
12. If you have any feedback on the survey, please leave it here. (Open-ended response)

### Icon-Text Combinations Evaluation

Please answer the following questions with regards to the [symbol/link/symbol and link] in the rectangular highlighted area near the bottom of the web page displayed. Make sure not to reveal any private or personally identifiable information about yourself or others in your responses to any open-ended questions.

*[Display the screenshot of the web page with the icon, link text, or icon-text combination that the participant was randomly assigned to. Below is an example of one study condition]*



*Close up of highlighted area:*



1. What do you think would happen if you clicked on the [symbol/link/symbol and link] in the highlighted area on this web page?

*[Display close up of highlighted area again.]*

2. What do you think [“sell” / “info” / “information” / “choices” / “options” / “opt-outs”] refers to in this link? (Open-ended response, displayed only if the participant saw a link text that includes the respective element)

*[Display close up of highlighted area again.]*

3. Which of the following do you think could happen if you clicked this symbol/link on a web page? *[For each statement below, participants were asked to choose*

*from a 5-point likert scale “Definitely” “Probably” “Not sure” “Probably not” and “Definitely not.” Statements were presented in randomized order.]*

- It will take me to a page where I can update the information in my user profile on the website.
  - It will take me to a page with choices about how my personal information is used and shared by the website.
  - It will take me to a page with choices about the sale of my personal information.
  - It will take me to a page with more information about how the company uses and shares the personal information it collects about me.
  - It will cause the website to send unwanted emails.
  - It will give the website permission to sell my personal information.
  - It will take me to a page with ads about privacy and security products.
  - It will take me to a page that steals my information or has a virus or malware.
4. Are you aware of any laws in the United States that require companies to provide a “do not sell my personal information” option? (See Appendix E.1 for answer options)
  5. What is your age? (See Appendix E.1 for answer options)
  6. What is your gender? (See Appendix E.1 for answer options)
  7. What is the highest level of education you have completed? (See Appendix E.1 for answer options)
  8. What was your total household income before taxes during the past 12 months? (See Appendix E.1 for answer options)

9. In which state do you currently reside? (Drop-down list)
10. Which of the following best describes your educational background or job field?  
(See Appendix E.1 for answer options)
11. Which of the following best describes your primary occupation? (See Appendix E.1 for answer options)
12. If you have any feedback on the survey, please leave it here. (Open-ended response)

### **CCPA Toggle Icon Evaluation**

Please answer the following questions with regards to the symbol and link in the rectangular highlighted area near the bottom of the web page displayed. Make sure not to reveal any private or personally identifiable information about yourself or others in your responses to any open-ended questions.

*[Display the screenshot of the web page with the icon and “Do Not Sell My Personal Information” link text that the participant was randomly assigned to.]*

*[Display close up of highlighted area.]*

1. What do you think would happen if you clicked on the symbol and link in the highlighted area on this web page?

*[Display close up of highlighted area again.]*

2. What do you think [“sell”/“information”] refers to in this link? (Open-ended response)

*[Display close up of highlighted area again.]*

3. Which of the following do you think could happen if you clicked this symbol/link on a web page? *[For each statement below, participants were asked to choose*

*from a 5-point likert scale “Definitely” “Probably” “Not sure” “Probably not” and “Definitely not.” Statements were presented in randomized order.]*

- It will immediately change the setting on this website from “Do Not Sell My Personal Information” to “Sell My Personal Information.”
  - It will immediately change the setting on this website from “Sell My Personal Information” to “Do Not Sell My Personal Information.”
  - It will take me to a page where I can choose whether or not the website can sell my personal information.
  - It will take me to a page where I can confirm that I do not want my personal information to be sold by the website.
  - It will take me to a page with more information about how the website uses and shares my personal information.
  - It will cause the website to send me unwanted emails.
  - It will take me to a page with ads about privacy and security products.
  - It will take me to a page that steals my information or has a virus or malware.
4. Are you aware of any laws in the United States that require companies to provide a “do not sell my personal information” option? (See Appendix E.1 for answer options)
  5. What is your age? (See Appendix E.1 for answer options)
  6. What is your gender? (See Appendix E.1 for answer options)
  7. What is the highest level of education you have completed? (See Appendix E.1 for answer options)

8. What was your total household income before taxes during the past 12 months?  
(See Appendix E.1 for answer options)
9. In which state do you currently reside? (Drop-down list)
10. Which of the following best describes your educational background or job field?  
(See Appendix E.1 for answer options)
11. Which of the following best describes your primary occupation? (See Appendix E.1 for answer options)
12. If you have any feedback on the survey, please leave it here. (Open-ended response)

## **E.2 Detailed Statistical Results**

	Icon Preliminary	Icon Refinement	Link Text Preliminary	Link Text Refinement	Combination	Toggle
<b>Sample Size</b>	240	280	140	400	1468	421
<b>Invalid Responses 11</b>	0	9	0	54	18	
<b>Age</b>						
18-24	5.00%	5.71%	8.57%	7.00%	10.29%	12.11%
25-34	45.00%	45.71%	52.14%	49.00%	35.76%	45.13%
35-44	29.58%	29.64%	22.86%	23.25%	25.95%	23.04%
45-54	12.08%	10.00%	8.57%	11.00%	15.74%	11.16%
55-64	7.08%	6.79%	4.29%	7.00%	8.72%	6.18%
>65	1.25%	2.14%	3.57%	2.75%	3.13%	1.90%
Prefer Not to Answer	0.00%	0.00%	0.00%	0.00%	0.41%	0.28%
<b>Gender</b>						
Female	41.25%	44.64%	34.29%	39.50%	46.87%	47.51%
Male	58.33%	55.36%	64.29%	60.00%	51.98%	50.83%
Non-binary	0.00%	0.00%	1.43%	0.50%	0.41%	0.95%
Self-described	0.00%	0.00%	0.00%	0.00%	0.14%	0.00%
Prefer Not to Answer	0.42%	0.00%	0.00%	0.00%	0.61%	0.71%
<b>Education</b>						
Less than High School	0.83%	0.71%	0.71%	0.25%	0.54%	0.48%
High School	9.58%	13.57%	7.86%	13.50%	8.92%	12.11%
Some College	15.83%	18.57%	17.14%	18.00%	21.93%	18.76%
Associate's, Occupational	5.83%	8.21%	4.29%	7.75%	6.81%	5.23%
Associate's, Academic	1.67%	5.00%	4.29%	4.75%	6.40%	5.46%
Bachelor	49.58%	42.86%	51.43%	43.75%	39.03%	43.71%
Master	13.33%	8.21%	13.57%	11.00%	11.92%	9.26%
Professional	2.50%	1.79%	0.71%	0.50%	1.63%	2.38%
Doctoral	0.42%	1.07%	0.00%	0.50%	2.18%	2.14%
Prefer Not to Answer	0.42%	0.00%	0.00%	0.00%	0.61%	0.48%
<b>State Residence</b>						
California	11.25%	14.29%	20.00%	9.75%	10.42%	16.39%
Non-California	88.75%	85.71%	80.00%	90.25%	89.58%	83.61%
<b>Educational/Job Background related to CS/IT</b>						
Yes	38.75%	27.50%	47.86%	33.00%	23.02%	30.64%
No	56.67%	66.79%	47.86%	62.75%	72.55%	63.66%
Prefer Not to Answer	4.58%	5.71%	4.29%	4.25%	4.43%	5.70%
<b>Awareness of any U.S. laws that require companies to provide a "do not sell my personal information" option</b>						
Yes	4.58%	4.64%	2.86%	3.00%	9.81%	7.13%
No	95.42%	95.36%	97.14%	97.00%	90.19%	92.87%

Table E.1: Age, gender, education, state residence demographics of participants as well as their familiarity with CCPA in each study.


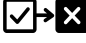










Name	Icon	Common Interpretations (# of Participants)
<i>Stylized-Toggle</i>		<b>accept/decline</b> (4); <b>activate/deactivate</b> (2); true/false (2); mark as completed (1)
<i>Changed-Choice</i>		okay/exit options (1); <b>accept/decline</b> (1); true/false (1); opposite is true (1); no guesses (2)
<i>DoNot-Checked</i>		<b>activate/deactivate</b> (2); mark as completed (2); completed downloads (2); <b>accept/decline</b> (1)
<i>Box-Arrow</i>		<b>removing something</b> (2); okay/exit options (2); email or message (1); no guesses (1)
<i>Circle-Arrow</i>		move forward/go (3); email or message (1); no guesses (2)
<i>Folder-Arrow</i>		folder/file (4); email or message (3)
<i>DoNot-Dollar</i>		cancel payment (2); losing money (2); low balance (2); money/paying (2); cash/dollars not accepted (1); something is free or requires no money (1)
<i>Slash-Dollar</i>		cash/dollars not accepted (4); something is free or requires no money (3); money/paying (1)
<i>Stop-Dollar</i>		money/paying (4); account balance (2); something costs money (2); something is free or requires no money (1); cash/dollars not accepted (1)
<i>ID-Card</i>		payment method (4); <b>something related to a person and money</b> (3); something costs money (2); account balance (1); no guesses (1)
<i>Profile</i>		money/paying (2); stop spending money (2); something costs money (2);
<i>DAA</i>		more information (3); move forward/go (2); play button (2)

Table E.2: Participants’ coded open-ended responses to “What does this symbol communicate to you?” from conditions in which the icon was shown without a link text in the icon preliminary testing, along with a code’s number of occurrences. Interpretations that align with the icon’s intended meaning are bolded.








Name	Icon	Common Interpretations (# of Participants)
<i>ID-Card</i>		something costs money (9); sending money to someone (5); money/paying (5); <b>something related to a person and money</b> (3); account balance (3); price related (2); payment methods accepted by website (2)
<i>Slash-Dollar</i>		something is free or requires no money (12); cash/dollars not accepted (6); money/paying (4); <b>selling is not allowed</b> (1)
<i>Stop-Dollar</i>		money/paying (10); stop spending money (5); something costs money (4); price related (3); sale/discount (3); no guesses (3)
<i>Stylized-Toggle</i>		<b>accept/decline something</b> (11); <b>activate/deactivate something</b> (4); true/false (4); okay/exit options (3)
<i>DAA</i>		more information (11); play button (7); move forward/go (3); ad related (2)

Table E.3: Participants’ coded open-ended responses to “What does this symbol communicate to you?” from conditions in which we showed the icon without a link text in the refined icons study, along with a code’s number of occurrences. Interpretations that align with the icon’s intended meaning are bolded.

	Choice			Privacy			Misconception			Do-Not-Sell		
	$\beta$	S.E.	Adj. $p$	$\beta$	S.E.	Adj. $p$	$\beta$	S.E.	Adj. $p$	$\beta$	S.E.	Adj. $p$
Intercept	0.52	0.31	1.0	2.0	0.41	<.001*	-3.1	0.74	.001*	-0.06	0.33	1.0
<b>Condition (ref = <i>Toggle-Privacy Options, None-Do Not Sell My Personal Information</i> for Do-Not-Sell)</b>												
<i>Toggle-None</i>	-1.0	0.40	.16	-1.2	0.48	.20	2.1	0.80	.21			
<i>Toggle-Do Not Sell My P.I.</i>	-0.62	0.40	1.0	-4.0	0.58	<.001*	2.4	0.79	.06	0.26	0.40	1.0
<i>Toggle-Personal Info Choices</i>	-1.2	0.40	.06	-2.4	0.47	<.001*	2.8	0.78	.009*	-3.2	0.66	<.001*
<i>Toggle-Do Not Sell My Info</i>	-0.33	0.40	1.0	-4.3	0.60	<.001*	1.5	0.82	.77	0.52	0.39	1.0
<i>Toggle-Privacy Choices</i>	-0.23	0.40	1.0	-0.86	0.48	.85	-1.6	0.82	.76	2.2	0.49	<.001*
<i>Toggle-Privacy Options</i>										-4.3	1.0	.001*
<i>DAA-None</i>	-3.6	0.60	<.001*	-3.0	0.49	<.001*	2.2	0.79	.14			
<i>DAA-Do Not Sell My P.I.</i>	-0.49	0.40	1.0	-2.6	0.47	<.001*	0.81	0.89	1.0	-0.09	0.38	1.0
<i>DAA-Privacy Options</i>	0.29	0.42	1.0	-0.26	0.52	1.0	-0.02	1.0	1.0	-4.3	1.0	.001*
<i>DAA-Personal Info Choices</i>	-1.6	0.41	.004*	-2.4	0.48	<.001*	2.5	0.79	.04*	-2.6	0.59	<.001*
<i>DAA-Do Not Sell My Info</i>	-0.59	0.40	1.0	-3.5	0.52	<.001*	1.3	0.84	1.0	0.25	0.39	1.0
<i>DAA-Privacy Choices</i>	0.10	0.41	1.0	-0.67	0.49	1.0	0.06	1.0	1.0	-2.4	0.51	<.001*
<i>Dollar-None</i>	-3.0	0.50	<.001*	-5.3	0.82	<.001*	4.2	0.78	<.001*			
<i>Dollar-Do Not Sell My P.I.</i>	-0.92	0.39	.32	-3.6	0.52	<.001*	1.6	0.82	.78	0.02	0.38	1.0
<i>Dollar-Privacy Options</i>	-0.28	0.40	1.0	-0.52	0.51	1.0	1.7	0.81	.53	-2.1	0.49	<.001*
<i>Dollar-Personal Info Choices</i>	-1.3	0.40	.03*	-2.0	0.46	<.001*	2.9	0.78	.005*	-2.8	0.59	<.001*
<i>Dollar-Do Not Sell My Info</i>	-0.47	0.40	1.0	-3.7	0.53	<.001*	-0.55	1.2	1.0	0.91	0.41	.36
<i>Dollar-Privacy Choices</i>	-1.1	0.39	.10	-0.82	0.48	.91	2.0	0.79	.21	-2.7	0.59	<.001*
<i>None-Do Not Sell My P.I.</i>	-0.94	0.40	.32	-3.3	0.51	<.001*	1.2	0.84	1.0			
<i>None-Privacy Options</i>	-0.48	0.40	1.0	-0.65	0.49	1.0	0.03	1.0	1.0	-2.8	.59	<.001*
<i>None-Personal Info Choices</i>	-1.3	0.40	.02*	-2.0	0.46	<.001*	2.8	0.78	.008*	-2.8	0.59	<.001*
<i>None-Do Not Sell My Info</i>	0.10	0.43	1.0	-3.4	0.52	<.001*	-0.58	1.2	1.0	0.62	0.40	1.0
<i>None-Privacy Choices</i>	-0.64	0.40	1.0	-0.94	0.49	.71	0.50	0.94	1.0	-2.4	0.55	<.001*
<i>Icon-None</i>										-4.6	0.76	<.001
<b>Age (ref = 18-34)</b>												
35-54	0.42	0.13	.02*	-0.07	0.15	1.0	-0.18	0.18	1.0	0.29	0.18	1.0
$\geq 55$	0.35	0.20	1.0	-0.42	0.22	.74	0.43	0.25	.97	0.58	0.27	.36
<b>Gender (ref = Female)</b>												
Male	0.10	0.12	1.0	0.18	0.14	-.16	0.17	1.0	1.0	-0.11	0.17	1.0
<b>Technical expertise (ref = None)</b>												
Yes	-0.42	0.15	.13	-0.11	0.17	1.0	0.46	0.20	.32	-0.68	0.22	.02*
<b>Highest obtained education (ref = No college degree)</b>												
College degree	0.33	0.14	.27	-0.13	0.15	1.0	-0.46	0.18	.22	0.20	0.19	1.0
Graduate degree	0.31	0.19	1.0	0.11	0.21	1.0	-0.61	0.26	.32	0.39	0.26	1.0

Due to perfect separation in the *DAA-none* and *Dollar-none* conditions, the icon-only conditions were collapsed together (*Icon-None*) for the do-not-sell choices regression. For each regression term we provide the estimate of the coefficient ( $\beta$ ), the standard error, and p-value adjusted with the Holm-Bonferroni correction. A significant negative coefficient indicates that participants in that group were less likely to have the expectation represented by the dependent variable, relative to the reference baseline. (\*) marks significant results for  $\alpha = .05$ .

Table E.4: Regression results for the four binary dependent variables (conveys *choice*, *privacy*, *misconceptions*, or *do-not-sell choices*) coded from participants' open-ended expectations.

	Do-Not-Sell Choices			Give Sell Permission			Privacy Choices		
	$\beta$	S.E.	Adj. $p$	$\beta$	S.E.	Adj. $p$	$\beta$	S.E.	Adj. $p$
Intercept	0.41	0.31	1.0	-2.0	0.44	<.001*	2.5	0.54	<.001*
<b>Condition (ref = None-Do Not Sell My Personal Information, Toggle-Privacy Options for Privacy Choices)</b>									
<i>Toggle-None</i>	-0.98	0.39	.33	0.56	0.53	1.0	-1.8	0.60	.07
<i>Toggle-Do Not Sell My P.I.</i>	-0.22	0.40	1.0	1.7	0.49	.02*	-2.1	0.59	.009*
<i>Toggle-Personal Info Choices</i>	-0.87	0.39	.60	0.96	0.50	1.0	-1.7	0.60	.10
<i>Toggle-Do Not Sell My Info</i>	-0.55	0.39	1.0	0.65	0.52	1.0	-1.8	0.60	.07
<i>Toggle-Privacy Choices</i>	-0.09	0.39	1.0	1.4	0.49	.12	-1.1	0.62	.99
<i>Toggle-Privacy Options</i>	0.07	0.39	1.0	0.87	0.51	1.0			
<i>DAA-None</i>	-2.9	0.52	<.001*	-0.45	0.62	1.0	-2.4	0.58	.001*
<i>DAA-Do Not Sell My P.I.</i>	-0.19	0.39	1.0	0.11	0.56	1.0	-0.73	0.65	1.0
<i>DAA-Privacy Options</i>	-0.56	0.38	1.0	0.93	0.50	1.0	-0.62	0.66	1.0
<i>DAA-Personal Info Choices</i>	-0.20	0.40	1.0	1.2	0.50	.47	-1.8	0.60	.06
<i>DAA-Do Not Sell My Info</i>	0.18	0.41	1.0	-0.22	0.59	1.0	-1.6	0.61	.19
<i>DAA-Privacy Choices</i>	-0.23	0.39	1.0	0.85	0.51	1.0	-0.25	0.70	1.0
<i>Dollar-None</i>	-1.3	0.39	.04*	-0.38	0.62	1.0	-3.6	0.60	<.001*
<i>Dollar-Do Not Sell My P.I.</i>	-0.29	0.39	1.0	0.12	0.56	1.0	-1.7	0.59	.06
<i>Dollar-Privacy Options</i>	-0.05	0.40	1.0	0.48	0.53	1.0	-0.69	0.66	1.0
<i>Dollar-Personal Info Choices</i>	-0.36	0.39	1.0	0.56	0.53	1.0	-1.1	0.63	1.0
<i>Dollar-Do Not Sell My Info</i>	-0.14	0.39	1.0	-0.58	0.66	1.0	-1.3	0.61	.52
<i>Dollar-Privacy Choices</i>	-0.65	0.38	1.0	0.57	0.51	1.0	-1.4	0.60	.36
<i>None-Privacy Options</i>	-0.20	0.39	1.0	0.82	0.51	1.0	-0.42	0.68	1.0
<i>None-Personal Info Choices</i>	-0.84	0.39	.72	-0.09	0.57	1.0	-1.7	0.60	.08
<i>None-Do Not Sell My Info</i>	0.84	0.45	1.0	0.11	0.57	1.0	-0.70	0.66	1.0
<i>None-Privacy Choices</i>	-0.24	0.40	1.0	0.55	0.53	1.0	-0.52	0.68	1.0
<i>None-Do Not Sell My P.I.</i>							-0.86	0.65	1.0
<b>Age (ref = 18-34)</b>									
35-54	0.09	0.12	1.0	-0.04	0.15	1.0	0.27	0.15	1.0
$\geq 55$	0.17	0.19	1.0	-0.008	0.23	1.0	0.16	0.23	1.0
<b>Gender (ref = Female)</b>									
Male	-0.02	0.12	1.0	-0.13	0.15	1.0	-0.21	0.14	1.0
<b>Technical expertise (ref = None)</b>									
Yes	0.03	0.15	1.0	0.72	0.17	<.001*	-0.03	0.18	1.0
<b>Highest obtained education (ref = No college degree)</b>									
College degree	0.31	0.13	.51	-0.09	0.16	1.0	0.15	0.16	1.0
Graduate degree	0.17	0.18	1.0	0.07	0.22	1.0	0.63	0.24	.15

For each regression term we provide the estimate of the coefficient ( $\beta$ ), the standard error, and p-value adjusted with the Holm-Bonferroni correction. A significant negative coefficient indicates that participants in that group were less likely to have the expectation represented by the dependent variable, relative to the reference baseline. (\*) marks significant results for  $\alpha = .05$ .

Table E.5: Regression results for the scenarios: “It will take me to a page with choices about the sale of my personal information” (Do-Not-Sell Choices), “It will give the website permission to sell my personal information” (Give Sell Permission), and “It will take me to a page with choices about how my personal information is used and shared by the website” (Privacy Choices).

	Misconception			Toggle Control		
	$\beta$	S.E.	$p$	$\beta$	S.E.	$p$
Intercept	-1.2	.34	<.001*	-1.6	.36	<.001*
<b>Condition (style ref = <i>Stylized Toggle</i>; color ref = <i>Blue</i>)</b>						
CalAG	.83	.28	.003*	.87	.30	.003*
CalAG-X	.86	.28	.003*	.87	.30	.004*
Red	.003	.22	.99	.28	.23	.23
<b>Age (ref = 18-34)</b>						
35-54	-.24	.24	.31	-.15	.25	.54
$\geq 55$	-.99	.49	.04*	-1.2	.58	.03*
<b>Gender (ref = Female)</b>						
Male	-.08	.23	.73	.03	.24	.89
<b>Technical expertise (ref = None)</b>						
Yes	-.47	.25	.06	-.86	.27	.002*
<b>Highest obtained education (ref = No college degree)</b>						
College degree	.65	.26	.01*	.68	.27	.01*
Graduate degree	.50	.36	.17	.71	.38	.06

We report results from the main effect model with icon style and color as the main independent variables, as the interaction between icon style and color was not significant. For each regression term we provide the estimate of the coefficient ( $\beta$ ), the standard error, and p-value adjusted with the Holm-Bonferroni correction. A significant positive coefficient indicates that participants in that group were more likely to have the expectation represented by the dependent variable, relative to the reference baseline. (\*) marks significant results for  $\alpha = .05$ .

Table E.6: Regression results for the binary dependent variables: conveys a *misconception* and perceived as a *toggle control*, coded from participants' open-ended expectations.

	Do-Not-Sell Switch			Do-Not-Sell Choices		
	$\beta$	S.E.	Adj. $p$	$\beta$	S.E.	Adj. $p$
Intercept	-.77	.35	.03*	.70	.31	.03*
<b>Condition (style ref = <i>Stylized Toggle</i>; color ref = <i>Blue</i>)</b>						
CalAG	.73	.38	.05	-.92	.27	<.001*
CalAG-X	1.0	.39	.009*	-.88	.27	.001*
Red	1.0	.37	.006*	.14	.22	.51
CalAG*Red	-1.2	.52	.02*			
CalAG-X*Red	-1.8	.53	<.001*			
<b>Age (ref = 18-34)</b>						
35-54	.24	.23	.31	.04	.23	.87
$\geq 55$	.29	.39	.46	.60	.41	.14
<b>Gender (ref = Female)</b>						
Male	.04	.22	.84	-.53	.22	.02*
<b>Technical expertise (ref = None)</b>						
Yes	.21	.24	.37	-.35	.24	.15
<b>Highest obtained education (ref = No college degree)</b>						
College degree	-.15	.24	.52	-.006	.24	.98
Graduate degree	-.07	.34	.83	-.08	.35	.83

We report results from the main effect model for Do-Not-Sell Choices with icon style and color as the main independent variables, as the interaction between icon style and color was not significant. For each regression term we provide the estimate of the coefficient ( $\beta$ ), the standard error, and p-value adjusted with the Holm-Bonferroni correction. A significant positive coefficient indicates that participants in that group were more likely to have the expectation represented by the dependent variable, relative to the reference baseline. (\*) marks significant results for  $\alpha = .05$ .

Table E.7: Regression results for the scenarios: “It will immediately change the setting on this website from ‘Sell My Personal Information’ to ‘Do Not Sell My Personal Information’ ” (Do-Not-Sell Switch), and “It will take me to a page with choices about the sale of my personal information” (Do-Not-Sell Choices).

## APPENDIX F

# Supplemental Materials: “Computer Security Customer Support for Survivors of Intimate Partner Violence”

### F.1 Focus Group Protocol: IPV Professionals

#### Part 1: Introduction

Thank you all for taking the time to talk to us. We’re researchers from [institution names]. [Company name] provides cybersecurity software and services like [product names].

[Company name] offers customer support hotlines and online chats to help their customers deal with tech-related issues. There are instances in which the caller appears to be in a dangerous situation, such as stalking and domestic violence. [Company name] wants to better assist these callers and understand the appropriate scope for their customer support team in doing so.

Today’s meeting will be primarily discussion-based with a few activities. There are no right or wrong answers to any of our questions. We’re simply interested in

your opinions based on your own experiences or perspectives. You can choose not to comment if you don't want to, and you can quit the session at any point.

We would also like to get your consent to audio record the workshop session as a backup of our notes. These will be transcribed, all identifying information will be removed, and we will destroy the original recordings once the transcription is done. Are you ok with us recording the meeting? Do you have any other questions before we get started?

- Let's go around the room with brief introductions. Please tell us your name, job title, and how many years you have been doing this job.
- Have you ever worked directly with clients? Have you encountered clients who have experienced IPV?
- Have you encountered clients who have experienced tech-related abuse? Can you give an example?

## **Part 2: Presenting and Discussing Customer Support Scenarios**

Now we'd like to present a few example customer support transcripts and get your expert opinions on these interactions. These transcripts are based on real chats, but have been shortened and identifying information removed. We'll let you read each scenario and ask a few follow-up questions.

*IPV professionals' own advice to this customer* Ignoring the technical aspect of this problem for a moment, imagine someone were to come to you with this problem...

- Are these problems similar to or different from the cases you normally receive at your organization? In what ways?
- What advice would you give this customer based on the available information?

*IPV professionals' advice to the support agent* Now let's think about this customer's interaction with customer support...

- In your opinion, what could the customer support agent offer this customer beyond assistance with [product name]?
- Are there additional questions that customer support should be asking?
- Are there resources customer support could have shared?
- In your opinion, should customer support point the caller to other organizations, such as family shelters or the police? Why or why not? If yes, how might it be done?
- In your opinion, should customer support provide specific advice about safety planning? Why or why not? If yes, how might it be done?

*Factors that might complicate advice* Let's discuss a few factors that make the situation trickier. For each case, should the support agent react differently in your opinion, why or why not?

- What if the support agent thinks the attacker is recording or listening to the chat?
- What if the customer is not alone when the call takes place?
- What if the attacker could be the person calling to gain more access to a victim's account?

### **Part 3: General Advice Going Beyond the Scenarios**

Now that we've looked at some examples of the problems that customer support gets, let's think about the broader role that customer support can play in providing support to victims of abuse.



- Under what circumstances, if any, do you think that customer support’s duty to help extends beyond addressing product-specific issues identified by the customer?
- In your opinion, should customer support try to identify situations in which the caller may need additional safety planning advice? Why or why not? If yes, what can customer support do to quickly identify such situations?
- Should customer support watch out for cues that indicate further questions would be unsafe (e.g., the conversation might be monitored)? Why or why not?
- How should customer support respond if a customer reveals personal, sensitive information about an assault or about suicide?
- What training or education do you think the support rep could have to help them avoid adverse outcomes? E.g., about empathetic language? About IPV and risks related to leaving an abuser? About resources to share?
- Any final thoughts?

## **F.2 Focus Group Protocol: Support Agents**

### **Introduction**

We are conducting a research study around technology and intimate partner abuse, or IPV. We are exploring how security companies can help IPV survivors through their customer support.

So far we’ve conducted five focus groups with about 20 experts in this space, such as social workers and legal advocates, to collect their feedback on this topic. We now want to talk to you as customer support practitioners and security experts, to understand how effective, efficient, and practical some of these ideas are.

After our talk today, we plan to develop recommendations from these insights and integrate them into guidelines and training materials for customer support representatives, and we're more than happy to share them with you.

Any questions so far?

## **Part 1: Study Background**

We conducted 5 focus groups with professionals to seek advice about how security companies can support IPV survivors. We presented three scenarios, created based on real chat transcripts from [Company], and asked participants how customer support could do better.

The ideas we elicited from IPV professionals are not final. Participants sometimes disagreed with each other, and also mentioned the challenges and constraints of some of the ideas.

Scenario A (see Figure 8.1) is an example of a real-world customer support chat we saw at [Company].

We'd like to ask some open questions about your organization:

- How is your customer support team organized?
- What are evaluation metrics for success for customer support representatives?

We'd now like to ask about your experiences with IPV at your company.

- Have you or your employees encountered similar cases that involve IPV/technology abuse?
- What are your company's current efforts for supporting IPV survivors, that you're aware of?

We plan to present our findings to you in four parts. During our presentation, please feel free to chime in anytime if you have any questions or comments.

At the end of each part we'll have a short summary and discussion to ask you some specific questions about what we shared with you and get your feedback. We expect each session to take about 12 minutes, and we'll leave a few minutes at the end of today's meeting to wrap up and discuss next steps.

## **Part 2: Interacting with customers**

Suggestions from IPV professionals:

- Explain why a product would be helpful
- Avoid overpromising
- Ask more probing questions

Questions:

- What are your initial reactions to these suggestions? Any comments or feedback?
- Do you think it is feasible?
- How much of this would you say is your team already doing?
- Would this create conflict with your evaluation metrics of support agents, such as the rate of "resolving issues?" If yes, Is there any way to mitigate such conflict?
- Do you see any challenges or concerns with these suggestions?
- [If a challenge is identified] Do you have ideas about how this could be done differently?

### **Part 3: Advice given to customers**

Open question before we present specific suggestions:

- What role do you think customer support should play in providing technical assistance vs. going beyond?

Suggestions from IPV professionals:

- Discuss potential consequences of given advice
- Provide resources for best security and safety practices

Questions:

- What are your initial reactions to these suggestions? Any comments or feedback?
- Have your employees already been discussing consequences of advice? If yes, could you give us an example?
- Is there any downside of discussing potential consequences of given advice?
- What resources do your support agents refer customers to about security and safety practices? How often do they do this?
- Do you see any challenges or concerns with these suggestions?
- [If a challenge is identified] Do you have ideas about how this could be done differently?

### **Part 4: Making referrals**

Suggestions from IPV professionals:

- Refer customers to a specialized team within the company

- Make external referrals based on trigger words

Questions:

- What are your initial reactions to these suggestions? Any comments or feedback?
- Do you already have a multi-tiered support system? What types of cases get transferred or escalated?
- How feasible do you think is it to have a specialized team within your company to deal with IPV/tech abuse?
- Do your support reps already refer customers to resources outside of the company? If yes, for what types of problems?
- From your experience, how difficult would it be to identify these cases? What are the challenges?
- Do you see any challenges or concerns with these suggestions?
- [If a challenge is identified] Do you have ideas about how this could be done differently?

## **Part 5: Training materials**

Suggestions from IPV professionals:

- Have agents be familiar with common tech abuse cases
- Train agents for trauma-informed responses
- Ensure the well-being of support agents

Questions:

- What are your initial reactions to these suggestions? Any comments or feedback?
- Have you embedded training for empathetic or trauma-informed responses in your current training materials/scripts?
- Are you already doing anything to prepare agents to handle difficult / traumatic customer issues?
- What things have you done to ensure the well-being of your employees? Could anything be done better?
- Do you see any challenges or concerns with these suggestions?
- [If a challenge is identified] Do you have ideas about how this could be done differently?

## **Closing**

We want to use the insights from both our work with IPV experts and customer support teams, like you, to develop guidelines and training materials for integrating IPV support into technical customer support. If you and your company are interested, we will follow up with you when we have drafted materials to share them with you.

## APPENDIX G

# Supplemental Materials: Nudges for Encouraging Password Changes After Data Breaches

## G.1 Survey Material

### G.1.1 Screening Survey

#### Informed consent

*[This informed consent form was for the screening survey. We used a very similar consent form for the main and follow-up surveys, except adjusting the estimated completion time and compensation.]*

**Study Title:** Consumers' Reactions Toward Data Breaches

**Principal Investigators:** REDACTED

**Purpose of this Study:** We are conducting a research study to understand how consumers perceive and react to data breaches.

#### Description of your involvement:

If you agree to be part of the research study, we will ask you to complete a screening survey that asks you to provide an email address. We will query this email

address in haveibeenpwned.com, a public database, to see if it has appeared in any data breaches.

Depending on the query results, we may send you an invitation to our main survey following your completion. The main survey will show more information about these breaches and ask you to answer a few questions about them.

**Compensation:**

We expect this screening survey to take **about two to three minutes**. You will be compensated **\$0.80** upon completing the survey.

You are free to withdraw at any time. However, you will not be compensated if you withdraw from the study.

**Benefits and Risks:** Although you may not directly benefit from participating in this study, the study will inform how to better protect consumers against data breaches.

The risks associated with your participation are similar to those normally encountered when using the Internet. We take strong measures to protect your personal information, as we outline in the “Confidentiality” section.

**Confidentiality:** By participating in the study, you understand and agree that the *REDACTED* may be required to disclose your consent form, data and other personally identifiable information as required by law, regulation, subpoena or court order.

Otherwise, your confidentiality will be maintained in the following manner:

- Your data will be stored in password-protected cloud services and will only be accessible to the study team.
- We plan to publish the results of this study, but we will not include any information that would identify you.
- Throughout the study, you will not be asked to provide any direct personal identifiers in the study apart from your email address. **We do not track your email address, and we will not be able to tie your email address to**



any results or analysis. All records of your email address will reside in temporary storage to facilitate the lookup of data breaches, and will be deleted following the completion of this task. We will never see your actual email address.

**Right to Ask Questions & Contact Information:** If you have questions about this research, you may contact the study team at *REDACTED*.

The *REDACTED* Institutional Review Board has determined that this study is exempt from IRB oversight.

**Voluntary Consent:** By proceeding to the next page, you are agreeing to participate in this study. You may print a copy of this consent form for your records. You may also contact the study team at any time by emailing *REDACTED* if you think of a question later.

Are you 18 years of age or older?  Yes  No

Are you physically living in the United States?  Yes  No

*[Participants need to answer “Yes” for both questions to proceed to the next page. If the answer is “No” to any question, show an error message “We’re sorry, but to participate in this survey you must be at least 18 and currently physically located in the United States. Thank you for your interest.”]*

### **Breach lookup**

We are going to ask you to enter your most commonly used email address at the bottom of this page. We will use your email address to look up whether your email address has appeared in any data breaches (also called “security breaches”), using the public lookup service for data breaches [haveibeenpwned.com](https://haveibeenpwned.com). Based on the results, we may invite you to our main survey in which we will show you more information about these breaches.

**Privacy Notice: We do not track or store your email address as part of this study, and we will not be able tie your email address to any results or analysis.** All records of your email address will reside in temporary storage to facilitate the lookup of data breaches, and will be deleted following the completion of this task. We will never see your actual email address.

To access information about breaches, your email address will be communicated to haveibeenpwned.com, a public service that maintains a database of data breaches involving email addresses. Communication with haveibeenpwned.com will occur on secure and encrypted channels. haveibeenpwned.com does not permanently store email addresses used in queries as described in their privacy policy.

If you have any further concerns about providing your email address, you may opt-out of the survey at this time. We will remove any record of your participation. Note that if you choose to opt out, you will not be compensated.

1. Please enter your most commonly used email address. After the task, you may search for another email address, but for now, we are primarily interested in breaches that may have involved your most commonly used email address. [free text]

### **Email-related questions**

Please tell us more about this email address.

2. Whose email address is it?  It is my own account / I have sole ownership of this account  It is my shared account / I share the account with someone else (e.g., a partner or family member)  It is someone else's account / someone else has sole ownership of this account  I made up an email address just for this study
3. How often do you check emails in this account?  Every day  A few times a week  A few times a month  A few times a year or less frequently

4. What do you use this email account for? Choose all that apply.  For personal correspondence (e.g., friends and family members)  For professional correspondence (e.g., with colleagues, business partners)  Sign up for sensitive accounts (e.g., banking, taxes)  Sign up for medium sensitive accounts (e.g., social media, online shopping)  Sign up for low-value accounts (I used it when I'm prompted to sign up but don't really care)  Other [free-text]
5. How long have you been using this email account? [number entry]  year(s)  month(s)  week(s)  day(s)
6. How many other email addresses/accounts do you regularly use? (Not counting the one you entered) [number entry]
7. Prior to our study, have you ever checked if this email address has appeared in any data breaches using Have I Been Pwned (haveibeenpwned.com) or other services?  Yes  No  Unsure

### **Demographics-related questions**

As the final part of this survey, please tell us a few things about yourself.

8. What is your age? [number entry or "prefer not to say"]
9. What is your gender?  Man  Woman  Non-Binary  Prefer to self-describe: [free text]  Prefer not to say
10. What is the highest level of education you have completed?  Less than high school  High school or equivalent  Some college, no degree  Associate's degree, occupational  Associate's degree, academic  Bachelor's Degree  Master's Degree  Professional degree  Doctoral degree  Prefer not to say
11. Have you studied or worked in the field of computer science or information technology?  Yes  No  Prefer not to answer
12. Have you studied or practiced law or other legal services?  Yes  No

Prefer not to answer

13. What was your total household income before taxes during the past 12 months?
- Under \$15,000
  - \$15,000 to \$24,999
  - \$25,000 to \$34,999
  - \$35,000 to \$49,999
  - \$50,000 to \$74,999
  - \$75,000 to \$99,999
  - \$100,000 to \$149,999
  - \$150,000 or above
  - Prefer not to say
14. Which of the following best describes you? Choose one or more.
- White
  - Black or African American
  - Asian
  - American Indian or Alaska Native
  - Native Hawaiian or other Pacific Islander
  - Middle Eastern or North African descent
  - Prefer to self-describe: [free text]
15. Are you Hispanic or Latino?  Yes  No  Prefer not to say

### **Next steps (only for eligible participants)**

We will send you the link to our main survey within the next week. In the main survey, we will show you more details about the data breach records associated with your provided email address.

Please click the “continue” button to proceed to the next page, where you will be automatically redirected to Prolific.

Stay tuned, and we look forward to seeing you back!

### **Debrief (only for ineligible participants)**

*[For participants who did not have any password breach but had other data breaches associated with their provided email address:]*

Thank you for completing our screening survey. We’re looking for participants who have at least one password breach (i.e., a data breach that exposes users’ account passwords) to take our main survey. According to haveibeenpwned.com, your email address has not appeared in any password breach.

That being said, your email address has appeared in the following data breaches,

and we suggest that you take necessary precautions. Please note that you can always obtain the same results by checking your email address on [haveibeenpwned.com](https://haveibeenpwned.com), which, in addition, provides records with sensitive breaches upon the verification of your email account. Please keep in mind that this list only reflects breaches that are registered in the [haveibeenpwned.com](https://haveibeenpwned.com) database, your information may have been exposed in other breaches.

*[For participants who did not have any data breach with their provided email address:]*

Thank you for completing our screening survey. We're looking for participants who have at least one password breach (i.e., a data breach that exposes users' account passwords) to take our main survey. According to [haveibeenpwned.com](https://haveibeenpwned.com), your email address has not appeared in any data breach.

That's great news! However, we still recommend that you monitor breach services like [haveibeenpwned.com](https://haveibeenpwned.com) in case new breaches come to light. In case you still want to learn about what to do when you're affected by a breach, you can find several resources below.

*[For all ineligible participants:]*

Below is a list of resources to help you better protect yourself from data breaches.

- Resources about recovering from a data breach:
  - Federal Trade Commission: Identity theft recovery steps
  - Federal Trade Commission: Credit Freeze FAQs
  - Firefox Monitor: What to do after a data breach
  - Norton: What to do after 5 types of data breaches
- Resources about protecting yourself against future breaches:
  - Firefox Monitor: How to create strong passwords
  - Firefox Monitor: Steps to protect your online identity

## G.1.2 Main Survey

### Show breach

In our previous survey, you provided an email address for querying the Have I Been Pwned (haveibeenpwned.com, referred to as “HIBP” onward) database for data breach records. Below is a data breach in which your provided email address has appeared. *[Show breach information.]*

### Breach-related questions

16. Prior to our study, have you ever heard of [site name]?  Yes  No  Unsure
17. Prior to our study, were you aware that you are affected by the [site name] data breach?  Yes  No  Unsure
18. To your knowledge, do you have an account with [site name]?  Yes, I created an account by providing an email address and a password  Yes, I created an account through a third-party service (e.g., “sign in with Google” or “sign in with Facebook”)  No, I did not have an account with [site name]
19. *[If “unsure” for Q16/Q17/Q18]* You selected unsure for [Q16/Q17/Q18]. Please explain why you selected “unsure.” [free text]
20. *[If “yes” for Q18]* To the best of your memory, how long have you been using your [site name] account? [number entry]  year(s)  month(s)  week(s)  day(s)
21. *[If “yes” for Q18]* How often do you use your [site name] account?  Every day  A few times a week  A few times a month  A few times a year or less frequently
22. *[If “yes” for Q18]* To the best of your memory, does your [company name] account have any of the following information about you? Please select all that apply.  Date of birth  Financial information (e.g., bank accounts or credit

card numbers     Gender     IP address     Name     Phone number      
Residential address

23. *[If “yes” for Q18]* Is there any other type of information the account may have about you? If you cannot think of any, please write “I don’t know.” [free text]
24. *[If “yes” for Q18]* How important or unimportant is your [site name] account to you?  Very unimportant     Unimportant     Somewhat unimportant      
Neither important nor unimportant     Somewhat important     Important      
Very important

### **Measuring password change intention**

*[Show intervention, see Figures 9.2 and 9.3 for the specific text.]*

25. After learning about this breach, do you intend to change the password of your [site name] account?  Yes, I intend to change the password of my [site name] account     No, I do not intend to change the password of my [site name] account     I have already changed the password of my [company name] account after this breach occurred (on [date]) and before taking this survey
26. Please explain why you selected this answer option. [free text]
27. Do you use your [site name] account’s password for any other online accounts?  
 Yes     No     I don’t know     I don’t have an account with [site name]

### **Measuring PMT constructs**

*[The order of questions in this section was randomized.]*

28. Please rate to what extent the following incidents would be a serious problem to you. [Answer options for each: not at all serious, slightly serious, somewhat serious, serious, extremely serious.]  Experience financial loss     Have my personal information sold to marketers     Have my online accounts hacked by someone     Have my identity stolen by someone     Receive more spam emails

- Please select “extremely serious” (this is an attention check)
29. As a result of the [site name] data breach, how likely or unlikely do you think you are to experience the following incidents? [Answer options for each: very unlikely, somewhat unlikely, neither likely nor unlikely, somewhat likely, very likely.]
- Experience financial loss
  - Have my personal information sold to marketers
  - Have my online accounts hacked by someone
  - Have my identity stolen by someone
  - Receive more spam emails
  - Please select “very unlikely” (this is an attention check)
30. Please rate your level of disagreement or agreement with the following statement: “Changing the password of my [site name] account will protect me from negative incidents as a result of the [site name] breach.”
- Strongly disagree
  - Somewhat disagree
  - Neither agree nor disagree
  - Somewhat agree
  - Strongly agree
31. How easy or difficult do you think it would be for you to change the password of your [site name] account?
- Very difficult
  - Somewhat difficult
  - Neither easy nor difficult
  - Somewhat easy
  - Very easy
32. If I were to change the password of my [site name] account, I would... [Answer options for each: not true at all, slightly true, somewhat true, true, extremely true.]
- Spend a lot of time changing the password
  - Changing the password requires me to learn new skills
  - Feel more anxious about the [site name] data breach
  - Forget the new password and be locked out of the account

### **SA-6 & prior negative experience**

You’re almost done! Just a few questions about yourself.

33. Please indicate your agreement with the following statements. [Answer options for each: strongly disagree, somewhat disagree, neither agree nor disagree, somewhat agree, strongly agree.]
- Generally, I diligently follow a routine about



security practices.     I always pay attention to experts' advice about the steps I need to take to keep my online data and accounts safe.     I am extremely knowledgeable about all the steps needed to keep my online data and accounts safe.     I am extremely motivated to take all the steps needed to keep my online data and accounts safe.     I often am interested in articles about security threats.     I seek out opportunities to learn about security measures that are relevant to me.

34. Has anyone ever gained unauthorized access to one of your online accounts? E.g., someone secretly changed your password without you noticing it.     Yes     No     Unsure
35. Have you ever learned that your information was exposed in a data breach before taking our survey?     Yes     No     Unsure
36. Have you ever been a victim of identity theft? E.g., someone secretly applied for a new credit card under your name.     Yes     No     Unsure
37. *[If "unsure" for Q34/Q35/Q36]* You selected unsure for [Q34/Q35/Q36]. Please explain why you selected "unsure." [free text]

### **Final remarks**

Thank you for completing this survey! Please note that the information about this data breach we showed to you is real. Your email address and potentially other personal information has appeared in this breach and could be used by criminals to steal your identity or access your online accounts.

You can check a full list of data breaches associated with your email address on [haveibeenpwned.com](https://haveibeenpwned.com), which, in addition, provides records with sensitive breaches upon the verification of your email account. Please keep in mind that this list only reflects breaches that are registered in the [haveibeenpwned.com](https://haveibeenpwned.com) database, your information may have been exposed in other breaches.

Please click the “continue” button to proceed to the next page, where you will be automatically redirected to Prolific.

### **G.1.3 Follow-up Survey**

#### **Reminder of breach**

In our previous survey, you provided an email address for querying the Have I Been Pwned (haveibeenpwned.com, referred to as “HIBP” onward) database for data breach records.

Below is a data breach in which your provided email address has appeared.

*[Show breach information.]*

38. Since taking our previous survey on [date], what did you do, if anything, after learning that your information was exposed in the [site name] data breach? Please explain why. [free text]

#### **Attention check**

39. This is an attention check. What is the name of the company that suffered a data breach and exposed your information, as we show you in the previous page?  correct answer  AKP emails  KnownCircle  Staminus

#### **Measuring password change behavior**

40. You took our previous survey on [date]. The survey showed that your information was exposed in the [site name] data breach. Since then, have you changed the password for your [site name] account?  Yes  No
41. Please explain why you changed or did not change your [site name] account’s password after taking our previous survey on [date]. [free text]
42. What did you do about passwords for other accounts?  I changed the password for every online account I have  I changed the password for other accounts

that use the same or similar passwords    ◦ I changed the password for really important accounts (e.g., bank accounts)    ◦ I kept using the same password for other accounts

43. *[If “Yes” for Q40]* To your best estimate, how soon after taking our previous survey on [date] did you change the password for your [site name] account? [number entry] days
44. *[If “Yes” for Q40]* What did you use for the new password for your [site name] account?    ◦ A password that I already use for other accounts    ◦ Something related to the old password but a few characters different, created by myself    ◦ Something completely unrelated to the old password, created by myself    ◦ A unique or random password, such as one generated by a password manager    ◦ Other: [free text]
45. *[If “Yes” for Q40]* What techniques did you use to remember the new password for your [site name] account? Please select all that apply.    ◦ I remembered my new password without writing it down or storing it digitally    ◦ I reset my password every time I log in rather than remembering my new password    ◦ I wrote my new password down on paper or other physical media    ◦ I stored my new password in a digital file or files    ◦ I saved my new password in the browser (e.g., passwords saved in Chrome)    ◦ I used a system-provided password manager (e.g., Apple’s Keychain)    ◦ I used a third-party password manager (e.g., 1Password or LastPass)    ◦ Other: [free text]

### **Screenshot upload**

*[Only display this page if “yes” for Q40.]*

Optional: Please upload a screenshot of the password reset confirmation email you received from [site name] after changing the password. Make sure the image you are uploading is in PNG format. We would greatly appreciate it if you do this, as it

will help us validate your responses.

If you upload a valid screenshot upon our verification, you will receive a \$1.00 bonus payment in addition to the \$1.20 base payment.

Steps to take:

- Sign in to your email account; make sure this account has the email address you checked for breaches in our earlier survey.
  - Do a keyword search of “[site name]” in your inbox and/or spam folder.
  - On a Windows PC, open the “Snipping Tool” program; on a Mac computer, press Shift + Command + 4 on your keyboard to take a screenshot.
  - Select the area you want to take a screenshot of. Make sure your screenshot includes the email’s subject line, sender, and date (see the example below).
  - **IMPORTANT: Do NOT upload a screenshot that includes your personal information, such as your actual email address, username, and new password.**
46. Which of the following options applies to you?     I have my screenshot and I am ready to upload it on the next page.     I cannot find the password reset confirmation email (please explain why): [free text]     I choose not to upload the screenshot (please explain why): [free text]

### **Final remarks**

Thank you for your participation! As we noted in our previous surveys, the information about the data breach we showed to you is real. Your email address and potentially other personal information has appeared in this breach and could be used by criminals to steal your identity or access your online accounts.

Below is the full list of breaches associated with the email address you provided. Please note that you can always obtain the same results by checking your email address on [haveibeenpwned.com](https://haveibeenpwned.com), which, in addition, provides records with sensitive

breaches upon the verification of your email account. Please keep in mind that this list only reflects breaches that are registered in the [haveibeenpwned.com](https://haveibeenpwned.com) database, your information may have been exposed in other breaches.

*[Show all breaches.]*

- Resources about recovering from a data breach:
  - Federal Trade Commission: Identity theft recovery steps
  - Federal Trade Commission: Credit Freeze FAQs
  - Firefox Monitor: What to do after a data breach
  - Norton: What to do after 5 types of data breaches
- Resources about protecting yourself against future breaches:
  - Firefox Monitor: How to create strong passwords
  - Firefox Monitor: Steps to protect your online identity

We greatly appreciate the insights you contributed, which would help us develop better technologies and interfaces that notify people of data breaches.

Would you like to be contacted by us to participate in our future research?  Yes, I am interested.  No, I am not interested.

Please click the “continue” button to proceed to the next page, where you will be automatically redirected to Prolific.

## G.2 Cognitive Walkthrough Protocol

### Introduction

Today we will be testing a series of three surveys for an experiment. To ensure you have a fresh experience, I will not reveal the experiment’s purpose, but I am happy to talk about it toward the end.

The first survey is a screening survey. You will be asked to provide an email address for querying a database and see if this email address has appeared in any

data breaches.

The second survey is the main survey. You will see more information about a breach associated with the email address you provided, and answer some questions about it.

The third survey is a follow-up survey. For real-world participants, they will receive a link to this survey two weeks after the second survey. Since we are doing pilot testing today, I will ask you to review the follow-up survey as well.

The link to all surveys is [URL]. If possible, could you please share your screen with me as you complete the surveys? We'll be using a think-aloud protocol, which means you will read out the survey questions on your screen line by line, tell me your answer and why you select it. Please also feel free to comment on your thoughts and reactions to the survey questions, especially if there's anything that's unclear or doesn't make sense to you.

For most of the time, I will be quietly sitting in the background and observe your interactions. I may chime in if there is a critical question I want to ask on the spot. I also have a list of questions I would like to get your feedback on toward the end after you complete all three surveys.

I will be taking notes as I observe your interactions. However, to make sure I capture everything, would you mind me recording our meeting today as well?

### **Questions to ask**

- Does the text under “What are the risks” describe the threats clearly? Does the text under “How to change your password” provide useful information? Do you feel motivated to change an exposed password after reading the text?
- Right now there's a 2-second delay after “What to do” and a 10-second delay after the threat/coping module. The delay seeks to nudge participants to pay close attention to the text while waiting. Is the delay too long, too short, or

just the right amount for you?

- For the follow-up survey, how's your experience of completing the screenshot upload question? Are the instructions clear? Is this question in line with ethical data collection? Are they written in a way that minimizes accidental information leaks from taking screenshots of emails?
- Any unclear or confusing wording for any survey questions?

## BIBLIOGRAPHY



## BIBLIOGRAPHY

- [1] Lillian Ablon, Paul Heaton, Diana Catherine Lavery, and Sasha Romanosky. Consumer attitudes toward data breach notifications and loss of personal information. Technical report, Rand Corporation, 2016. [https://www.rand.org/content/dam/rand/pubs/research\\_reports/RR1100/RR1187/RAND\\_RR1187.pdf](https://www.rand.org/content/dam/rand/pubs/research_reports/RR1100/RR1187/RAND_RR1187.pdf).
- [2] Ruba Abu-Salma, M Angela Sasse, Joseph Bonneau, Anastasia Danilova, Alena Naiakshina, and Matthew Smith. Obstacles to the adoption of secure communication tools. In *IEEE Symposium on Security and Privacy*, pages 137–153, 2017. doi:10.1109/SP.2017.65.
- [3] Substance Abuse and Mental Health Services Administration. Trauma training for criminal justice professionals, 2020. <https://www.samhsa.gov/gains-center/trauma-training-criminal-justice-professionals>.
- [4] Alessandro Acquisti, Idris Adjerid, Rebecca Balebako, Laura Brandimarte, Lorie Faith Cranor, Saranga Komanduri, Pedro Giovanni Leon, Norman Sadeh, Florian Schaub, Manya Sleeper, Yang Wang, and Shomir Wilson. Nudges for privacy and security: Understanding and assisting users’ choices online. *ACM Computing Surveys*, 50(3):44:1–44:41, 2017. doi:10.1145/3054926.
- [5] Alessandro Acquisti, Laura Brandimarte, and George Loewenstein. Privacy and human behavior in the age of information. *Science*, 347(6221):509–514, 2015. doi:10.1126/science.aaa1465.
- [6] Alessandro Acquisti, Laura Brandimarte, and George Loewenstein. Reasons for both pessimism and optimism: A response to the commentaries. *Journal of Consumer Psychology*, 30(4):780–783, 2020. doi:10.1002/jcpy.1187.
- [7] Alessandro Acquisti, Laura Brandimarte, and George Loewenstein. Secrets and likes: The drive for privacy and the difficulty of achieving it in the digital age. *Journal of Consumer Psychology*, 30(4):736–758, 2020. doi:10.1002/jcpy.1191.
- [8] Alessandro Acquisti and Christina Fong. An experiment in hiring discrimination via online social networks. *Management Science*, 66(3):1005–1024, 2020. doi:10.1287/mnsc.2018.3269.

- [9] Alessandro Acquisti and Jens Grossklags. Privacy and rationality in individual decision making. *IEEE Security & Privacy*, 3(1):26–33, 2005. doi:10.1109/MSP.2005.22.
- [10] Alessandro Acquisti, Leslie K John, and George Loewenstein. The impact of relative standards on the propensity to disclose. *Journal of Marketing Research*, 49(2):160–174, 2012. doi:10.1509/jmr.09.0215.
- [11] Alessandro Acquisti, Leslie K John, and George Loewenstein. What is privacy worth? *The Journal of Legal Studies*, 42(2):249–274, 2013. <https://www.journals.uchicago.edu/doi/pdf/10.1086/671754>.
- [12] Idris Adjerid, Alessandro Acquisti, Laura Brandimarte, and George Loewenstein. Sleights of privacy: Framing, disclosures, and the limits of transparency. In *Symposium on Usable Privacy and Security*, pages 9:1–9:11, 2013. doi:10.1145/2501604.2501613.
- [13] Idris Adjerid, Alessandro Acquisti, and George Loewenstein. Choice architecture, framing, and cascaded privacy choices. *Management Science*, 65(5):2267–2290, 2019. doi:10.1287/mnsc.2018.3028.
- [14] Icek Ajzen. From intentions to actions: A theory of planned behavior. In *Action Control*, pages 11–39. Springer, 1985. doi:10.1007/978-3-642-69746-3\_2.
- [15] Icek Ajzen. The theory of planned behavior. *Organizational Behavior and Human Decision Processes*, 50(2):179–211, 1991. doi:10.1016/0749-5978(91)90020-T.
- [16] Devdatta Akhawe and Adrienne Porter Felt. Alice in warningland: A large-scale field study of browser security warning effectiveness. In *USENIX Security Symposium*, pages 257–272, 2013. [https://www.usenix.org/system/files/conference/usenixsecurity13/sec13-paper\\_akhawe.pdf](https://www.usenix.org/system/files/conference/usenixsecurity13/sec13-paper_akhawe.pdf).
- [17] Elham Al Qahtani, Mohamed Shehab, and Abrar Aljohani. The effectiveness of fear appeals in increasing smartphone locking behavior among saudi arabians. In *Symposium on Usable Privacy and Security*, pages 31–46, 2018. <https://www.usenix.org/system/files/conference/soups2018/soups2018-qahtani.pdf>.
- [18] Yusuf Albayram, Mohammad Maifi Hasan Khan, Theodore Jensen, and Nhan Nguyen. “... better to use a lock screen than to worry about saving a few seconds of time”: Effect of Fear Appeal in the Context of Smartphone Locking Behavior. In *Symposium on Usable Privacy and Security*, pages 49–63, 2017. <https://www.usenix.org/system/files/conference/soups2017/soups2017-albayram.pdf>.
- [19] Nora Alkaldi and Karen Renaud. Why do people adopt, or reject, smartphone password managers? In *European Workshop on Usable Security*, pages 5:1–5:14, 2016. doi:10.14722/eurosec.2016.23011.

- [20] Nora Alkaldi, Karen Renaud, and Lewis Mackenzie. Encouraging password manager adoption by meeting adopter self-determination needs. In *Hawaii International Conference on System Sciences*, pages 4824–4833, 2019. doi:10.24251/HICSS.2019.582.
- [21] Hazim Almuhiemedi, Florian Schaub, Norman Sadeh, Idris Adjerid, Alessandro Acquisti, Joshua Gluck, Lorrie Faith Cranor, and Yuvraj Agarwal. Your location has been shared 5,398 times! A field study on mobile app privacy nudging. In *ACM Conference on Human Factors in Computing Systems*, pages 787–796, 2015. doi:10.1145/2702123.2702210.
- [22] Irwin Altman. *The environment and social behavior: privacy, personal space, territory, and crowding*. ERIC, 1975.
- [23] James Amirkhan. A factor analytically derived measure of coping: The coping strategy indicator. *Journal of Personality and Social Psychology*, 59(5):1066–1074, 1990. doi:10.1037/0022-3514.59.5.1066.
- [24] Catherine L Anderson and Ritu Agarwal. Practicing safe computing: A multimethod empirical examination of home computer user security behavioral intentions. *MIS Quarterly*, 34(3):613–643, 2010. doi:10.2307/25750694.
- [25] Catherine L Anderson and Ritu Agarwal. The digitization of health-care: boundary risks, emotion, and consumer willingness to disclose personal health information. *Information Systems Research*, 22(3):469–490, 2011. doi:10.1287/isre.1100.0335.
- [26] J. Craig Anderson. Identity theft growing, costly to victims, 2013. <https://www.usatoday.com/story/money/personalfinance/2013/04/14/identity-theft-growing/2082179/>.
- [27] Edmund L Andrews. How flawed data aggravates inequality in credit, 2021. <https://hai.stanford.edu/news/how-flawed-data-aggravates-inequality-credit>.
- [28] Julio Angulo and Martin Ortlieb. “WTH..!?” Experiences, reactions, and expectations related to online privacy panic situations. In *Symposium On Usable Privacy and Security*, pages 19–38, 2015. <https://www.usenix.org/system/files/conference/soups2015/soups15-paper-angulo.pdf>.
- [29] Annie I Antón, Julia Brande Earp, Qingfeng He, William Stufflebeam, Davide Bolchini, and Carlos Jensen. Financial privacy policies and the need for standardization. *IEEE Security & Privacy*, 2(2):36–45, 2004. doi:10.1109/MSECP.2004.1281243.
- [30] Apple Inc. What is sign in with apple?, 2021. <https://support.apple.com/en-us/HT210318>.

- [31] Budi Arief, Kovila PL Coopamootoo, Martin Emms, and Aad van Moorsel. Sensible privacy: How we can protect domestic violence survivors without facilitating misuse. In *Workshop on Privacy in the Electronic Society*, pages 201–204, 2014. doi:10.1145/2665943.2665965.
- [32] Richard A Armstrong. When to use the bonferroni correction. *Ophthalmic and Physiological Optics*, 34(5):502–508, 2014. doi:10.1111/opo.12131.
- [33] Brooke Auxier, Lee Rainie, Monica Anderson, Andrew Perrin, Madhu Kumar, and Erica Turner. Americans and privacy: Concerned, confused and feeling lack of control over their personal information. Technical report, Pew Research Center, 2019. [https://www.pewresearch.org/internet/wp-content/uploads/sites/9/2019/11/Pew-Research-Center\\_PI\\_2019.11.15\\_Privacy\\_FINAL.pdf](https://www.pewresearch.org/internet/wp-content/uploads/sites/9/2019/11/Pew-Research-Center_PI_2019.11.15_Privacy_FINAL.pdf).
- [34] Michael F Baber. *Integrated Business Leadership Through Cross Marketing*. Warren H. Green, 1986.
- [35] Babich, Nick. Ux design: Checkbox and toggle in forms, 2016. <https://uxplanet.org/checkbox-and-toggle-in-forms-f0de6086ac41>.
- [36] Young Min Baek, Eun-mee Kim, and Young Bae. My privacy is okay, but theirs is endangered: Why comparative optimism matters in online privacy concerns. *Computers in Human Behavior*, 31:48–56, 2014. doi:10.1016/j.chb.2013.10.010.
- [37] Rebecca Balebako, Florian Schaub, Idris Adjerid, Alessandro Acquisti, and Lorie Cranor. The impact of timing on the salience of smartphone app privacy notices. In *Workshop on Security and Privacy in Smartphones and Mobile Devices*, pages 63–74, 2015. doi:10.1145/2808117.2808119.
- [38] Albert Bandura. Self-efficacy: toward a unifying theory of behavioral change. *Psychological Review*, 84(2):191–215, 1977. doi:10.1037/0033-295X.84.2.191.
- [39] Vinayshekhar Bannihatti Kumar, Roger Iyengar, Namita Nisal, Yuanyuan Feng, Hana Habib, Peter Story, Sushain Cherivirala, Margaret Hagan, Lorie Faith Cranor, Shomir Wilson, Florian Schaub, and Norman Sadeh. Finding a choice in a haystack: Automatic extraction of opt-out statements from privacy policy text. In *The Web Conference*, pages 1943–1954, 2020. doi:10.1145/3366423.3380262.
- [40] Gaurav Bansal. Distinguishing between privacy and security concerns: An empirical examination and scale validation. *Journal of Computer Information Systems*, 57(4):330–343, 2017. doi:10.1080/08874417.2016.1232981.
- [41] Susanne Barth and Menno DT De Jong. The privacy paradox—investigating discrepancies between expressed privacy concerns and actual online behavior—a systematic literature review. *Telematics and Informatics*, 34(7):1038–1058, 2017. doi:10.1016/j.tele.2017.04.013.

- [42] Rosanna Bellini, Alexander Wilson, and Jan David Smeddinck. Fragments of the past: curating peer support with perpetrators of domestic violence. In *ACM Conference on Human Factors in Computing Systems*, pages 708:1–708:14, 2021. doi:10.1145/3411764.3445611.
- [43] Buster Benson. Cognitive bias cheat sheet, 2016. <https://medium.com/better-humans/cognitive-bias-cheat-sheet-55a472476b18>.
- [44] Adam J Berinsky, Gregory A Huber, and Gabriel S Lenz. Evaluating online labor markets for experimental research: Amazon.com’s mechanical turk. *Political Analysis*, 20(3):351–368, 2012. doi:10.1093/pan/mpr057.
- [45] Andrew Besmer, Jason Watson, and Heather Richter Lipford. The impact of social navigation on privacy policy configuration. In *Symposium on Usable Privacy and Security*, pages 7:1–7:10, 2010. doi:10.1145/1837110.1837120.
- [46] Sruti Bhagavatula, Lujo Bauer, and Apu Kapadia. (how) do people change their passwords after a breach? In *USENIX ;login.*, pages 1–7, 2020. <https://www.usenix.org/sites/default/files/password-change-after-breach.pdf>.
- [47] Jaspreet Bhatia, Travis D Breaux, Joel R Reidenberg, and Thomas B Norton. A theory of vagueness and privacy risk perception. In *IEEE International Requirements Engineering Conference*, pages 26–35, 2016. doi:10.1109/RE.2016.20.
- [48] Fabio Bisogni. Proving limits of state data breach notification laws: Is a federal law the most adequate solution? *Journal of Information Policy*, 6(1):154–205, 2016. doi:10.5325/jinfopoli.6.2016.0154.
- [49] Jan Boehmer, Robert LaRose, Nora Rifon, Saleem Alhabash, and Shelia Cotten. Determinants of online safety behaviour: Towards an intervention strategy for college students. *Behaviour & Information Technology*, 34(10):1022–1035, 2015. doi:10.1080/0144929X.2015.1028448.
- [50] Christoph Bösch, Benjamin Erb, Frank Kargl, Henning Kopp, and Stefan Pfattheicher. Tales from the dark side: Privacy dark strategies and privacy dark patterns. *Proceedings on Privacy Enhancing Technologies*, 2016(4):237–254, 2016. doi:10.1515/popets-2016-0038.
- [51] Scott R Boss, Dennis F Galletta, Paul Benjamin Lowry, Gregory D Moody, and Peter Polak. What do systems users have to fear? using fear appeals to engender threats and fear that motivate protective security behaviors. *MIS Quarterly*, 39(4):837–864, 2015. <https://www.jstor.org/stable/26628654>.
- [52] Laura Brandimarte, Alessandro Acquisti, and George Loewenstein. Misplaced confidences: Privacy and the control paradox. *Social Psychological and Personality Science*, 4(3):340–347, 2013. doi:10.1177/1948550612455931.

- [53] Petter Bae Brandtzæg, Marika Lüders, and Jan Håvard Skjetne. Too many Facebook “friends”? Content sharing and sociability versus the need for privacy in social network sites. *International Journal of Human–Computer Interaction*, 26(11-12):1006–1030, 2010. doi:10.1080/10447318.2010.516719.
- [54] Cristian Bravo-Lillo, Lorrie Faith Cranor, Julie Downs, and Saranga Komanduri. Bridging the gap in computer security warnings: A mental model approach. *IEEE Security & Privacy*, 9(2):18–26, 2011. doi:10.1109/MSP.2010.198.
- [55] Patrick Brown. A nudge in the right direction? towards a sociological engagement with libertarian paternalism. *Social policy and society*, 11(3):305–317, 2012. doi:10.1017/S1474746412000061.
- [56] Scott Brown. Did you forget to reply to an email? the new gmail will remind you, 2018. <https://www.androidauthority.com/gmail-nudges-feature-865435/>.
- [57] AJ Bernheim Brush and Kori M Inkpen. Yours, mine and ours? Sharing and use of technology in domestic environments. In *International Conference on Ubiquitous Computing*, pages 109–126, 2007. doi:10.1007/978-3-540-74853-3\_7.
- [58] Marc Brysbaert and Michaël Stevens. Power analysis and effect size in mixed effects models: A tutorial. *Journal of Cognition*, 1(1):1–20, 2018. doi:10.5334/joc.10.
- [59] Tom Buchanan, Carina Paine, Adam N Joinson, and Ulf-Dietrich Reips. Development of measures of online privacy concern and protection for use on the internet. *Journal of the American Society for Information Science and Technology*, 58(2):157–165, 2007. doi:10.1002/asi.20459.
- [60] Buffalo Center for Social Research. What is trauma-informed care, 2020. <https://socialwork.buffalo.edu/social-research/institutes-centers/institute-on-trauma-and-trauma-informed-care/what-is-trauma-informed-care.html>.
- [61] Daniel Bühler, Fabian Hemmert, and Jörn Hurtienne. Universal and intuitive? scientific guidelines for icon design. In *ACM Conference on Mensch und Computer*, pages 91–103, 2020. doi:10.1145/3404983.3405518.
- [62] Mohamad Adam Bujang, Nadiah Sa’at, Tg Mohd Ikhwan Tg Abu Bakar, et al. Sample size guidelines for logistic regression from observational studies with large population. *The Malaysian Journal of Medical Sciences*, 25(4):122–130, 2018. doi:10.21315/mjms2018.25.4.12.
- [63] AJ Burns, Clay Posey, Tom L Roberts, and Paul Benjamin Lowry. Examining the relationship of organizational insiders’ psychological capital with information security threat and coping appraisals. *Computers in Human Behavior*, 68:190–209, 2017. doi:10.1016/j.chb.2016.11.018.

- [64] Karoline Busse, Julia Schäfer, and Matthew Smith. Replication: no one can hack my mind revisiting a study on expert and non-expert security practices and advice. In *Symposium on Usable Privacy and Security*, pages 117–136, 2019. <https://www.usenix.org/system/files/soups2019-busse.pdf>.
- [65] John T Cacioppo and Wendi L Gardner. Emotion. *Annual Review of Psychology*, 50(1):191–214, 1999. doi:10.1146/annurev.psych.50.1.191.
- [66] Kelly Caine. *Exploring everyday privacy behaviors and disclosures*. PhD thesis, Georgia Institute of Technology, 2009. <https://smartech.gatech.edu/handle/1853/31665>.
- [67] Caroline Cakebread. You’re not alone, no one reads terms of service agreements. Business Insider, 2017. <https://www.businessinsider.com/deloitte-study-91-percent-agree-terms-of-service-without-reading-2017-11>.
- [68] Ryan Calo. Code, nudge, or notice. *Iowa Law Review*, 99(2):773–802, 2013. <https://heinonline.org/HOL/LandingPage?handle=hein.journals/ilr99&div=22>.
- [69] L Jean Camp. Mental models of privacy and security. *IEEE Technology and society magazine*, 28(3):37–46, 2009. doi:10.1109/MTS.2009.934142.
- [70] Xavier De Carné De Carnavalet and Mohammad Mannan. A large-scale evaluation of high-impact password strength meters. *ACM Transactions on Information and System Security*, 18(1):1–32, 2015. doi:10.1145/2739044.
- [71] Ronald P Carver. Is reading rate constant or flexible? *Reading Research Quarterly*, 18(2):190–215, 1983. doi:10.2307/747517.
- [72] Fred H Cate. The limits of notice and choice. *IEEE Security & Privacy*, 8(2):59–62, 2010. doi:10.1109/MSP.2010.84.
- [73] Ann Cavoukian. Privacy by design: The 7 foundational principles, 2009. <https://dataprotection.industries/wp-content/uploads/2017/10/privacy-by-design.pdf>.
- [74] Centers for Disease Control and Prevention. Intimate partner violence, 2018. <https://www.cdc.gov/violenceprevention/intimatepartnerviolence/index.html>.
- [75] Daphne Chang, Erin L Krupka, Eytan Adar, and Alessandro Acquisti. Engineering information disclosure: Norm shaping designs. In *ACM Conference on Human Factors in Computing Systems*, pages 587–597, 2016. doi:10.1145/2858036.2858346.
- [76] Rahul Chatterjee, Periwinkle Doerfler, Hadas Orgad, Sam Havron, Jackeline Palmer, Diana Freed, Karen Levy, Nicola Dell, Damon McCoy, and Thomas Ristenpart. The spyware used in intimate partner violence. In *IEEE Symposium on Security and Privacy*, pages 441–458, 2018. doi:10.1109/SP.2018.00061.

- [77] Angela Chen. Most Americans are fine with cops using facial recognition on them, 2019. <https://www.technologyreview.com/f/614267/facial-recognition-police-law-enforcement-surveillance-privacy-pew-research-survey/>.
- [78] Henian Chen, Patricia Cohen, and Sophie Chen. How big is a big odds ratio? interpreting the magnitudes of odds ratios in epidemiological studies. *Communications in Statistics—simulation and Computation*, 39(4):860–864, 2010. doi:10.1080/03610911003650383.
- [79] Hsuan-Ting Chen and Wenhong Chen. Couldn’t or wouldn’t? The influence of privacy concerns and self-efficacy in privacy management on privacy protection. *Cyberpsychology, Behavior, and Social Networking*, 18(1):13–19, 2015. doi:10.1089/cyber.2014.0456.
- [80] Janet X Chen, Allison McDonald, Yixin Zou, Emily Tseng, Kevin A Roundy, Acar Tamersoy, Florian Schaub, Thomas Ristenpart, and Nicola Dell. Trauma-informed computing: Towards safer technology experiences for all. In *ACM Conference on Human Factors in Computing Systems*, pages 544:1–544:20, 2022. doi:10.1145/3491102.3517475.
- [81] Hichang Cho, Jae-Shin Lee, and Siyoung Chung. Optimistic bias about on-line privacy risks: Testing the moderating effects of perceived controllability and prior experience. *Computers in Human Behavior*, 26(5):987–995, 2010. doi:10.1016/j.chb.2010.02.012.
- [82] Shin-Ho Chung and Richard J Herrnstein. Choice and delay of reinforcement. *Journal of the Experimental Analysis of Behavior*, 10(1):67–74, 1967. doi:10.1901/jeab.1967.10-67.
- [83] Danielle Keats Citron and Daniel J Solove. Privacy harms. *Boston University Law Review*, 102:793–863, 2022. <https://heinonline.org/HOL/LandingPage?handle=hein.journals/bulr102&div=20>.
- [84] Clinic to End Tech Abuse. Homepage, 2021. <https://www.ceta.tech.cornell.edu/>.
- [85] Jacob Cohen. *Statistical Power Analysis for the Behavioral Sciences*. Academic press, 2013.
- [86] Roger Coleman, John Clarkson, and Julia Cassim. *Design for inclusivity: A practical guide to accessible, innovative and user-centred design*. CRC Press, 2016.
- [87] Jessica Colnago, Lorrie Faith Cranor, Alessandro Acquisti, and Kate Hazel Jain. Is it a concern or a preference? an investigation into the ability of privacy scales to capture and distinguish granular privacy constructs. In *Symposium on Usable Privacy and Security*, pages 331–346, 2022. <https://www.usenix.org/system/files/soups2022-colnago.pdf>.



- [88] Jessica Colnago, Yuanyuan Feng, Tharangini Palanivel, Sarah Pearman, Megan Ung, Alessandro Acquisti, Lorrie Faith Cranor, and Norman Sadeh. Informing the design of a personalized privacy assistant for the internet of things. In *ACM Conference on Human Factors in Computing Systems*, pages 262:1–262:13, 2020. doi:10.1145/3313831.3376389.
- [89] Colorado Coalition Against Sexual Assault. Sexual assault advocacy & crisis line training guide, 2011. <https://www.ccasa.org/wp-content/uploads/2014/01/Sexual-Assault-Advocacy-and-Crisis-Line-Training-Guide.pdf>.
- [90] The Federal Trade Commission. Fraud alerts & credit freezes: What’s the difference?, 2020. <https://www.consumer.ftc.gov/blog/2020/02/fraud-alerts-credit-freezes-whats-difference>.
- [91] Nathan S Consedine, Shulamit Sabag-Cohen, and Yulia S Krivoshekova. Ethnic, gender, and socioeconomic differences in young adults’ self-disclosure: Who discloses what and to whom? *Cultural Diversity and Ethnic Minority Psychology*, 13(3):254–263, 2007. doi:10.1037/1099-9809.13.3.254.
- [92] Consumer Financial Protection Bureau. Measuring financial well-being: A guide to using the CFPB Financial Well-Being Scale, 2015. <https://www.consumerfinance.gov/data-research/research-reports/financial-well-being-scale/>.
- [93] Kovila PL Coopamootoo. Usage patterns of privacy-enhancing technologies. In *ACM Conference on Computer and Communications Security*, pages 1371–1390, 2020. doi:10.1145/3372297.3423347.
- [94] Veracities Public Benefit Corporation. Certified FETI | the official forensic experiential trauma interview, 2021. <https://www.certifiedfeti.com/>.
- [95] Michelle Cottle. The adultery arms race, 2014. <https://www.theatlantic.com/magazine/archive/2014/11/the-adultery-arms-race/380794/>.
- [96] Claudia Coulton and Julian Chow. Interaction effects in multiple regression. *Journal of Social Service Research*, 16(1-2):179–199, 1993. doi:10.1300/J079v16n01\_09.
- [97] Joseph Cox. Security companies and activists launch ‘coalition against stalkerware’, 2019. [https://www.vice.com/en\\_us/article/ywa7xv/coalition-against-stalkerware-launches-eff-kaspersky](https://www.vice.com/en_us/article/ywa7xv/coalition-against-stalkerware-launches-eff-kaspersky).
- [98] William D Crano and Radmila Prislin. Attitudes and persuasion. *Annual Review of Psychology*, 57:345–374, 2006. doi:10.1146/annurev.psych.57.102904.190034.
- [99] Lorrie Cranor. Time to rethink mandatory password changes, 2016. <https://www.ftc.gov/policy/advocacy-research/tech-at-ftc/2016/03/time-rethink-mandatory-password-changes>.

- [100] Lorrie Faith Cranor. Necessary but not sufficient: Standardized mechanisms for privacy notice and choice. *Journal on Telecommunications and High Technology Law*, 10:273–308, 2012. <https://heinonline.org/HOL/LandingPage?handle=hein.journals/jtelhtel10&div=22>.
- [101] Kimberlé Crenshaw. Demarginalizing the intersection of race and sex: A black feminist critique of antidiscrimination doctrine, feminist theory, and antiracist politic. In *Feminist Legal Theories*, pages 57–80. Routledge, 2018. <https://www.taylorfrancis.com/chapters/edit/10.4324/9781315051536-2/demarginalizing-intersection-race-sex-kimberle-crenshaw>.
- [102] Robert E Crossler, James H Long, Tina M Loraas, and Brad S Trinkle. Understanding compliance with bring your own device policies utilizing protection motivation theory: Bridging the intention-behavior gap. *Journal of Information Systems*, 28(1):209–226, 2014. doi:10.2308/isys-50704.
- [103] Dana Cuomo and Natalie Dolci. Gender-based violence and technology-enabled coercive control in seattle: Challenges & opportunities. Technical report, Technology-Enabled Coercive Control Working Group, Whitepaper Series, 2019. <https://teccworkinggroup.org/research-2>.
- [104] Paul Curzon and Ann Blandford. Formally justifying user-centred design rules: a case study on post-completion errors. In *International Conference on Integrated Formal Methods*, pages 461–480, 2004. doi:10.1007/978-3-540-24756-2\_25.
- [105] Anupam Das, Joseph Bonneau, Matthew Caesar, Nikita Borisov, and XiaoFeng Wang. The tangled web of password reuse. In *Network and Distributed System Security Symposium*, pages 26:1–26:15, 2014. doi:10.14722/NDSS.2014.23357.
- [106] Sanchari Das, Andrew Dingman, and L Jean Camp. Why Johnny Doesn’t Use Two Factor A Two-Phase Usability Study of the FIDO U2F Security Key. In *International Conference on Financial Cryptography and Data Security*, pages 160–179, 2018. doi:10.1007/978-3-662-58387-6\_9.
- [107] Sauvik Das, Laura A Dabbish, and Jason I Hong. A typology of perceived triggers for end-user security and privacy behaviors. In *Symposium on Usable Privacy and Security*, pages 97–115, 2019. <https://www.usenix.org/system/files/soups2019-das.pdf>.
- [108] Sauvik Das, Tiffany Hyun-Jin Kim, Laura A Dabbish, and Jason I Hong. The effect of social influence on security sensitivity. In *Symposium on Usable Privacy and Security*, pages 143–157, 2014. <https://www.usenix.org/system/files/conference/soups2014/soups14-paper-das.pdf>.
- [109] Sauvik Das, Adam DI Kramer, Laura A Dabbish, and Jason I Hong. The role of social influence in security feature adoption. In *ACM Conference on*

*Computer Supported Cooperative Work & Social Computing*, pages 1416–1426, 2015. doi:10.1145/2675133.2675225.

- [110] Sauvik Das, Joanne Lo, Laura Dabbish, and Jason I Hong. Breaking! A Typology of Security and Privacy News and How It’s Shared. In *ACM Conference on Human Factors in Computing Systems*, pages 1:1–1:12, 2018. doi:10.1145/3173574.3173575.
- [111] Dashlane. World Password Day: How to Improve Your Passwords, 2018. <https://blog.dashlane.com/world-password-day/>.
- [112] Fred D Davis, Richard P Bagozzi, and Paul R Warshaw. User acceptance of computer technology: A comparison of two theoretical models. *Management science*, 35(8):982–1003, 1989. doi:10.1287/mnsc.35.8.982.
- [113] Behnam Dayanim and Edward George. Data breach litigation and regulatory enforcement: A survey of our present and how to prepare for the future. *Cyber Security*, 1(4):301–315, 2018. <https://www.henrystewartpublications.com/sites/default/files/CSJDayanim.pdf>.
- [114] Joe DeBlasio, Stefan Savage, Geoffrey M Voelker, and Alex C Snoeren. Tripwire: Inferring internet site compromise. In *ACM Internet Measurement Conference*, pages 341–354, 2017. doi:10.1145/3131365.3131391.
- [115] Sergej Dechand, Alena Naiakshina, Anastasia Danilova, and Matthew Smith. In encryption we don’t trust: The effect of end-to-end encryption to the masses on user perception. In *IEEE European Symposium on Security and Privacy*, pages 401–415, 2019. doi:10.1109/EuroSP.2019.00037.
- [116] Lina Dencik and Jonathan Cable. The advent of surveillance realism: Public opinion and activist responses to the snowden leaks. *International Journal of Communication*, 11:763–781, 2017. <https://ijoc.org/index.php/ijoc/article/view/5524/1939>.
- [117] Morton Deutsch and Harold B Gerard. A study of normative and informational social influences upon individual judgment. *Journal of Abnormal and Social Psychology*, 51(3):629–636, 1955. doi:10.1037/h0046408.
- [118] Jayati Dev, Emilee Rader, and Sameer Patil. Why johnny can’t unsubscribe: Barriers to stopping unwanted email. In *ACM Conference on Human Factors in Computing Systems*, pages 38:1–38:12, 2020. doi:10.1145/3313831.3376165.
- [119] Kerry DeVito. How to improve your Watchtower score in 1Password, 2022. <https://blog.1password.com/improve-watchtower-score-1password/>.
- [120] Tobias Dienlin and Sabine Trepte. Is the privacy paradox a relic of the past? an in-depth analysis of privacy attitudes and privacy behaviors. *European journal of social psychology*, 45(3):285–297, 2015. doi:10.1002/ejsp.2049.

- [121] Tawanna R Dillahunt and Amelia R Malone. The promise of the sharing economy among disadvantaged communities. In *ACM Conference on Human Factors in Computing Systems*, pages 2285–2294, 2015. doi:10.1145/2702123.2702189.
- [122] Jill P Dimond, Casey Fiesler, and Amy S Bruckman. Domestic violence and information communication technologies. *Interacting with Computers*, 23(5):413–421, 2011. doi:10.1016/j.intcom.2011.04.006.
- [123] Tamara Dinev and Paul Hart. An extended privacy calculus model for e-commerce transactions. *Information Systems Research*, 17(1):61–80, 2006. doi:10.1287/isre.1060.0080.
- [124] Disconnect, Inc. Disconnect privacy icons, 2014. <https://github.com/disconnectme/privacy-icons>.
- [125] Matthew Dixon, Karen Freeman, and Nicholas Toman. Stop trying to delight your customers, 2010. <https://hbr.org/2010/07/stop-trying-to-delight-your-customers>.
- [126] Alexandra Dmitrienko, Christopher Liebchen, Christian Rossow, and Ahmad-Reza Sadeghi. On the (in) security of mobile two-factor authentication. In *International Conference on Financial Cryptography and Data Security*, pages 365–383, 2014. doi:10.1007/978-3-662-45472-5\_24.
- [127] Heather Douglas, Bridget A Harris, and Molly Dragiewicz. Technology-facilitated domestic and family violence: Women’s experiences. *The British Journal of Criminology*, 59(3):551–570, 2019. doi:10.1093/bjc/azy068.
- [128] Anthony Downs. *An economic theory of democracy*. Harper New York, 1957.
- [129] Nora Draper and Joseph Turow. The corporate cultivation of digital resignation. *New Media & Society*, 21(8):1824–1839, 2019. doi:10.1177/1461444819833331.
- [130] Adam Drewnowski and Bennet B Murdock. The role of auditory features in memory span for words. *Journal of Experimental Psychology: Human Learning and Memory*, 6(3):319–332, 1980. doi:10.1037/0278-7393.6.3.319.
- [131] Janna Lynn Dupree, Richard Devries, Daniel M Berry, and Edward Lank. Privacy personas: Clustering users via attitudes and behaviors toward security practices. In *ACM Conference on Human Factors in Computing Systems*, pages 5228–5239, 2016. doi:10.1145/2858036.2858214.
- [132] Marc Dupuis and Mercy Ebenezer. Help wanted: Consumer privacy behavior and smart home internet of things (IoT) devices. In *ACM International Conference on Information Technology in Education*, pages 117–122, 2018. doi:10.1145/3241815.3241869.

- [133] Serge Egelman, Lorrie Faith Cranor, and Jason Hong. You’ve been warned: an empirical study of the effectiveness of web browser phishing warnings. In *ACM Conference on Human Factors in Computing Systems*, pages 1065–1074, 2008. doi:10.1145/1357054.1357219.
- [134] Serge Egelman, Raghudeep Kannavara, and Richard Chow. Is this thing on? crowdsourcing privacy indicators for ubiquitous sensing platforms. In *ACM Conference on Human Factors in Computing Systems*, pages 1669–1678, 2015. doi:10.1145/2702123.2702251.
- [135] Serge Egelman, Andreas Sotirakopoulos, Ildar Muslukhov, Konstantin Beznosov, and Cormac Herley. Does my password go up to eleven? the impact of password meters on password selection. In *ACM Conference on Human Factors in Computing Systems*, pages 2379–2388, 2013. doi:10.1145/2470654.2481329.
- [136] Serge Egelman, Janice Tsai, Lorrie Faith Cranor, and Alessandro Acquisti. Timing is everything? the effects of timing and placement of online privacy indicators. In *ACM Conference on Human Factors in Computing Systems*, pages 319–328, 2009. doi:10.1145/1518701.1518752.
- [137] Nicole B Ellison, Jessica Vitak, Charles Steinfield, Rebecca Gray, and Cliff Lampe. Negotiating privacy concerns and social capital needs in a social media environment. In *Privacy online*, pages 19–32. Springer, 2011. doi:10.1007/978-3-642-21521-6\_3.
- [138] Pardis Emami Naeini, Martin Degeling, Lujó Bauer, Richard Chow, Lorrie Faith Cranor, Mohammad Reza Haghighat, and Heather Patterson. The influence of friends and experts on privacy decision making in IoT scenarios. *Proceedings of the ACM on Human-Computer Interaction*, 2(CSCW):48:1–48:26, 2018. doi:10.1145/3274317.
- [139] Pardis Emami-Naeini, Henry Dixon, Yuvraj Agarwal, and Lorrie Faith Cranor. Exploring how privacy and security factor into iot device purchase behavior. In *ACM Conference on Human Factors in Computing Systems*, pages 534:1–534:12, 2019. doi:10.1145/3290605.3300764.
- [140] Martin Emms, Budi Arief, and Aad van Moorsel. Electronic footprints in the sand: Technologies for assisting domestic violence survivors. In *Annual Privacy Forum*, pages 203–214. Springer, 2012. doi:10.1007/978-3-642-54069-1\_13.
- [141] Samson Esayas, Tobias Mahler, and Kevin McGillivray. Is a Picture Worth a Thousand Terms? Visualising Contract Terms and Data Protection Requirements for Cloud Computing Users. In *International Conference on Web Engineering*, pages 39–56, 2016. doi:10.1007/978-3-319-46963-8\_4.
- [142] Gabriele Esposito, Penélope Hernández, René van Bavel, and José Vila. Nudging to prevent the purchase of incompatible digital products online: An experimental study. *PloS one*, 12(3):e0173333, 2017. doi:10.1371/journal.pone.0173333.

- [143] European Data Protection Supervisor. EDPS opinion on the legislative package “A New Deal for Consumers”, 2018. [https://edps.europa.eu/sites/edp/files/publication/18-10-05\\_opinion\\_consumer\\_law\\_en.pdf](https://edps.europa.eu/sites/edp/files/publication/18-10-05_opinion_consumer_law_en.pdf).
- [144] European Parliament. Regulation (eu) 2016/679 of the european parliament and of the council, 2016. <https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:32016R0679>.
- [145] Karin M Eyrich-Garg. Sheltered in cyberspace? computer use among the unsheltered ‘street’ homeless. *Computers in Human Behavior*, 27(1):296–303, 2011. doi:10.1016/j.chb.2010.08.007.
- [146] Michael Fagan and Mohammad Maifi Hasan Khan. Why do they do what they do?: A study of what motivates users to (not) follow computer security advice. In *Symposium on Usable Privacy and Security*, pages 59–75, 2016. <https://www.usenix.org/system/files/conference/soups2016/soups2016-paper-fagan.pdf>.
- [147] Cori Faklaris, Laura A Dabbish, and Jason I Hong. A self-report measure of end-user security attitudes (sa-6). In *Symposium on Usable Privacy and Security*, pages 61–77, 2019. <https://www.usenix.org/system/files/soups2019-faklaris.pdf>.
- [148] Adrienne Porter Felt, Robert W Reeder, Alex Ainslie, Helen Harris, Max Walker, Christopher Thompson, Mustafa Embre Acer, Elisabeth Morant, and Sunny Consolvo. Rethinking connection security indicators. In *Symposium on Usable Privacy and Security*, pages 1–14, 2016. <https://www.usenix.org/system/files/conference/soups2016/soups2016-paper-porter-felt.pdf>.
- [149] Adrienne Porter Felt, Robert W Reeder, Hazim Almuhiemedi, and Sunny Consolvo. Experimenting At Scale With Google Chrome’s SSL Warning. In *ACM Conference on Human Factors in Computing Systems*, pages 2667–2670, 2014. doi:10.1145/2556288.2557292.
- [150] Yuanyuan Feng, Yaxing Yao, and Norman Sadeh. A Design Space for Privacy Chocies: Towards Meaningful Privacy Control in the Internet of Things. In *ACM Conference on Human Factors in Computing Systems*, pages 64:1–64:16, 2021. doi:10.1145/3411764.3445148.
- [151] A Finset, S Steine, L Haugli, E Steen, and E Laerum. The brief approach/avoidance coping questionnaire: Development and validation. *Psychology, Health & Medicine*, 7(1):75–85, 2002. doi:10.1080/13548500120101577.
- [152] Melissa L Finucane, Ali Alhakami, Paul Slovic, and Stephen M Johnson. The affect heuristic in judgments of risks and benefits. *Journal of Behavioral Decision Making*, 13(1):1–17, 2000. doi:10.1002/(SICI)1099-0771(200001/03)13:1<1::AID-BDM333>3.0.CO;2-S.

- [153] Simone Fischer-Hübner, Erik Wästlund, and Harald Zwingelberg. Ui prototypes: Policy administration and presentation—version 1. Technical report, Privacy and Identity Management in Europe for Life, 2009. [http://primelife.ercim.eu/images/stories/deliverables/d4.3.1-ui\\_prototypes-policy\\_administration\\_and\\_presentation\\_v1.pdf](http://primelife.ercim.eu/images/stories/deliverables/d4.3.1-ui_prototypes-policy_administration_and_presentation_v1.pdf).
- [154] Robert J Fisher. Social desirability bias and the validity of indirect questioning. *Journal of Consumer Research*, 20(2):303–315, 1993. doi:10.1086/209351.
- [155] Kim Flaherty. Top 10 design mistakes in the unsubscribe experience, 2018. <https://www.nngroup.com/articles/unsubscribe-mistakes/>.
- [156] Joseph L Fleiss, Bruce Levin, and Myunghee Cho Paik. *Statistical methods for rates and proportions*. John Wiley & Sons, 2013.
- [157] Rudolf Franz Flesch et al. *Art of readable writing*. Harper, 1949.
- [158] Dinei Florencio and Cormac Herley. A large-scale study of web password habits. In *International Conference on World Wide Web*, pages 657–666, 2007. doi:10.1145/1242572.1242661.
- [159] James Flores. What Is Social Login and Is It Worth Implementing?, 2020. <https://www.okta.com/blog/2020/08/social-login/>.
- [160] Donna L Floyd, Steven Prentice-Dunn, and Ronald W Rogers. A meta-analysis of research on protection motivation theory. *Journal of Applied Social Psychology*, 30(2):407–429, 2000. doi:10.1111/j.1559-1816.2000.tb02323.x.
- [161] Asbjørn Følstad and Marita Skjuve. Chatbots for customer service: User experience and motivation. In *ACM International Conference on Conversational User Interfaces*, pages 1:1–1:9, 2019. doi:10.1145/3342775.3342784.
- [162] Forbrukerrådet. Deceived by design: How tech companies use dark patterns to discourage us from exercising our rights to privacy. Technical report, Norwegian Consumer Council, 2018. <https://fil.forbrukerradet.no/wp-content/uploads/2018/06/2018-06-27-deceived-by-design-final.pdf>.
- [163] Alain Forget, Sarah Pearman, Jeremy Thomas, Alessandro Acquisti, Nicolas Christin, Lorrie Faith Cranor, Serge Egelman, Marian Harbach, and Rahul Telang. Do or do not, there is no try: user engagement may not improve security outcomes. In *Symposium on Usable Privacy and Security*, pages 97–111, 2016. <https://www.usenix.org/system/files/conference/soups2016/soups2016-paper-forget.pdf>.
- [164] Lorenzo Franceschi-Bicchierai. Kaspersky Lab Will Now Alert Users to ‘Stalkerware’ Used In Domestic Abuse, 2019. [https://www.vice.com/en\\_us/article/vbw9g8/kaspersky-lab-alert-stalkerware-domestic-abuse](https://www.vice.com/en_us/article/vbw9g8/kaspersky-lab-alert-stalkerware-domestic-abuse).

- [165] Cynthia Fraser, Erica Olsen, Kaofeng Lee, Cindy Southworth, and Sarah Tucker. The new age of stalking: Technological implications for stalking. *Juvenile and Family Court Journal*, 61(4):39–55, 2010. doi:10.1111/j.1755-6988.2010.01051.x.
- [166] Diana Freed, Sam Havron, Emily Tseng, Andrea Gallardo, Rahul Chatterjee, Thomas Ristenpart, and Nicola Dell. “Is my phone hacked?” Analyzing clinical computer security interventions with survivors of intimate partner violence. *Proceedings of the ACM on Human-Computer Interaction*, 3(CSCW):202:1–202:24, 2019. doi:10.1145/3359304.
- [167] Diana Freed, Jackeline Palmer, Diana Minchala, Karen Levy, Thomas Ristenpart, and Nicola Dell. “A Stalker’s Paradise”: How Intimate Partner Abusers Exploit Technology. In *ACM Conference on Human Factors in Computing Systems*, pages 667:1–667:13, 2018. doi:10.1145/3173574.3174241.
- [168] Diana Freed, Jackeline Palmer, Diana Elizabeth Minchala, Karen Levy, Thomas Ristenpart, and Nicola Dell. Digital technologies and intimate partner violence: A qualitative analysis with multiple stakeholders. *Proceedings of the ACM on Human-Computer Interaction*, 1(CSCW):46:1–46:22, 2017. doi:10.1145/3134681.
- [169] Rose de Fremery. Breaking the Cycle of Password Reuse, 2021. <https://blog.lastpass.com/2021/09/breaking-the-cycle-of-password-reuse/>.
- [170] Alisa Frik, Nathan Malkin, Marian Harbach, Eyal Peer, and Serge Egelman. A promise is a promise: the effect of commitment devices on computer security intentions. In *ACM Conference on Human Factors in Computing Systems*, pages 604:1–604:12, 2019. doi:10.1145/3290605.3300834.
- [171] Alisa Frik, Leysan Nurgalieva, Julia Bernd, Joyce Lee, Florian Schaub, and Serge Egelman. Privacy and security threat models and mitigation strategies of older adults. In *Symposium on Usable Privacy and Security*, 2019. <https://www.usenix.org/system/files/soups2019-frik.pdf>.
- [172] Kevin Gallagher, Sameer Patil, and Nasir Memon. New me: Understanding expert and non-expert perceptions and usage of the Tor anonymity network. In *Symposium on Usable Privacy and Security*, pages 385–398, 2017. <https://www.usenix.org/system/files/conference/soups2017/soups2017-gallagher.pdf>.
- [173] Stacia Garlach and Daniel Suthers. I’m Supposed to See That? AdChoices Usability in the Mobile Environment. In *Hawaii International Conference on System Sciences*, 2018. doi:10.24251/hicss.2018.476.
- [174] Shirley Gaw and Edward W Felten. Password management strategies for online accounts. In *Symposium on Usable Privacy and Security*, pages 44–55, 2006. doi:10.1145/1143120.1143127.



- [175] Christine Geeng, Mike Harris, Elissa Redmiles, and Franziska Roesner. “Like Lesbians Walking the Perimeter”: Experiences of US LGBTQ+ Folks With Online Security, Safety, and Privacy Advice. In *USENIX Security Symposium*, pages 305–322, 2022. <https://www.usenix.org/system/files/sec22-geeng.pdf>.
- [176] General Assembly of Maryland. Md. code ann. comm. law 14-3504: Maryland’s personal information protection act, 2018. <https://mgaleg.maryland.gov/mgawebsite/Laws/StatuteText?article=gcl&section=14-3504>.
- [177] Nina Gerber, Paul Gerber, and Melanie Volkamer. Explaining the privacy paradox: A systematic review of literature investigating privacy attitude and behavior. *Computers & Security*, 77:226–261, 2018. doi:10.1016/j.cose.2018.04.002.
- [178] Reza Ghaiumy Anaraky, Kaileigh Angela Byrne, Pamela J Wisniewski, Xinru Page, and Bart Knijnenburg. To disclose or not to disclose: examining the privacy decision-making processes of older vs. younger adults. In *ACM Conference on Human Factors in Computing Systems*, pages 686:1–686:14, 2021. doi:10.1145/3411764.3445204.
- [179] Susan Meyer Goldstein, Robert Johnston, JoAnn Duffy, and Jay Rao. The service concept: the missing link in service design research? *Journal of Operations management*, 20(2):121–134, 2002. doi:10.1016/S0272-6963(01)00090-0.
- [180] Maximilian Golla, Miranda Wei, Juliette Hainline, Lydia Filipe, Markus Dürmuth, Elissa Redmiles, and Blase Ur. What was that site doing with my facebook password?: Designing password-reuse notifications. In *ACM Conference on Computer and Communications Security*, pages 1549–1566, 2018. doi:10.1145/3243734.3243767.
- [181] Peter M Gollwitzer. Implementation intentions: strong effects of simple plans. *American Psychologist*, 54(7):493–503, 1999. doi:10.1037/0003-066X.54.7.493.
- [182] John Goodman. *Strategic Customer Service*. Amacom, 2019.
- [183] Google. Cleaning up after password dumps, 2014. <https://security.googleblog.com/2014/09/cleaning-up-after-password-dumps.html>.
- [184] Mark A Graber, Donna M D Alessandro, and Jill Johnson-West. Reading level of privacy policies on internet health web sites. *Journal of Family Practice*, 51(7):642–642, 2002. [https://cdn.mdedge.com/files/s3fs-public/Document/September-2017/5107JFP\\_BriefReport.pdf](https://cdn.mdedge.com/files/s3fs-public/Document/September-2017/5107JFP_BriefReport.pdf).
- [185] Paul A Grassi, Michael E Garcia, and James L Fenton. Nist special publication 800-63-3: Digital identity guidelines. Technical report, National Institute of Standards and Technology, U.S. Department of Commerce, 2020. doi:10.6028/NIST.SP.800-63-3.

- [186] Colin M Gray, Yubo Kou, Bryan Battles, Joseph Hoggatt, and Austin L Toombs. The dark (patterns) side of UX design. In *ACM Conference on Human Factors in Computing Systems*, pages 534:1–534:14, 2018. doi:10.1145/3173574.3174108.
- [187] Claire Greene and Joanna Stavins. Did the target data breach change consumer assessments of payment card security? *Journal of Payments Strategy & Systems*, 11(2):121–133, 2017. <https://www.bostonfed.org/-/media/Documents/researchdatareport/pdf/rdr1601.pdf>.
- [188] Sara S Greene. Stealing (identity) from the poor. *Minnesota Law Review*, 106:59–124, 2021. <https://heinonline.org/HOL/LandingPage?handle=hein.journals/mnlr106&div=7>.
- [189] Graham Greenleaf. Global Data Privacy Laws 2021: Despite COVID Delays, 145 Laws Show GDPR Dominance. *Privacy Laws & Business International Report*, 1:3–5, 2021. doi:10.2139/ssrn.3836348.
- [190] Mark Gridley and William J Jenkins. *An Analysis of Stanley Milgram’s Obedience to Authority: An Experimental View*. CRC Press, 2017.
- [191] Technology-Enabled Coercive Control Working Group. Same tactics, new tools, 2021. <https://teccworkinggroup.org/>.
- [192] Tamy Guberek, Allison McDonald, Sylvia Simioni, Abraham H Mhaidli, Kentaro Toyama, and Florian Schaub. Keeping a low profile? technology, risk and privacy among undocumented immigrants. In *ACM Conference on Human Factors in Computing Systems*, pages 114:1–114:15, 2018. doi:10.1145/3173574.3173688.
- [193] Robert Gunning. The fog index after twenty years. *Journal of Business Communication*, 6(2):3–13, 1969. doi:10.1177/002194366900600202.
- [194] Quang-An Ha, Jengchung Victor Chen, Ha Uy Uy, and Erik Paolo Capistrano. Exploring the privacy concerns in using intelligent virtual assistants under perspectives of information sensitivity and anthropomorphism. *International Journal of Human-Computer Interaction*, 37(6):1–16, 2020. doi:10.1080/10447318.2020.1834728.
- [195] Hana Habib, Jessica Colnago, Vidya Gopalakrishnan, Sarah Pearman, Jeremy Thomas, Alessandro Acquisti, Nicolas Christin, and Lorrie Faith Cranor. Away from prying eyes: analyzing usage and understanding of private browsing. In *Symposium on Usable Privacy and Security*, pages 159–175, 2018. <https://www.usenix.org/system/files/conference/soups2018/soups2018-habib-prying.pdf>.
- [196] Hana Habib, Pardis Emami Naeini, Summer Devlin, Maggie Oates, Chelse Swoopes, Lujo Bauer, Nicolas Christin, and Lorrie Faith Cranor. User behaviors

- and attitudes under password expiration policies. In *Symposium on Usable Privacy and Security*, pages 13–30, 2018. <https://www.usenix.org/system/files/conference/soups2018/soups2018-habib-password.pdf>.
- [197] Hana Habib, Sarah Pearman, Jiamin Wang, Yixin Zou, Alessandro Acquisti, Lorrie Faith Cranor, Norman Sadeh, and Florian Schaub. “It’s a scavenger hunt”: Usability of Websites’ Opt-Out and Data Deletion Choices. In *ACM Conference on Human Factors in Computing Systems*, pages 384:1–384:12, 2020. doi:10.1145/3313831.3376511.
- [198] Hana Habib, Yixin Zou, Aditi Jannu, Neha Sridhar, Chelse Swoopes, Alessandro Acquisti, Lorrie Faith Cranor, Norman Sadeh, and Florian Schaub. An empirical analysis of data deletion and opt-out choices on 150 websites. In *Symposium on Usable Privacy and Security*, 2019. <https://www.usenix.org/system/files/soups2019-habib.pdf>.
- [199] Hana Habib, Yixin Zou, Yaxing Yao, Alessandro Acquisti, Lorrie Faith Cranor, Joel Reidenberg, Norman Sadeh, and Florian Schaub. Toggles, Dollar Signs, and Triangles: How to (In)Effectively Convey Privacy Choices with Icons and Link Texts. In *ACM Conference on Human Factors in Computing Systems*, pages 63:1–63:25, 2021. doi:10.1145/3411764.3445387.
- [200] Tzipora Halevi, James Lewis, and Nasir Memon. A pilot study of cyber security and privacy related behavior and personality traits. In *International Conference on World Wide Web*, pages 737–744, 2013. doi:10.1145/2487788.2488034.
- [201] Sven Hammann, Michael Crabb, Saša Radomirović, Ralf Sasse, and David Basin. “I’m Surprised So Much Is Connected”: A Study on Users’ Online Accounts. In *ACM Conference on Human Factors in Computing Systems*, pages 620:1–620:13, 2022. doi:10.1145/3491102.3502125.
- [202] Bartłomiej Hanus and Yu “Andy” Wu. Impact of users’ security awareness on desktop security behavior: A protection motivation theory perspective. *Information Systems Management*, 33(1):2–16, 2016. doi:10.1080/10580530.2015.1117842.
- [203] SM Taiabul Haque, MD Romael Haque, Swapnil Nandy, Priyank Chandra, Mahdi Nasrullah Al-Ameen, Shion Guha, and Syed Ishtiaque Ahmed. Privacy vulnerabilities in public digital service centers in dhaka, bangladesh. In *ACM International Conference on Information and Communication Technologies and Development*, pages 14:1–14:12, 2020. doi:10.1145/3392561.3394642.
- [204] David Harborth and Sebastian Pape. Examining technology use factors of privacy-enhancing technologies: the role of perceived anonymity and trust. In *Americas Conference on Information Systems*, pages 1–10, 2018. <https://pape.science/files/publications/HP18amcis.pdf>.

- [205] Elizabeth Liz Harding, Jarno J Vanto, Reece Clark, L Hannah Ji, and Sara C Ainsworth. Understanding the scope and impact of the California Consumer Privacy Act of 2018. *Journal of Data Protection & Privacy*, 2(3):234–253, 2019. <https://www.ingentaconnect.com/content/hsp/jdpp/2019/00000002/00000003/art00007>.
- [206] Seth Hardy, Masashi Crete-Nishihata, Katharine Kleemola, Adam Senft, Byron Sonne, Greg Wiseman, Phillipa Gill, and Ronald J Deibert. Targeted threat index: Characterizing and quantifying politically-motivated targeted malware. In *USENIX Security Symposium*, pages 527–541, 2014. <https://www.usenix.org/system/files/conference/usenixsecurity14/sec14-paper-hardy.pdf>.
- [207] Eszter Hargittai and Alice Marwick. “What can I really do?” Explaining the privacy paradox with online apathy. *International Journal of Communication*, 10:3737–3757, 2016. <https://ijoc.org/index.php/ijoc/article/view/4655>.
- [208] Benjamin Harkin, Thomas L Webb, Betty PI Chang, Andrew Prestwich, Mark Conner, Ian Kellar, Yael Benn, and Paschal Sheeran. Does monitoring goal progress promote goal attainment? A meta-analysis of the experimental evidence. *Psychological Bulletin*, 142(2):198–229, 2016. doi:10.1037/bul0000025.
- [209] Diarmaid Harkin, Adam Molnar, and Erica Vowles. The commodification of mobile phone surveillance: An analysis of the consumer spyware industry. *Crime, Media, Culture*, 16(1):33–60, 2020. doi:10.1177/1741659018820562.
- [210] Erika Harrell. Victims of identity theft, 2016. Technical report, U.S. Department of Justice, 2019. <https://bjs.ojp.gov/content/pub/pdf/vit16.pdf>.
- [211] Zahra Hassanzadeh, Sky Marsen, and Robert Biddle. We’re here to help: Company image repair and user perception of data breaches. In *Proceedings of Graphics Interface*, pages 236–245, 2020. doi:10.20380/GI2020.24.
- [212] Sam Havron, Diana Freed, Rahul Chatterjee, Damon McCoy, Nicola Dell, and Thomas Ristenpart. Clinical computer security for victims of intimate partner violence. In *USENIX Security Symposium*, pages 105–122, 2019. <https://www.usenix.org/system/files/sec19-havron.pdf>.
- [213] Fritz Heider and Marianne Simmel. An experimental study of apparent behavior. *The American Journal of Psychology*, 57(2):243–259, 1944. doi:10.2307/1416950.
- [214] Joseph Henrich, Steven J Heine, and Ara Norenzayan. The weirdest people in the world? *Behavioral and Brain Sciences*, 33(2-3):61–83, 2010. doi:10.1017/S0140525X0999152X.

- [215] Nicola Henry and Anastasia Powell. Technology-facilitated sexual violence: A literature review of empirical research. *Trauma, violence, & abuse*, 19(2):195–208, 2018. doi:10.1177/1524838016650189.
- [216] Tejaswini Herath and H Raghav Rao. Protection motivation and deterrence: a framework for security policy compliance in organisations. *European Journal of Information Systems*, 18(2):106–125, 2009. doi:10.1057/ejis.2009.6.
- [217] Kristen Herhold. How people view facebook after the cambridge analytica data breach, 2019. <https://themanifest.com/social-media/how-people-view-facebook-after-cambridge-analytica-data-breach>.
- [218] Cormac Herley. So long, and no thanks for the externalities: the rational rejection of security advice by users. In *Workshop on New Security Paradigms*, pages 133–144, 2009. doi:10.1145/1719030.1719050.
- [219] Cormac Herley. More is not the answer. *IEEE Security & Privacy*, 12(1):14–19, 2013. doi:10.1109/MSP.2013.134.
- [220] Mindy Herman-Stabl, Mark Stemmler, and Anne Petersen. Approach and avoidant coping: Implications for adolescent mental health. *Journal of Youth and Adolescence*, 24(6):649–665, 1995. doi:10.1007/BF01536949.
- [221] E Tory Higgins. Self-discrepancy: a theory relating self and affect. *Psychological Review*, 94(3):319–340, 1987. doi:10.1037/0033-295X.94.3.319.
- [222] Mark Hochhauser. Lost in the fine print: Readability of financial privacy notices, 2001. <https://www.privacyrights.org/blog/lost-fine-print-readability-financial-privacy-notices-hochhauser>.
- [223] Leif-Erik Holtz, Katharina Nocun, and Marit Hansen. Towards displaying privacy information with icons. In *IFIP PrimeLife International Summer School on Privacy and Identity Management for Life*, pages 338–348, 2010. doi:10.1007/978-3-642-20769-3\_27.
- [224] Safe Horizon. Tech safety: For victims of crime, abuse, domestic violence and stalking, 2016. [https://www.safehorizon.org/wp-content/uploads/2016/06/1412609869\\_Tech-Safety-for-Victims-of-Abuse\\_Safe-Horizon.pdf](https://www.safehorizon.org/wp-content/uploads/2016/06/1412609869_Tech-Safety-for-Victims-of-Abuse_Safe-Horizon.pdf).
- [225] William K Horton. *The Icon Book: Visual Symbols for Computer Systems and Documentation*. John Wiley & Sons, Inc., 1994.
- [226] Patrick Houston. Why biometrics are about to put an end to password-only authentication, 2018. <https://www.symantec.com/blogs/feature-stories/why-biometrics-are-about-put-end-password-only-authentication>.
- [227] Adele E Howe, Indrajit Ray, Mark Roberts, Malgorzata Urbanska, and Zinta Byrne. The psychology of security for the home computer user. In *IEEE Symposium on Security and Privacy*, pages 209–223, 2012. doi:10.1109/SP.2012.23.

- [228] Mariea Grubbs Hoy and George Milne. Gender differences in privacy-related measures for young adult Facebook users. *Journal of Interactive Advertising*, 10(2):28–45, 2010. doi:10.1080/15252019.2010.10722168.
- [229] Yue Huang, Borke Obada-Obieh, and Konstantin Beznosov. Users’ perceptions of chrome compromised credential notification. In *Symposium on Usable Privacy and Security*, pages 155–174, 2022. [https://www.usenix.org/system/files/soups2022-huang\\_1.pdf](https://www.usenix.org/system/files/soups2022-huang_1.pdf).
- [230] Connor Huff and Dustin Tingley. “Who Are These People?” Evaluating the Demographic Characteristics and Political Preferences of MTurk Survey Respondents. *Research & Politics*, 2(3), 2015. doi:10.1177/2053168015604648.
- [231] Jun Ho Huh, Hyoungshick Kim, Swathi SVP Rayala, Rakesh B Bobba, and Konstantin Beznosov. I’m too busy to reset my LinkedIn password: On the effectiveness of password reset emails. In *ACM Conference on Human Factors in Computing Systems*, pages 387–391, 2017. doi:10.1145/3025453.3025788.
- [232] Troy Hunt. Have I Been Pwned: Check if you have an account that has been compromised in a data breach, 2022. <https://haveibeenpwned.com/>.
- [233] Troy Hunt. Pwned websites, 2022. <https://haveibeenpwned.com/PwnedWebsites>.
- [234] Renato Iannella and Adam Finden. Privacy Awareness: Icons and Expression for Social Networks. In *International Workshop for Technical, Economic and Legal Aspects of Business Models for Virtual Goods*, pages 1:1–1:15, 2010. [http://virtualgoods.org/2010/VirtualGoodsBook2010\\_13.pdf](http://virtualgoods.org/2010/VirtualGoodsBook2010_13.pdf).
- [235] Identity Theft Resource Center. The aftermath: the non-economic impacts of identity theft. Technical report, Identity Theft Resource Center, 2018. [https://www.idtheftcenter.org/wp-content/uploads/2018/09/ITRC\\_Aftermath-2018\\_Web\\_FINAL.pdf](https://www.idtheftcenter.org/wp-content/uploads/2018/09/ITRC_Aftermath-2018_Web_FINAL.pdf).
- [236] Identity Theft Resource Center. 2021 Annual Data Breach Report – Identity Compromises: From the Era of Identity Theft to the Age of Identity Fraud. Technical report, Identity Theft Resource Center, 2022. <https://www.idtheftcenter.org/post/identity-theft-resource-center-2021-annual-data-breach-report-sets-new-record-for-number-of-compromises/>.
- [237] Princely Ifinedo. Understanding information systems security policy compliance: An integration of the theory of planned behavior and the protection motivation theory. *Computers & Security*, 31(1):83–95, 2012. doi:10.1016/j.cose.2011.10.007.
- [238] Sarah E Igo. *The known citizen: A history of privacy in modern America*. Harvard University Press, 2018.

- [239] NortonLifeLock Inc. Norton internet security center, 2020. <https://us.norton.com/internetsecurity>.
- [240] Information Commissioner’s Office. Guide to the General Data Protection Regulation (GDPR), 2019. <https://ico.org.uk/for-organisations/guide-to-data-protection/guide-to-the-general-data-protection-regulation-gdpr/>.
- [241] Philip G Inglesant and M Angela Sasse. The true cost of unusable password policies: password use in the wild. In *ACM Conference on Human Factors in Computing Systems*, pages 383–392, 2010. doi:10.1145/1753326.1753384.
- [242] Iulia Ion, Rob Reeder, and Sunny Consolvo. “... no one can hack my mind”: Comparing Expert and Non-Expert Security Practices. In *Symposium On Usable Privacy and Security*, pages 327–346, 2015. <https://www.usenix.org/system/files/conference/soups2015/soups15-paper-ion.pdf>.
- [243] Sarah J Isherwood, Siné JP McDougall, and Martin B Curry. Icon identification in context: The changing role of icon characteristics with user experience. *Human Factors*, 49(3):465–476, 2007. doi:10.1518/001872007X200102.
- [244] Alexander Jenkins, Murugan Anandarajan, and Rob D’Ovidio. ‘All that glitters is not gold’: The role of impression management in data breach notification. *Western Journal of Communication*, 78(3):337–357, 2014. doi:10.1080/10570314.2013.866686.
- [245] Jeffrey Jenkins, Alexandra Durcikova, and Jay F Nunamaker Jr. Mitigating the Security Intention-Behavior Gap: The Moderating Role of Required Effort on the Intention-Behavior Relationship. *Journal of the Association for Information Systems*, 22(1):246–272, 2021. doi:10.17705/1jais.00660.
- [246] Jeffrey L Jenkins, Mark Grimes, Jeffrey Gainer Proudfoot, and Paul Benjamin Lowry. Improving password cybersecurity through inexpensive and minimally invasive means: Detecting and deterring password reuse through keystroke-dynamics monitoring and just-in-time fear appeals. *Information Technology for Development*, 20(2):196–213, 2014.
- [247] Carlos Jensen, Colin Potts, and Christian Jensen. Privacy practices of internet users: Self-reports versus observed behavior. *International Journal of Human-Computer Studies*, 63(1-2):203–227, 2005. doi:10.1016/j.ijhcs.2005.04.019.
- [248] Haiyan Jia, Pamela J Wisniewski, Heng Xu, Mary Beth Rosson, and John M Carroll. Risk-taking as a learning process for shaping teen’s online information privacy behaviors. In *ACM Conference on Computer Supported Cooperative Work & Social Computing*, pages 583–599, 2015. doi:10.1145/2675133.2675287.
- [249] Zhenhui Jiang, Cheng Suang Heng, and Ben CF Choi. Privacy concerns and privacy-protective behavior in synchronous online social interactions. *Information Systems Research*, 24(3):579–595, 2013. doi:10.1287/isre.1120.0441.

- [250] Alison Grace Johansen. What to do after 5 types of data breaches, 2021. <https://us.norton.com/internetsecurity-emerging-threats-what-to-do-after-a-data-breach.html#>.
- [251] Alex Johnson. Equifax breaks down just how bad last year's data breach was, 2018. <https://www.nbcnews.com/news/us-news/equifax-breaks-down-just-how-bad-last-year-s-data-n872496>.
- [252] Allen C Johnston and Merrill Warkentin. Fear appeals and information security behaviors: An empirical study. *MIS Quarterly*, 34(3):549–566, 2010. doi:10.2307/25750691.
- [253] Penny Jones and Jerry Powell. Gary anderson has been found! *Resource Recycling*, 18:25–27, 1999. <https://discardstudies.com/wp-content/uploads/2012/07/garyandersonfound.pdf>.
- [254] Yoonhyuk Jung and Jonghwa Park. An investigation of relationships among privacy concerns, affective responses, and coping behaviors in location-based services. *International Journal of Information Management*, 43:15–24, 2018. doi:10.1016/j.ijinfomgt.2018.05.007.
- [255] Daniel Kahneman, Jack L Knetsch, and Richard H Thaler. Anomalies: The endowment effect, loss aversion, and status quo bias. *Journal of Economic Perspectives*, 5(1):193–206, 1991. doi:10.1257/jep.5.1.193.
- [256] Daniel Kahneman and Amos Tversky. Subjective probability: A judgment of representativeness. *Cognitive Psychology*, 3(3):430–454, 1972. doi:10.1016/0010-0285(72)90016-3.
- [257] Daniel Kahneman and Amos Tversky. Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2):263–292, 1979. doi:10.2307/1914185.
- [258] Margot E Kaminski. Standing after snowden: lessons on privacy harm from national security surveillance litigation. *DePaul Law Review*, 66:413–438, 2016. <https://heinonline.org/HOL/LandingPage?handle=hein.journals/deplr66&div=19>.
- [259] Saraschandra Karanam, Janhavi Viswanathan, Anand Theertha, Bipin Indurkha, and Herre Van Oostendorp. Impact of placing icons next to hyperlinks on information-retrieval tasks on the web. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, pages 2834–2839, 2010. <https://escholarship.org/content/qt27w0n9kc/qt27w0n9kc.pdf>.
- [260] Sowmya Karunakaran, Kurt Thomas, Elie Bursztein, and Oxana Comanescu. Data breaches: user comprehension, expectations, and concerns with handling exposed data. In *Symposium on Usable Privacy and Security*, pages 217–234, 2018. <https://www.usenix.org/system/files/conference/soups2018/soups2018-karunakaran.pdf>.



- [261] Flavius Kehr, Tobias Kowatsch, Daniel Wentzel, and Elgar Fleisch. Blissfully ignorant: the effects of general privacy concerns, general institutional trust, and affect in the privacy calculus. *Information Systems Journal*, 25(6):607–635, 2015. doi:10.1111/isj.12062.
- [262] Mark J Keith, Samuel C Thompson, Joanne Hale, Paul Benjamin Lowry, and Chapman Greer. Information disclosure on mobile devices: Re-examining privacy calculus with actual user behavior. *International Journal of Human-Computer Studies*, 71(12):1163–1173, 2013.
- [263] Patrick Gage Kelley, Lucian Cesca, Joanna Bresee, and Lorrie Faith Cranor. Standardizing privacy notices: an online study of the nutrition label approach. In *ACM Conference on Human factors in Computing Systems*, pages 1573–1582, 2010. doi:10.1145/1753326.1753561.
- [264] Timothy Kelley and Bennett I Bertenthal. Attention and past behavior, not security knowledge, modulate users’ decisions to login to insecure websites. *Information & Computer Security*, 24(2):164–176, 2016. doi:10.1108/ICS-01-2016-0002.
- [265] Jeanette Kerr, Carolyn Whyte, and Heather Strang. Targeting escalation and harm in intimate partner violence: Evidence from northern territory police, australia. *Cambridge Journal of Evidence-Based Policing*, 1:143–159, 2017. doi:10.1007/s41887-017-0005-z.
- [266] Murat Kezer, Barış Sevi, Zeynep Cemalcilar, and Lemi Baruh. Age differences in privacy attitudes, literacy and privacy management on facebook. *Cyberpsychology: Journal of Psychosocial Research on Cyberspace*, 10(1):2:1–2:20, 2016. doi:10.5817/CP2016-1-2.
- [267] Mohammad Taha Khan, Maria Hyun, Chris Kanich, and Blase Ur. Forgotten but not gone: Identifying the need for longitudinal data management in cloud storage. In *ACM Conference on Human Factors in Computing Systems*, pages 543:1–543:12, 2018. doi:10.1145/3173574.3174117.
- [268] Bokyoung Kim, Kristine Johnson, and Sun-Young Park. Lessons from the five data breaches: Analyzing framed crisis response strategies and crisis severity. *Cogent Business & Management*, 4(1):1–15, 2017. doi:10.1080/23311975.2017.1354525.
- [269] Jennifer King. CCPA Comments, Round Two: The Battle of the Do-Not-Sell Button, 2020. <https://cyberlaw.stanford.edu/blog/2020/02/ccpa-comments-round-two-battle-do-not-sell-button>.
- [270] Keita Kinjo and Takeshi Ebina. Paradox of choice and consumer nonpurchase behavior. *AI & Society*, 30(2):291–297, 2015. doi:10.1007/s00146-014-0546-7.

- [271] Amanda Kippert. Empowering survivors: Why domestic violence advocates say the best way to help survivors is to give them back control, 2015. <https://www.domesticshelters.org/articles/escaping-violence/empowering-survivors>.
- [272] Kleimann Communication Group. Evolution of a prototype financial privacy notice: A report on the form development project. Technical report, The Federal Trade Commission, 2006. [https://www.ftc.gov/sites/default/files/documents/reports/evolution-prototype-financial-privacy-notice-report-form-development-project/evolution\\_prototype\\_financial\\_privacy\\_notice.pdf](https://www.ftc.gov/sites/default/files/documents/reports/evolution-prototype-financial-privacy-notice-report-form-development-project/evolution_prototype_financial_privacy_notice.pdf).
- [273] Aaron Klein. America's poor subsidize wealthier consumers in a vicious income inequality cycle, 2018. <https://www.nbcnews.com/think/opinion/america-s-poor-subsidize-wealthier-consumers-vicious-income-inequality-cycle-ncna845091>.
- [274] Spyros Kokolakis. Privacy attitudes and privacy behaviour: A review of current research on the privacy paradox phenomenon. *Computers & Security*, 64:122–134, 2017. doi:10.1016/j.cose.2015.07.002.
- [275] Stefan Korff and Rainer Böhme. Too much choice: End-user privacy decisions in the context of choice proliferation. In *Symposium On Usable Privacy and Security*, pages 69–87, 2014. <https://www.usenix.org/system/files/soups14-paper-korff.pdf>.
- [276] Brian Krebs. Are credit monitoring services worth it?, 2014. <https://krebsonsecurity.com/2014/03/are-credit-monitoring-services-worth-it/>.
- [277] Lauren I Labrecque, Ereni Markos, and Aron Darmody. Addressing online behavioral advertising and privacy implications: A comparison of passive versus active learning approaches. *Journal of Marketing Education*, 43(1):1–16, 2019. doi:10.1177/0273475319828788.
- [278] Ravie Lakshmanan. Chrome and Firefox will now alert you about data breaches involving your accounts, 2019. <https://thenextweb.com/security/2019/10/23/chrome-and-firefox-will-now-alert-you-about-data-breaches-involving-your-accounts/>.
- [279] Issie Lapowsky. How cambridge analytica sparked the great privacy awakening, 2019. <https://www.wired.com/story/cambridge-analytica-facebook-privacy-awakening/>.
- [280] Robert LaRose, Nora J Rifon, and Richard Enbody. Promoting personal responsibility for internet safety. *Communications of the ACM*, 51(3):71–76, 2008. doi:10.1145/1325555.1325569.

- [281] Elmer Lastdrager, Inés Carvajal Gallardo, Pieter Hartel, and Marianne Junger. How effective is anti-phishing training for children? In *Symposium on Usable Privacy and Security*, pages 229–239, 2017. <https://www.usenix.org/system/files/conference/soups2017/soups2017-lastdrager.pdf>.
- [282] LastPass. Digital security dashboard, 2022. <https://www.lastpass.com/features/security-dashboard>.
- [283] Josephine Lau, Benjamin Zimmerman, and Florian Schaub. Alexa, are you listening? privacy perceptions, concerns and privacy-seeking behaviors with smart speakers. *Proceedings of the ACM on Human-Computer Interaction*, 2(CSCW):102:1–102:31, 2018. doi:10.1145/3274371.
- [284] Kristin Lauter, Sreekanth Kannepalli, Kim Laine, and Radames Cruz Moreno. Password Monitor: Safeguarding passwords in Microsoft Edge, 2021. <https://www.microsoft.com/en-us/research/blog/password-monitor-safeguarding-passwords-in-microsoft-edge/>.
- [285] Roslyn Layton. The 10 Problems of the GDPR: The US can learn from the EU’s mistakes and leapfrog its policy, 2019. <https://www.judiciary.senate.gov/imo/media/doc/Layton%20Testimony1.pdf>.
- [286] Jonathan Lazar, Jinjuan Heidi Feng, and Harry Hochheiser. *Research methods in human-computer interaction*. Morgan Kaufmann, 2017.
- [287] Joseph Lazzarotti, Jason Gavejian, and Maya Atrakchi. Security breach notification laws, 2018. <http://www.ncsl.org/research/telecommunications-and-information-technology/security-breach-notification-laws.aspx>.
- [288] Stevens Le Blond, Alejandro Cuevas, Juan Ramón Troncoso-Pastoriza, Philipp Jovanovic, Bryan Ford, and Jean-Pierre Hubaux. On enforcing the digital immunity of a large humanitarian organization. In *IEEE Symposium on Security and Privacy*, pages 424–440, 2018. doi:10.1109/SP.2018.00019.
- [289] Doohwang Lee, Robert Larose, and Nora Rifon. Keeping our network safe: a model of online protection behaviour. *Behaviour & Information Technology*, 27(5):445–454, 2008. doi:10.1080/01449290600879344.
- [290] Rainie Lee, Sara Kiesler, Ruogu Kang, and Mary Madden. Anonymity, privacy, and security online. Technical report, Pew Research Center, 2013. <https://www.pewresearch.org/internet/2013/09/05/anonymity-privacy-and-security-online/>.
- [291] Roxanne Leitão. Anticipating smart home security and privacy threats with survivors of intimate partner abuse. In *ACM Conference on Designing Interactive Systems*, pages 527–539, 2019. doi:10.1145/3322276.3322366.

- [292] Roxanne Leitão. Technology-facilitated intimate partner abuse: A qualitative analysis of data from online domestic abuse forums. *Human-Computer Interaction*, 36(3):1–40, 2019. doi:10.1080/07370024.2019.1685883.
- [293] Amanda Lenhart, Mary Madden, Sandra Cortesi, Urs Gasser, and Aaron Smith. Where teens seek online privacy advice. Technical report, Pew Research Center, 2013. <https://core.ac.uk/download/pdf/30676198.pdf>.
- [294] Pedro Leon, Justin Cranshaw, Lorrie Faith Cranor, Jim Graves, Manoj Hastak, Blase Ur, and Guzi Xu. What do online behavioral advertising privacy disclosures communicate to users? In *Workshop on Privacy in the Electronic Society*, pages 19–30, 2012. doi:10.1145/2381966.2381970.
- [295] Pedro Leon, Blase Ur, Richard Shay, Yang Wang, Rebecca Balebako, and Lorrie Cranor. Why johnny can’t opt out: a usability evaluation of tools to limit online behavioral advertising. In *ACM Conference on Human Factors in Computing Systems*, pages 589–598, 2012. doi:10.1145/2207676.2207759.
- [296] Karen Levy and Bruce Schneier. Privacy threats in intimate relationships. *Journal of Cybersecurity*, 6(1):1–13, 2020. doi:10.1093/cybsec/tyaa006.
- [297] Kevin Lewis, Jason Kaufman, and Nicholas Christakis. The taste for privacy: An analysis of college student privacy settings in an online social network. *Journal of Computer-Mediated Communication*, 14(1):79–100, 2008. doi:10.1111/j.1083-6101.2008.01432.x.
- [298] Han Li, Rathindra Sarathy, and Heng Xu. The role of affect and cognition on online consumers’ decision to disclose personal information to unfamiliar online vendors. *Decision Support Systems*, 51(3):434–445, 2011. doi:10.1016/j.dss.2011.01.017.
- [299] Johnny Lieu. Terms and Conditions are too long, just ask a guy who read Amazon’s for 9 hours, 2017. <https://mashable.com/2017/03/15/reading-amazons-terms-conditions/#IQDa1u7Bs0q0>.
- [300] Eric Lin, Saul Greenberg, Eileah Trotter, David Ma, and John Aycock. Does domain highlighting help people identify phishing sites? In *ACM Conference on Human Factors in Computing Systems*, pages 2075–2084, 2011. doi:10.1145/1978942.1979244.
- [301] Janne Lindqvist, Justin Cranshaw, Jason Wiese, Jason Hong, and John Zimmerman. I’m the mayor of my house: examining why people use foursquare - a social-driven location sharing application. In *ACM Conference on Human Factors in Computing Systems*, pages 2409–2418, 2011. doi:10.1145/1978942.1979295.

- [302] Bin Liu, Mads Schaarup Andersen, Florian Schaub, Hazim Almuhiemedi, Shikun Aerin Zhang, Norman Sadeh, Yuvraj Agarwal, and Alessandro Acquisti. Follow my recommendations: A personalized privacy assistant for mobile app permissions. In *Symposium on Usable Privacy and Security*, pages 27–41, 2016. <https://www.usenix.org/system/files/conference/soups2016/soups2016-paper-liu.pdf>.
- [303] Katie Lobosco. Here’s what happened when I tried to freeze my credit, 2017. <https://money.cnn.com/2017/09/14/pf/equifax-freeze-my-credit/index.html>.
- [304] George F Loewenstein, Elke U Weber, Christopher K Hsee, and Ned Welch. Risk as feelings. *Psychological Bulletin*, 127(2):267–286, 2001. doi:10.1037/0033-2909.127.2.267.
- [305] Isabel Lopez-Neira, Trupti Patel, Simon Parkin, George Danezis, and Leonie Tanczer. ‘Internet of Things’: How Abuse is Getting Smarter. *Safe – The Domestic Abuse Quarterly*, 63:22–26, 2019. doi:10.2139/ssrn.3350615.
- [306] Marika Lüders and Petter Bae Brandtzæg. ‘My children tell me it’s so simple’: a mixed-methods approach to understand older non-users’ perceptions of social networking sites. *New Media & Society*, 19(2):181–198, 2017. doi:10.1177/1461444814554064.
- [307] Annamaria Lusardi and Olivia S Mitchell. How ordinary consumers make complex economic decisions: Financial literacy and retirement readiness. *Quarterly Journal of Finance*, 7(3):1750008, 2017. doi:10.1142/S2010139217500082.
- [308] Eleanor Lyon, Jill Bradshaw, and Anne Menard. Meeting survivor needs through non-residential domestic violence services & supports: Results of a multi-state study. Technical report, National Resource Center on Domestic Violence, 2012. <https://vawnet.org/sites/default/files/materials/files/2016-07/DVServicesStudy-FINALReport2011.pdf>.
- [309] Mary Madden, Susannah Fox, Aaron Smith, and Jessica Vitak. Americans’ attitudes about privacy, security and surveillance. Technical report, Pew Internet & American Life Project, 2007. <https://www.pewtrusts.org/en/research-and-analysis/reports/2007/12/16/digital-footprints-online-identity-management-and-search-in-the-age-of-transparency>.
- [310] Mary Madden, Michele Gilman, Karen Levy, and Alice Marwick. Privacy, poverty, and big data: A matrix of vulnerabilities for poor americans. *Washington University Law Review*, 95:53–126, 2017. <https://heinonline.org/HOL/LandingPage?handle=hein.journals/walq95&div=5>.
- [311] Mary Madden and Lee Rainie. Americans’ attitudes about privacy, security and surveillance. Technical report, Pew Research Center, 2015. [https://www.pewresearch.org/internet/wp-content/uploads/sites/9/2015/05/Privacy-and-Security-Attitudes-5.19.15\\_FINAL.pdf](https://www.pewresearch.org/internet/wp-content/uploads/sites/9/2015/05/Privacy-and-Security-Attitudes-5.19.15_FINAL.pdf).

- [312] Maureen Mahoney. California Consumer Privacy Act: Are Consumers' Digital Rights Protected? Technical report, Consumer Reports, 2020. [https://advocacy.consumerreports.org/wp-content/uploads/2020/09/CR\\_CCPA-Are-Consumers-Digital-Rights-Protected\\_092020\\_vf.pdf](https://advocacy.consumerreports.org/wp-content/uploads/2020/09/CR_CCPA-Are-Consumers-Digital-Rights-Protected_092020_vf.pdf).
- [313] Malwarebytes. Malwarebytes lab, 2020. <https://blog.malwarebytes.com/>.
- [314] Bernard Mar. GDPR: The Biggest Data Breaches And The Shocking Fines (That Would Have Been), 2018. <https://www.forbes.com/sites/bernardmarr/2018/06/11/gdpr-the-biggest-data-breaches-and-the-shocking-fines-that-would-have-been/#199b5b4b6c10>.
- [315] Meghan L Marsac, Nancy Kassam-Adams, Aimee K Hildenbrand, Elizabeth Nicholls, Flaura K Winston, Stephen S Leff, and Joel Fein. Implementing a trauma-informed approach in pediatric health care networks. *JAMA pediatrics*, 170(1):70–77, 2016. doi:10.1001/jamapediatrics.2015.2206.
- [316] Kirsten E Martin. Transaction costs, privacy, and trust: The laudable goals and ultimate failure of notice and choice to respect privacy online. *First Monday*, 18(12-2), 2013. doi:10.5210/fm.v18i12.4838.
- [317] Alice Marwick and Eszter Hargittai. Nothing to hide, nothing to lose? incentives and disincentives to sharing information with institutions online. *Information, Communication & Society*, 22(12):1697–1713, 2019. doi:10.1080/1369118X.2018.1450432.
- [318] Maryland Office of the Attorney General. Guidelines for Businesses to Comply with the Maryland Personal Information Protection Act. <https://www.marylandattorneygeneral.gov/Pages/IdentityTheft/businessGL.aspx>, 2021.
- [319] Michael Massimi, Jill P Dimond, and Christopher A Le Dantec. Finding a new normal: the role of technology in life disruptions. In *ACM Conference on Computer Supported Cooperative Work & Social Computing*, pages 719–728, 2012. doi:10.1145/2145204.2145314.
- [320] Manfred Massironi. *The Psychology of Graphic Images: Seeing, Drawing, Communicating*. Psychology Press, 2001.
- [321] Arunesh Mathur, Gunes Acar, Michael J Friedman, Elena Lucherini, Jonathan Mayer, Marshini Chetty, and Arvind Narayanan. Dark patterns at scale: Findings from a crawl of 11k shopping websites. *Proceedings of the ACM on Human-Computer Interaction*, 3(CSCW):81:1–81:32, 2019. doi:10.1145/3359183.
- [322] Christian Matt and Philipp Peckelsen. Sweet idleness, but why? how cognitive factors and personality traits affect privacy-protective behavior. In *Hawaii International Conference on System Sciences*, pages 4832–4841, 2016. doi:10.1109/HICSS.2016.599.

- [323] Célestin Matte, Nataliia Bielova, and Cristiana Santos. Do Cookie Banners Respect my Choice?: Measuring Legal Compliance of Banners from IAB Europe’s Transparency and Consent Framework. In *IEEE Symposium on Security and Privacy*, pages 791–809, 2020. doi:10.1109/SP40000.2020.00076.
- [324] Tara Matthews, Kerwell Liao, Anna Turner, Marianne Berkovich, Robert Reeder, and Sunny Consolvo. “She’ll just grab any device that’s closer” A Study of Everyday Device & Account Sharing in Households. In *ACM Conference on Human Factors in Computing Systems*, pages 5921–5932, 2016. doi:10.1145/2858036.2858051.
- [325] Tara Matthews, Kathleen O’Leary, Anna Turner, Manya Sleeper, Jill Palzkill Woelfer, Martin Shelton, Cori Manthorne, Elizabeth F Churchill, and Sunny Consolvo. Stories from survivors: Privacy & security practices when coping with intimate partner abuse. In *ACM Conference on Human Factors in Computing Systems*, pages 2189–2201, 2017. doi:10.1145/3025453.3025875.
- [326] Peter Mayer, Alexandra Kunz, and Melanie Volkamer. Reliable behavioural factors in the information security context. In *ACM International Conference on Availability, Reliability and Security*, pages 9:1–9:10, 2017. doi:10.1145/3098954.3098986.
- [327] Peter Mayer, Collins W Munyendo, Michelle L Mazurek, and Adam J Aviv. Why users (don’t) use password managers at a large educational institution. In *USENIX Security Symposium*, pages 1849–1866, 2022. <https://www.usenix.org/system/files/sec22-mayer.pdf>.
- [328] Peter Mayer, Yixin Zou, Byron M Lowens, Hunter A Dyer, Khue Le, Florian Schaub, and Adam J Aviv. Awareness, Intention, (In)Action: Individuals’ Reactions to Data Breaches, 2022. Under review.
- [329] Peter Mayer, Yixin Zou, Florian Schaub, and Adam Aviv. “Now I’m a bit angry:” User Awareness, Perception, and Responses to Data Breaches. In *USENIX Security Symposium*. USENIX Association, 2021. <https://www.usenix.org/system/files/sec21fall-mayer.pdf>.
- [330] Michelle L Mazurek, Saranga Komanduri, Timothy Vidas, Lujo Bauer, Nicolas Christin, Lorrie Faith Cranor, Patrick Gage Kelley, Richard Shay, and Blase Ur. Measuring password guessability for an entire university. In *ACM Conference on Computer and Communications Security*, pages 173–186, 2013. doi:10.1145/2508859.2516726.
- [331] Agata McCormac, Tara Zwaans, Kathryn Parsons, Dragana Calic, Marcus Butavicius, and Malcolm Pattinson. Individual differences and information security awareness. *Computers in Human Behavior*, 69:151–156, 2017. doi:10.1016/j.chb.2016.11.065.

- [332] Aleecia M McDonald and Lorrie Faith Cranor. The cost of reading privacy policies. *Journal of Law and Policy for the Information Society*, 4(3):543–568, 2008. <https://heinonline.org/HOL/LandingPage?handle=hein.journals/isjlp soc4&div=27>.
- [333] Aleecia M McDonald, Robert W Reeder, Patrick Gage Kelley, and Lorrie Faith Cranor. A comparative study of online privacy policies and formats. In *International Symposium on Privacy Enhancing Technologies*, pages 37–55, 2009. doi:10.1007/978-3-642-03168-7\_3.
- [334] Allison McDonald. *Advancing Digital Safety for High-Risk Communities*. PhD thesis, University of Michigan, 2022.
- [335] Allison McDonald, Catherine Barwulor, Michelle L Mazurek, Florian Schaub, and Elissa M Redmiles. “It’s stressful having all these phones”: Investigating Sex Workers’ Safety Goals, Risks, and Practices Online. In *USENIX Security Symposium*, 2021. <https://www.usenix.org/system/files/sec21fall-mcdonald.pdf>.
- [336] Nora McDonald, Sarita Schoenebeck, and Andrea Forte. Reliability and interrater reliability in qualitative research: Norms and guidelines for CSCW and HCI practice. *Proceedings of the ACM on Human-Computer Interaction*, 3(CSCW):72:1–72:23, 2019. doi:10.1145/3359174.
- [337] Kenneth O McGraw and Seok P Wong. Forming inferences about some intraclass correlation coefficients. *Psychological Methods*, 1(1):30–46, 1996. doi:10.1037/1082-989X.1.1.30.
- [338] B Dawn Medlin and Joseph A Cazier. An Empirical Investigation: Health Care Employee Passwords and Their Crack Times in Relationship to HIPAA Security Standards. *International Journal of Healthcare Information Systems and Informatics*, 2(3):39–48, 2007. doi:10.4018/jhisi.2007070104.
- [339] Jean-Frederick Menard. A ‘nudge’ for public health ethics: libertarian paternalism as a framework for ethical analysis of public health interventions? *Public Health Ethics*, 3(3):229–238, 2010. doi:10.1093/phe/phq024.
- [340] Tamir Mendel and Eran Toch. Susceptibility to social influence of privacy behaviors: Peer versus authoritative sources. In *ACM Conference on Computer Supported Cooperative Work & Social Computing*, pages 581–593, 2017. doi:10.1145/2998181.2998323.
- [341] Helena M Mentis, Galina Madjaroff, Aaron Massey, and Zoya Trendafilova. The illusion of choice in discussing cybersecurity safeguards between older adults with mild cognitive impairment and their caregivers. *Proceedings of the ACM on Human-Computer Interaction*, 4(CSCW):164:1–164:19, 2020. doi:10.1145/3415235.



- [342] Helena M Mentis, Galina Madjaroff, and Aaron K Massey. Upside and downside risk in online security for older adults with mild cognitive impairment. In *ACM Conference on Human Factors in Computing Systems*, pages 343:1–343:13, 2019. doi:10.1145/3290605.3300573.
- [343] Besty Mikel. Designing a Content-First Experience on Firefox Monitor, 2020. <https://medium.com/firefox-ux/designing-a-content-first-experience-on-firefox-monitor-9b8875d44386>.
- [344] Vyacheslav Mikhed and Michael Vogan. Out of sight, out of mind: consumer reaction to news on data breaches and identity theft. *SSRN*, 2015. [https://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=2691902](https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2691902).
- [345] Vyacheslav Mikhed and Michael Vogan. How data breaches affect consumer credit. *Journal of Banking & Finance*, 88:192–207, 2018. doi:10.1016/j.jbankfin.2017.12.002.
- [346] Matthew B Miles, A Michael Huberman, and Johnny Saldaña. *Qualitative data analysis: A methods sourcebook*. Sage publications, 2018.
- [347] George R Milne and Mary J Culnan. Strategies for reducing online privacy risks: Why consumers read (or don’t read) online privacy notices. *Journal of Interactive Marketing*, 18(3):15–29, 2004. doi:10.1002/dir.20009.
- [348] Drew Mitnick. No more waiting: it’s time for a federal data breach law in the U.S., 2018. <https://www.accessnow.org/no-more-waiting-its-time-for-a-federal-data-breach-law-in-the-u-s/>.
- [349] Reham Ebada Mohamed and Sonia Chiasson. Online privacy and aging of digital artifacts. In *Symposium on Usable Privacy and Security*, pages 177–195, 2018. <https://www.usenix.org/system/files/conference/soups2018/soups2018-mohamed.pdf>.
- [350] Saif M. Mohammad and Peter D. Turney. Crowdsourcing a word-emotion association lexicon. *Computational Intelligence*, 29(3):436–465, 2013. doi:10.1111/j.1467-8640.2012.00460.x.
- [351] Don A Moore and Paul J Healy. The trouble with overconfidence. *Psychological Review*, 115(2):502–517, 2008. doi:10.1037/0033-295X.115.2.502.
- [352] Mozilla. What to do after a data breach, 2019. <https://blog.mozilla.org/firefox/what-to-do-after-a-data-breach/>.
- [353] Mozilla. Privacy icons, 2020. [https://wiki.mozilla.org/Privacy\\_Icons](https://wiki.mozilla.org/Privacy_Icons).
- [354] Mozilla. Firefox monitor, 2021. <https://monitor.firefox.com/>.
- [355] Mozilla. Firefox relay, 2021. <https://relay.firefox.com/>.

- [356] Deirdre K Mulligan, Priscilla M Regan, and Jennifer King. The fertile dark matter of privacy takes on the dark patterns of surveillance. *Journal of Consumer Psychology*, 30(4):767–773, 2020. doi:10.1002/jcpy.1190.
- [357] Ambar Murillo, Andreas Kramm, Sebastian Schnorf, and Alexander De Luca. “If I press delete, it’s gone”: User Understanding of Online Data Deletion and Expiration. In *Symposium on Usable Privacy and Security*, pages 329–339, 2018. <https://www.usenix.org/system/files/conference/soups2018/soups2018-murillo.pdf>.
- [358] Christine E Murray, G Evette Horton, Catherine Higgins Johnson, Lori Notestine, Bethany Garr, Allison Marsh Pow, Paulina Flasch, and Elizabeth Doom. Domestic violence service providers’ perceptions of safety planning: A focus group study. *Journal of Family Violence*, 30(3):381–392, 2015. doi:10.1007/s10896-015-9674-1.
- [359] Steven Muzatko and Gaurav Bansal. Timing of data breach announcement and e-commerce trust. In *Midwest Association for Information Systems Conference*, pages 7:1–7:7, 2018. <https://core.ac.uk/download/pdf/301374905.pdf>.
- [360] Moses Namara, Daricia Wilkinson, Kelly Caine, and Bart P Knijnenburg. Emotional and practical considerations towards the adoption and abandonment of vpns as a privacy-enhancing technology. *Proceedings on Privacy Enhancing Technologies*, 2020(1):83–102, 2020. doi:10.2478/popets-2020-0006.
- [361] National Conference of State Legislators. 2018 security breach legislation, 2018. <http://www.ncsl.org/research/telecommunications-and-information-technology/2018-security-breach-legislation.aspx>.
- [362] National Financial Educators Council. Test your knowledge with the nfec financial foundation test, 2017. <https://www.financialeducatorsCouncil.org/financial-foundation-test/>.
- [363] Boon-Yuen Ng, Atreyi Kankanhalli, and Yunjie Calvin Xu. Studying users’ computer security behavior: A health belief perspective. *Decision Support Systems*, 46(4):815–825, 2009. doi:10.1016/j.dss.2008.11.010.
- [364] Helen Nissenbaum. *Privacy in context: Technology, policy, and the integrity of social life*. Stanford University Press, 2009.
- [365] Patricia A Norberg, Daniel R Horne, and David A Horne. The privacy paradox: Personal information disclosure intentions versus behaviors. *Journal of Consumer Affairs*, 41(1):100–126, 2007. doi:10.1111/j.1745-6606.2006.00070.x.
- [366] Don Norman. *The design of everyday things: Revised and expanded edition*. Constellation, 2013.
- [367] Northnode Inc. Domestic violence training for new staff & volunteers, 2008. [http://www.healthrecovery.org/images/products/34\\_full.pdf](http://www.healthrecovery.org/images/products/34_full.pdf).

- [368] Midas Nouwens, Ilaria Liccardi, Michael Veale, David Karger, and Lalana Kagal. Dark patterns after the GDPR: Scraping consent pop-ups and demonstrating their influence. In *ACM Conference on Human Factors in Computing Systems*, pages 194:1–194:13, 2020. doi:10.1145/3313831.3376321.
- [369] Maggie Oates, Yama Ahmadullah, Abigail Marsh, Chelse Swoopes, Shikun Zhang, Rebecca Balebako, and Lorrie Faith Cranor. Turtles, locks, and bathrooms: Understanding mental models of privacy through illustration. *Proceedings on Privacy Enhancing Technologies*, 2018(4):5–32, 2018. doi:10.1515/popets-2018-0029.
- [370] Jonathan A Obar and Anne Oeldorf-Hirsch. The biggest lie on the internet: Ignoring the privacy policies and terms of service policies of social networking services. *Information, Communication & Society*, 23(1):128–147, 2020. doi:10.1080/1369118X.2018.1486870.
- [371] Sarah O’Brien. What consumers are doing to protect their data a year after huge Equifax breach, 2018. <https://www.cnbc.com/2018/09/07/equifax-anniversary.html>.
- [372] Global Network of Women’s Shelters. Provide information, 2020. <https://www.gnws.org/index.php/women-s-helplines/provide-information>.
- [373] Office of the California Attorney General. California Consumer Privacy Act (CCPA): Proposed Text of Regulations, 2019. <https://www.oag.ca.gov/sites/all/files/agweb/pdfs/privacy/ccpa-proposed-regs.pdf>.
- [374] Office of the California Attorney General. The california privacy rights and enforcement act of 2020, 2019. <https://oag.ca.gov/system/files/initiatives/pdfs/19-0017%20%28Consumer%20Privacy%20%29.pdf>.
- [375] Office of the California Attorney General. California Consumer Privacy Act (CCPA): Final Text of Proposed Regulations, 2020. <https://oag.ca.gov/sites/all/files/agweb/pdfs/privacy/oal-sub-final-text-of-regs.pdf>.
- [376] Office of the California Attorney General. California Consumer Privacy Act (CCPA): First Set of Modified Regulations, 2020. <https://oag.ca.gov/sites/all/files/agweb/pdfs/privacy/ccpa-text-of-mod-clean-020720.pdf>.
- [377] Office of the California Attorney General. California Consumer Privacy Act (CCPA): Fourth Set of Modified Regulations, 2020. <https://oag.ca.gov/sites/all/files/agweb/pdfs/privacy/ccpa-prop-mods-text-of-regs-4th.pdf>.
- [378] Office of the California Attorney General. Attorney General Berra Announces Approval of Additional Regulations That Empower

- Data Privacy Under the California Consumer Privacy Act, 2021. <https://oag.ca.gov/news/press-releases/attorney-general-becerra-announces-approval-additional-regulations-empower-data>.
- [379] Office of the California Attorney General. CCPA Regulations, 2021. <https://oag.ca.gov/privacy/ccpa/regs>.
- [380] Jeremiah Onalapo, Enrico Mariconti, and Gianluca Stringhini. What happens after you are pwnd: Understanding the use of leaked webmail credentials in the wild. In *ACM Internet Measurement Conference*, pages 65–79, 2016. doi:10.1145/2987443.2987475.
- [381] Isabelle Oomen and Ronald Leenes. Privacy risk perceptions and privacy protection strategies. In *Policies and Research in Identity Management*, pages 121–138. Springer, 2008. doi:10.1007/978-0-387-77996-6\_10.
- [382] Stefan Palan and Christian Schitter. Prolific.ac—a subject pool for online experiments. *Journal of Behavioral and Experimental Finance*, 17:22–27, 2018. doi:10.1016/j.jbef.2017.12.004.
- [383] Odysseas Papadimitriou. Identity theft: What it is, how it happens & the best protection, 2018. <https://wallethub.com/edu/identity-theft/17120/>.
- [384] A Parsu Parasuraman, Valarie A Zeithaml, and Leonard L Berry. Serqual: A multiple-item scale for measuring consumer perceptions of service quality. *Journal of Retailing*, 64(1):12–40, 1988.
- [385] Cheul Young Park, Cori Faklaris, Siyan Zhao, Alex Sciuto, Laura Dabbish, and Jason Hong. Share and share alike? An exploration of secure behaviors in romantic relationships. In *Symposium on Usable Privacy and Security*, pages 83–102, 2018. <https://www.usenix.org/system/files/conference/soups2018/soups2018-park.pdf>.
- [386] Yong Jin Park. Do men and women differ in privacy? gendered privacy and (in) equality in the internet. *Computers in Human Behavior*, 50:252–258, 2015. doi:10.1016/j.chb.2015.04.011.
- [387] Sameer Patil, Xinru Page, and Alfred Kobsa. With a little help from my friends: Can social navigation inform interpersonal privacy preferences? In *ACM Conference on Computer Supported Cooperative Work & Social Computing*, pages 391–394, 2011. doi:10.1145/1958824.1958885.
- [388] Sarah Pearman, Jeremy Thomas, Pardis Emami Naeini, Hana Habib, Lujo Bauer, Nicolas Christin, Lorrie Faith Cranor, Serge Egelman, and Alain Forget. Let’s go in for a closer look: Observing passwords in their natural habitat. In *ACM Conference on Computer and Communications Security*, pages 295–310, 2017. doi:10.1145/3133956.3133973.

- [389] Sarah Pearman, Shikun Aerin Zhang, Lujo Bauer, Nicolas Christin, and Lorie Faith Cranor. Why people (don't) use password managers effectively. In *Symposium on Usable Privacy and Security*, pages 319–338, 2019. <https://www.usenix.org/system/files/soups2019-pearman.pdf>.
- [390] Peter Peduzzi, John Concato, Elizabeth Kemper, Theodore R Holford, and Alvan R Feinstein. A simulation study of the number of events per variable in logistic regression analysis. *Journal of Clinical Epidemiology*, 49(12):1373–1379, 1996. doi:10.1016/S0895-4356(96)00236-3.
- [391] Eyal Peer, Laura Brandimarte, Sonam Samat, and Alessandro Acquisti. Beyond the Turk: Alternative platforms for crowdsourcing behavioral research. *Journal of Experimental Social Psychology*, 70:153–163, 2017. doi:10.1016/j.jesp.2017.01.006.
- [392] Eyal Peer, Serge Egelman, Marian Harbach, Nathan Malkin, Arunesh Mathur, and Alisa Frik. Nudge me right: Personalizing online security nudges to people's decision-making styles. *Computers in Human Behavior*, 109:e106347, 2020. doi:10.1016/j.chb.2020.106347.
- [393] Peng Peng, Chao Xu, Luke Quinn, Hang Hu, Bimal Viswanath, and Gang Wang. What happens after you leak your password: Understanding credential sharing on phishing sites. In *ACM Asia Conference on Computer and Communications Security*, pages 181–192, 2019. doi:10.1145/3321705.3329818.
- [394] Supavich Pengnate and Pavlo Antonenko. A multimethod evaluation of online trust and its interaction with metacognitive awareness: an emotional design perspective. *International Journal of Human-Computer Interaction*, 29(9):582–593, 2013. doi:10.1080/10447318.2012.735185.
- [395] Scott R Peppet. Regulating the internet of things: first steps toward managing discrimination, privacy, security and consent. *Texas Law Review*, 93:85–176, 2014. <https://heinonline.org/HOL/LandingPage?handle=hein.journals/tlr93&div=5>.
- [396] Justin Petelka, Yixin Zou, and Florian Schaub. Put your warning where your link is: Improving and evaluating email phishing warnings. In *ACM Conference on Human Factors in Computing Systems*, pages 518:1–518:15, 2019. doi:10.1145/3290605.3300748.
- [397] Rachael M Peters. So you've been notified, now what: The problem with current data-breach notification laws. *Arizona Law Review*, 56(4):1171–1202, 2014. <https://heinonline.org/HOL/LandingPage?handle=hein.journals/arz56&div=40>.
- [398] Katie Petrillo. Protect Your Accounts with Breach Alerts Through LastPass, 2018. <https://blog.lastpass.com/2018/11/protect-your-accounts-with-breach-alerts-through-lastpass/>.

- [399] Emiko Petrosky, Janet M Blair, Carter J Betz, Katherine A Fowler, Shane PD Jack, and Bridget H Lyons. Racial and Ethnic Differences in Homicides of Adult Women and the Role of Intimate Partner Violence—United States, 2003–2014. *Morbidity and Mortality Weekly Report*, 66(28):741–746, 2017. doi:10.15585/mmwr.mm6628a1.
- [400] Bettina Piko. Gender differences and similarities in adolescents’ ways of coping. *The Psychological Record*, 51(2):223–235, 2001. doi:10.1007/BF03395396.
- [401] Mikaela Pitcan, Alice E Marwick, and Danah Boyd. Performing a vanilla self: Respectability politics, social class, and the digital world. *Journal of Computer-Mediated Communication*, 23(3):163–179, 2018. doi:10.1093/jcmc/zmy008.
- [402] Scott Plous. *The psychology of judgment and decision making*. McGraw-Hill Book Company, 1993.
- [403] Ponemon Institute. The aftermath of a data breach: Consumer sentiment. Technical report, Ponemon Institute, 2014. <https://www.ponemon.org/local/upload/file/Consumer%20Study%20on%20Aftermath%20of%20a%20Breach%20FINAL%202.pdf>.
- [404] Rebecca S Portnoff, Linda N Lee, Serge Egelman, Pratyush Mishra, Derek Leung, and David Wagner. Somebody’s watching me? assessing the effectiveness of webcam indicator lights. In *ACM Conference on Human Factors in Computing Systems*, pages 1649–1658, 2015. doi:10.1145/2702123.2702164.
- [405] Privacy Rights Clearinghouse. Data breaches, 2021. <https://privacyrights.org/data-breaches>.
- [406] Jennifer Pullman, Kurt Thomas, and Elie Bursztein. Protect your accounts from data breaches with Password Checkup, 2019. <https://security.googleblog.com/2019/02/protect-your-accounts-from-data.html>.
- [407] Ashwini Rao, Florian Schaub, Norman Sadeh, Alessandro Acquisti, and Ruogu Kang. Expecting the unexpected: Understanding mismatched privacy expectations online. In *Symposium on Usable Privacy and Security*, pages 77–96, 2016. <https://www.usenix.org/system/files/conference/soups2016/soups2016-paper-rao.pdf>.
- [408] Hirak Ray, Flynn Wolf, Ravi Kuber, and Adam J Aviv. Why older adults (don’t) use password managers. In *USENIX Security Symposium*, 2021. [https://www.usenix.org/system/files/sec21summer\\_ray.pdf](https://www.usenix.org/system/files/sec21summer_ray.pdf).
- [409] Elissa M Redmiles. “Should I Worry?” A Cross-Cultural Examination of Account Security Incident Response. In *IEEE Symposium on Security and Privacy*, pages 920–934, 2019. doi:10.1109/SP.2019.00059.

- [410] Elissa M Redmiles, Jessica Bodford, and Lindsay Blackwell. “I Just Want to Feel Safe”: A Diary Study of Safety Perceptions on Social Media. In *International AAAI Conference on Web and Social Media*, volume 13, pages 405–416, 2019. <https://ojs.aaai.org/index.php/ICWSM/article/view/3356/3224>.
- [411] Elissa M Redmiles, Sean Kross, and Michelle L Mazurek. How I learned to be secure: a census-representative survey of security advice sources and behavior. In *ACM Conference on Computer and Communications Security*, pages 666–677, 2016. doi:10.1145/2976749.2978307.
- [412] Elissa M Redmiles, Sean Kross, and Michelle L Mazurek. Where is the digital divide? A survey of security, privacy, and socioeconomics. In *ACM Conference on Human Factors in Computing Systems*, pages 931–936, 2017. doi:10.1145/3025453.3025673.
- [413] Elissa M Redmiles, Sean Kross, and Michelle L Mazurek. How well do my results generalize? comparing security and privacy survey results from mturk, web, and telephone samples. In *IEEE Symposium on Security and Privacy*, pages 1326–1343, 2019. doi:10.1109/SP.2019.00014.
- [414] Elissa M Redmiles, Amelia R Malone, and Michelle L Mazurek. I think they’re trying to tell me something: Advice sources and selection for digital security. In *IEEE Symposium on Security and Privacy*, pages 272–288, 2016. doi:10.1109/SP.2016.24.
- [415] Elissa M Redmiles, Noel Warford, Amritha Jayanti, Aravind Koneru, Sean Kross, Miraida Morales, Rock Stevens, and Michelle L Mazurek. A comprehensive quality evaluation of security and privacy advice on the web. In *USENIX Security Symposium*, pages 89–108, 2020. <https://www.usenix.org/system/files/sec20-redmiles.pdf>.
- [416] Robert W Reeder, Iulia Ion, and Sunny Consolvo. 152 simple steps to stay safe online: security advice for non-tech-savvy users. *IEEE Security & Privacy*, 15(5):55–64, 2017. doi:10.1109/MSP.2017.3681050.
- [417] Joel R Reidenberg, Jaspreet Bhatia, Travis D Breaux, and Thomas B Norton. Ambiguity in privacy policies and the impact of regulation. *The Journal of Legal Studies*, 45(S2):S163–S190, 2016. doi:10.1086/688669.
- [418] Joel R Reidenberg, N Cameron Russell, Alexander J Callen, Sophia Qasir, and Thomas B Norton. Privacy harms and the effectiveness of the notice and choice framework. *Journal of Law and Policy for the Information Society*, 11:485–524, 2015. <https://heinonline.org/HOL/LandingPage?handle=hein.journals/isjlp soc11&div=19>.
- [419] Karen Renaud and Verena Zimmermann. Ethical guidelines for nudging in information security & privacy. *International Journal of Human-Computer Studies*, 120:22–35, 2018. doi:10.1016/j.ijhcs.2018.05.011.

- [420] Alex Reynolds. GDPR matchup: US state data breach laws, 2017. <https://iapp.org/news/a/gdpr-match-up-u-s-state-data-breach-laws/>.
- [421] Markus Riek, Rainer Bohme, and Tyler Moore. Measuring the influence of perceived cybercrime risk on online service avoidance. *IEEE Transactions on Dependable and Secure Computing*, 13(2):261–273, 2015. doi:10.1109/TDSC.2015.2410795.
- [422] Ronald W Rogers. A protection motivation theory of fear appeals and attitude change. *The Journal of Psychology*, 91(1):93–114, 1975. doi:10.1080/00223980.1975.9915803.
- [423] Ronald W Rogers. Cognitive and psychological processes in fear appeals and attitude change: A revised theory of protection motivation. In John T Cacioppo and Richard E Petty, editors, *Social psychophysiology: A sourcebook*, chapter 6, pages 153–176. Guilford Press, 1983.
- [424] Sasha Romanosky. Examining the costs and causes of cyber incidents. *Journal of Cybersecurity*, 2(2):121–135, 2016. doi:10.1093/cybsec/tyw001.
- [425] Sasha Romanosky, David Hoffman, and Alessandro Acquisti. Empirical analysis of data breach litigation. *Journal of Empirical Legal Studies*, 11(1):74–104, 2014. doi:10.1111/jels.12035.
- [426] Sasha Romanosky, Rahul Telang, and Alessandro Acquisti. Do data breach disclosure laws reduce identity theft? *Journal of Policy Analysis and Management*, 30(2):256–286, 2011. doi:10.1002/pam.20567.
- [427] Matthew Rosenberg, Nicholas Confessore, and Carole Cadwaladr. How trump consultants exploited the facebook data of millions, 2018. <https://www.nytimes.com/2018/03/17/us/politics/cambridge-analytica-trump-campaign.html>.
- [428] Arianna Rossi and Monica Palmirani. Dapis: a data protection icon set to improve information transparency under the gdpr. *Knowledge of the Law in the Big Data Age*, 252:181–195, 2019. [http://gdprbydesign.cirsfid.unibo.it/wp-content/uploads/2019/01/report\\_DaPIS\\_jan19.pdf](http://gdprbydesign.cirsfid.unibo.it/wp-content/uploads/2019/01/report_DaPIS_jan19.pdf).
- [429] Susan Roth and Lawrence J Cohen. Approach, avoidance, and coping with stress. *American Psychologist*, 41(7):813–819, 1986. doi:10.1037/0003-066X.41.7.813.
- [430] Jeffrey N Rouder, Christopher R Engelhardt, Simon McCabe, and Richard D Morey. Model Comparison in ANOVA. *Psychonomic Bulletin & Review*, 23(6):1779–1786, 2016. doi:10.3758/s13423-016-1026-5.
- [431] Kevin Alejandro Roundy, Paula Barmaimon Mendelberg, Nicola Dell, Damon McCoy, Daniel Nissani, Thomas Ristenpart, and Acar Tamersoy. The many



- kinds of creepware used for interpersonal attacks. In *IEEE Symposium on Security and Privacy*, pages 626–643, 2020. doi:10.1109/sp40000.2020.00069.
- [432] Manuel Rudolph, Denis Feth, and Svenja Polst. Why users ignore privacy policies—a survey and intention model for explaining user privacy behavior. In *International Conference on Human-Computer Interaction*, pages 587–598, 2018. doi:10.1007/978-3-319-91238-7\_45.
- [433] Robert AC Ruiter, Loes TE Kessels, Gjalt-Jorn Y Peters, and Gerjo Kok. Sixty years of fear appeal research: Current state of the evidence. *International journal of psychology*, 49(2):63–70, 2014. doi:10.1002/ijop.12042.
- [434] Johnny Saldaña. *The Coding Manual for Qualitative Researchers*. Sage, 2015.
- [435] Shruti Sannon, Billie Sun, and Dan Cosley. Privacy, surveillance, and power in the gig economy. In *ACM Conference on Human Factors in Computing Systems*, pages 619:1–619:15, 2022. doi:10.1145/3491102.3502083.
- [436] David Satcher. *Methods in community-based participatory research for health*. John Wiley & Sons, 2005.
- [437] Florian Schaub, Rebecca Balebako, Adam L Durity, and Lorrie Faith Cranor. A design space for effective privacy notices. In *Symposium On Usable Privacy and Security*, pages 1–17, 2015. <https://www.usenix.org/system/files/conference/soups2015/soups15-paper-schaub.pdf>.
- [438] Florian Schaub and Lorrie Faith Cranor. Usable and useful privacy interfaces. In Travis D Breaux, editor, *An Introduction to Privacy for Technology Professionals*, chapter 5, pages 176–238. International Association of Privacy Professionals, 2020. <https://iapp.org/media/pdf/certification/IAPP-Intro-to-Privacy-for-Tech-Prof-SAMPLE.pdf>.
- [439] Stuart E Schechter, Rachna Dhamija, Andy Ozment, and Ian Fischer. The emperor’s new security indicators. In *IEEE Symposium on Security and Privacy*, pages 51–65, 2007. doi:10.1109/SP.2007.35.
- [440] Jillian R Scheer and V Paul Poteat. Trauma-informed care and health among lgbtq intimate partner violence survivors. *Journal of Interpersonal Violence*, pages 1–23, 2018. doi:10.1177/0886260518820688.
- [441] KlausR Scherer. Profiles of emotion-antecedent appraisal: Testing theoretical predictions across cultures. *Cognition & Emotion*, 11(2):113–150, 1997. doi:10.1080/026999397379962.
- [442] Ari Schlesinger, W Keith Edwards, and Rebecca E Grinter. Intersectional hci: Engaging identity through gender, race, and class. In *ACM Conference on Human Factors in Computing Systems*, pages 5412–5427, 2017. doi:10.1145/3025453.3025766.

- [443] Christophe Olivier Schneble, Bernice Simone Elger, and David Shaw. The cambridge analytica affair and internet-mediated research. *EMBO reports*, 19(8):e46579, 2018. doi:10.15252/embr.201846579.
- [444] Theodor Schnitzler, Shujaat Mirza, Markus Dürmuth, and Christina Pöpper. Sok: Managing longitudinal privacy of publicly shared personal online data. *Proceedings on Privacy Enhancing Technologies (PoPETs)*, 2021(1):229–249, 2021. doi:10.2478/popets-2021-0013.
- [445] Sarita Schoenebeck, Nicole B Ellison, Lindsay Blackwell, Joseph B Bayer, and Emily B Falk. Playful backstalking and serious impression management: How young adults reflect on their past identities on facebook. In *ACM Conference on Computer Supported Cooperative Work & Social Computing*, pages 1475–1487, 2016. doi:10.1145/2818048.2819923.
- [446] Robert Schoshinski. Equifax data breach: Pick free credit monitoring, 2019. <https://www.consumer.ftc.gov/blog/2019/07/equifax-data-breach-pick-free-credit-monitoring>.
- [447] Mahmood Sharif, Kevin Alejandro Roundy, Matteo Dell’Amico, Christopher Gates, Daniel Kats, Lujo Bauer, and Nicolas Christin. A field study of computer-security perceptions using anti-virus customer-support chats. In *ACM Conference on Human Factors in Computing Systems*, pages 78:1–78:12, 2019. doi:10.1145/3290605.3300308.
- [448] Tali Sharot. The optimism bias. *Current biology*, 21(23):941–945, 2011. doi:10.1016/j.cub.2011.10.030.
- [449] Richard Shay, Iulia Ion, Robert W Reeder, and Sunny Consolvo. “My religious aunt asked why I was trying to sell her viagra”: experiences with account hijacking. In *ACM Conference on Human Factors in Computing Systems*, pages 2657–2666, 2014. doi:10.1145/2556288.2557330.
- [450] Richard Shay, Patrick Gage Kelley, Saranga Komanduri, Michelle L Mazurek, Blase Ur, Timothy Vidas, Lujo Bauer, Nicolas Christin, and Lorrie Faith Cranor. Correct horse battery staple: Exploring the usability of system-assigned passphrases. In *Symposium on Usable Privacy and Security*, pages 7:1–7:20, 2012. doi:10.1145/2335356.2335366.
- [451] Richard Shay, Saranga Komanduri, Adam L Durity, Phillip Huh, Michelle L Mazurek, Sean M Segreti, Blase Ur, Lujo Bauer, Nicolas Christin, and Lorrie Faith Cranor. Designing password policies for strength and usability. *ACM Transactions on Information and System Security*, 18(4):13:1–13:34, 2016. doi:10.1145/2891411.
- [452] Kim Bartel Sheehan and Mariea Grubbs Hoy. Flaming, complaining, abstaining: How online users respond to privacy concerns. *Journal of Advertising*, 28(3):37–51, 1999. doi:10.1080/00913367.1999.10673588.

- [453] Paschal Sheeran and Thomas L Webb. The intention–behavior gap. *Social and Personality Psychology Compass*, 10(9):503–518, 2016. doi:10.1111/spc3.12265.
- [454] Steve Sheng, Mandy Holbrook, Ponnurangam Kumaraguru, Lorrie Faith Cranor, and Julie Downs. Who falls for phish? a demographic analysis of phishing susceptibility and effectiveness of interventions. In *ACM Conference on Human Factors in Computing Systems*, pages 373–382, 2010. doi:10.1145/1753326.1753383.
- [455] Ruth Shillair, Shelia R Cotten, Hsin-Yi Sandy Tsai, Saleem Alhabash, Robert LaRose, and Nora J Rifon. Online safety begins with you and me: Convincing internet users to protect themselves. *Computers in Human Behavior*, 48:199–207, 2015. doi:10.1016/j.chb.2015.01.046.
- [456] Jeff Shiner. Finding compromised passwords with 1Password, 2022. <https://blog.1password.com/finding-pwned-passwords-with-1password/>.
- [457] Fatemeh Shirazi and Melanie Volkamer. What deters jane from preventing identification and tracking on the web? In *Workshop on Privacy in the Electronic Society*, pages 107–116, 2014. doi:10.1145/2665943.2665963.
- [458] Irina Shklovski, Scott D Mainwaring, Halla Hrunn Skúladóttir, and Höskuldur Borgthorsson. Leakiness and creepiness in app space: Perceptions of privacy and mobile app use. In *ACM Conference on Human Factors in Computing Systems*, pages 2347–2356, 2014. doi:10.1145/2556288.2557421.
- [459] Hu Shuijing and Jiang Tao. An empirical study on digital privacy risk of senior citizens. In *IEEE International Conference on Robots & Intelligent System*, pages 19–24, 2017. doi:10.1109/ICRIS.2017.13.
- [460] Johanna M Silvennoinen, Tuomo Kujala, and Jussi PP Jokinen. Semantic distance as a critical factor in icon design for in-car infotainment systems. *Applied Ergonomics*, 65:369–381, 2017. doi:10.1016/j.apergo.2017.07.014.
- [461] Lucy Simko. *Humans and Vulnerability During Times of Change: Computer Security Needs, Practices, Challenges, and Opportunities*. PhD thesis, University of Washington, 2022.
- [462] Natasha Singer. When Websites Won’t Take No for an Answer, 2016. [https://www.nytimes.com/2016/05/15/technology/personaltech/when-websites-wont-take-no-for-an-answer.html?mcubz=0&\\_r=0](https://www.nytimes.com/2016/05/15/technology/personaltech/when-websites-wont-take-no-for-an-answer.html?mcubz=0&_r=0).
- [463] Mikko Siponen, M Adam Mahmood, and Seppo Pahnala. Employees’ adherence to information security policies: An exploratory field study. *Information & Management*, 51(2):217–224, 2014. doi:10.1016/j.im.2013.08.006.
- [464] Suzanne M Slattery and Lisa A Goodman. Secondary trauma stress among domestic violence advocates: Workplace risk and protective factors. *Violence Against Women*, 15(11):1358–1379, 2009. doi:10.1177/1077801209347469.

- [465] Manya Sleeper, Rebecca Balebako, Sauvik Das, Amber McConahy, Jason Wiese, and Lorrie Cranor. The post that wasn't: exploring self-censorship on facebook. In *ACM Conference on Computer Supported Cooperative Work & Social Computing*, pages 793–802, 2013. doi:10.1145/2441776.2441865.
- [466] Robert H Sloan and Richard Warner. Beyond notice and choice: Privacy, norms, and consent. *Journal of High Technology Law*, 14(2):370–414, 2014. <https://heinonline.org/HOL/LandingPage?handle=hein.journals/jhtl114&div=11>.
- [467] Paul Slovic. Perception of risk. *Science*, 236(4799):280–285, 1987. doi:10.1126/science.3563507.
- [468] Paul Slovic, Melissa L Finucane, Ellen Peters, and Donald G MacGregor. Risk as analysis and risk as feelings: Some thoughts about affect, reason, risk, and rationality. *Risk Analysis*, 24(2):311–322, 2004. doi:10.1111/j.0272-4332.2004.00433.x.
- [469] Edith Smit, Guda Van Noort, and Hilde Voorveld. Understanding online behavioural advertising: User knowledge, privacy concerns and online coping behaviour in europe. *Computers in Human Behavior*, 32:15–22, 2014. doi:10.1016/j.chb.2013.11.008.
- [470] Aaron Smith. More Than Half of U.S. Adults Trust Law Enforcement to Use Facial Recognition Responsibly. Technical report, Pew Research Center, 2019. [https://www.pewresearch.org/internet/wp-content/uploads/sites/9/2019/09/09.05.19.facial\\_recognition\\_FULLREPORT\\_update.pdf](https://www.pewresearch.org/internet/wp-content/uploads/sites/9/2019/09/09.05.19.facial_recognition_FULLREPORT_update.pdf).
- [471] Rob Sobers. 64% of Americans Don't Know What to Do After a Data Breach — Do You?, 2020. <https://www.varonis.com/blog/data-breach-literacy-survey/>.
- [472] Daniel J Solove. I've got nothing to hide and other misunderstandings of privacy. *San Diego Law Review*, 44(4):745–772, 2007. <https://heinonline.org/HOL/LandingPage?handle=hein.journals/sanlr44&div=40>.
- [473] Daniel J Solove. Privacy self-management and the consent dilemma. *Harvard Law Review*, 126(7):1880–1903, 2013. <https://heinonline.org/HOL/LandingPage?handle=hein.journals/hlr126&div=87>.
- [474] Daniel J Solove. The myth of the privacy paradox. *George Washington Law Review*, 89(1):1–51, 2021. <https://heinonline.org/HOL/LandingPage?handle=hein.journals/gwlr89&div=4>.
- [475] Daniel J Solove and Danielle Keats Citron. Risk and anxiety: A theory of data-breach harms. *Texas Law Review*, 96:737–786, 2017. <https://heinonline.org/HOL/LandingPage?handle=hein.journals/tlr96&div=28>.

- [476] Daniel J Solove and Woodrow Hartzog. *Breached!: Why Data Security Law Fails and How to Improve it*. Oxford University Press, 2022.
- [477] Jai-Yeol Son and Sung Kim. Internet users’ information privacy-protective responses: A taxonomy and a nomological model. *MIS Quarterly*, 32(3):503–529, 2008. doi:10.2307/25148854.
- [478] Cynthia Southworth, Jerry Finn, Shawndell Dawson, Cynthia Fraser, and Sarah Tucker. Intimate partner violence, technology, and stalking. *Violence Against Women*, 13(8):842–856, 2007. doi:10.1177/1077801207302045.
- [479] The Coalition Against Stalkerware. Founding partners, 2021. <https://stopstalkerware.org/partners/>.
- [480] Stanford Legal Design Lab. Icons for legal help, 2021. <https://betterinternet.law.stanford.edu/design-guide/icons-for-legal-help/>.
- [481] State of California. California Civil Code 1798.82, 2021. [https://leginfo.legislature.ca.gov/faces/codes\\_displaySection.xhtml?lawCode=CIV&sectionNum=1798.82](https://leginfo.legislature.ca.gov/faces/codes_displaySection.xhtml?lawCode=CIV&sectionNum=1798.82).
- [482] Gregg Steinhafel. A message from ceo gregg steinhafel about target’s payment card issues, 2013. <https://corporate.target.com/article/2013/12/target-ceo-gregg-steinhafel-message>.
- [483] Elizabeth Stobert and Robert Biddle. The password life cycle: user behaviour in managing passwords. In *Symposium On Usable Privacy and Security*, pages 243–255, 2014. <https://www.usenix.org/system/files/conference/soups2014/soups14-paper-stobert.pdf>.
- [484] Peter Story, Daniel Smullen, Alessandro Acquisti, Lorrie Faith Cranor, Norman Sadeh, and Florian Schaub. From intent to action: Nudging users towards secure mobile payments. In *Symposium on Usable Privacy and Security*, pages 379–415, 2020. <https://www.usenix.org/system/files/soups2020-story.pdf>.
- [485] Peter Story, Daniel Smullen, Rex Chen, Yaxing Yao, Alessandro Acquisti, Lorrie Faith Cranor, Norman Sadeh, and Florian Schaub. Increasing adoption of tor browser using informational and planning nudges. *Proceedings on Privacy Enhancing Technologies*, 2022(2):152–183, 2022. doi:10.2478/popets-2022-0040.
- [486] Fred Stutzman and Jacob Kramer-Duffield. Friends only: examining a privacy-enhancing behavior in facebook. In *ACM Conference on Human Factors in Computing Systems*, pages 1553–1562, 2010. doi:10.1145/1753326.1753559.
- [487] Qizhang Sun, Martijn C Willemsen, and Bart P Knijnenburg. Unpacking the intention-behavior gap in privacy decision making for the internet of

- things (IoT) using aspect listing. *Computers & Security*, 97:e101924, 2020. doi:10.1016/j.cose.2020.101924.
- [488] Madiha Tabassum, Jess Kropczynski, Pamela Wisniewski, and Heather Richter Lipford. Smart home beyond the home: A case for community-based access control. In *ACM Conference on Human Factors in Computing Systems*, pages 128:1–128:12, 2020. doi:10.1145/3313831.3376255.
- [489] Jenny Tang, Eleanor Birrell, and Ada Lerner. Replication: How Well Do My Results Generalize Now? The External Validity of Online Privacy and Security Surveys. In *Symposium on Usable Privacy and Security*, pages 367–385, 2022. <https://www.usenix.org/system/files/soups2022-tang.pdf>.
- [490] Richard H Thaler and Cass R Sunstein. *Nudge: Improving decisions about health, wealth, and happiness*. Penguin, 2008.
- [491] The Consumer Financial Protection Bureau. What is a credit report?, 2017. <https://www.consumerfinance.gov/ask-cfpb/what-is-a-credit-report-en-309/>.
- [492] The Digital Advertising Alliance. Self-regulatory principles for online behavioral advertising, 2009. <http://digitaladvertisingalliance.org/principles>.
- [493] The Digital Advertising Alliance. Campaign informs consumers about interest-based advertising and encourages them to take control of their online privacy, 2012. <https://digitaladvertisingalliance.org/press-release/digital-advertising-alliance-daa-announces-%E2%80%98your-adchoices%E2%80%99-consumer-education>.
- [494] The Digital Advertising Alliance. License the Privacy Rights Icon, 2021. <https://digitaladvertisingalliance.org/license-pricon>.
- [495] The Digital Advertising Alliance. Your AdChoices, 2021. <https://youradchoices.com/>.
- [496] The European Union Agency for Cybersecurity. Privacy enhancing technologies, 2021. <https://www.enisa.europa.eu/topics/data-protection/privacy-enhancing-technologies>.
- [497] The Federal Trade Commission. CAN-SPAM Act: A compliance guide for business, 2009. <https://www.ftc.gov/tips-advice/business-center/guidance/can-spam-act-compliance-guide-business>.
- [498] The Federal Trade Commission. Final model privacy form under the Gramm-Leach-Bliley Act, 2009. [https://www.ftc.gov/system/files/documents/rules/privacy-consumer-financial-information-financial-privacy-rule/privacymodelform\\_optout.pdf](https://www.ftc.gov/system/files/documents/rules/privacy-consumer-financial-information-financial-privacy-rule/privacymodelform_optout.pdf).

- [499] The Federal Trade Commission. .com Disclosures: How to Make Effective Disclosures in Digital Advertising, 2013. <https://www.ftc.gov/sites/default/files/attachments/press-releases/ftc-staff-revises-online-advertising-disclosure-guidelines/130312dotcomdisclosures.pdf>.
- [500] The Federal Trade Commission. Children’s online privacy protection rule: A six-step compliance plan for your business, 2017. <https://www.ftc.gov/tips-advice/business-center/guidance/childrens-online-privacy-protection-rule-six-step-compliance>.
- [501] The Federal Trade Commission. The Equifax Data Breach: What to Do, 2017. <https://www.consumer.ftc.gov/blog/2017/09/equifax-data-breach-what-do>.
- [502] The Federal Trade Commission. Starting Today, New Federal Law Allows Consumers to Place Free Credit Freezes And Yearlong Fraud Alerts, 2018. <https://www.ftc.gov/news-events/press-releases/2018/09/starting-today-new-law-allows-consumers-place-free-credit-freezes>.
- [503] The Federal Trade Commission. The consumer sentinel network data book 2018, 2019. [https://www.ftc.gov/system/files/documents/reports/consumer-sentinel-network-data-book-2018/consumer\\_sentinel\\_network\\_data\\_book\\_2018\\_0.pdf](https://www.ftc.gov/system/files/documents/reports/consumer-sentinel-network-data-book-2018/consumer_sentinel_network_data_book_2018_0.pdf).
- [504] The Federal Trade Commission. FTC brings first case against developers of “stalking” apps, 2019. <https://www.ftc.gov/news-events/press-releases/2019/10/ftc-brings-first-case-against-developers-stalking-apps>.
- [505] The Federal Trade Commission. Equifax Data Breach Settlement, 2020. <https://www.ftc.gov/enforcement/cases-proceedings/refunds/equifax-data-breach-settlement>.
- [506] The Federal Trade Commission. When information is lost or exposed, 2020. <https://www.identitytheft.gov/databreach>.
- [507] The Federal Trade Commission. Identity theft, 2021. <https://www.consumer.ftc.gov/topics/identity-theft>.
- [508] The Interactive Advertising Bureau Technology Laboratory. GDPR transparency and consent framework, 2019. <https://iabtechlab.com/standards/gdpr-transparency-and-consent-framework/>.
- [509] The National Coalition Against Domestic Violence. Statistics, 2021. <https://ncadv.org/statistics>.
- [510] The Network Advertising Initiative. NAI code of conduct, 2018. [https://www.networkadvertising.org/sites/default/files/nai\\_code2018.pdf](https://www.networkadvertising.org/sites/default/files/nai_code2018.pdf).

- [511] The Online Trust Alliance. Email marketing & unsubscribe audit, 2018. <https://www.internetsociety.org/wp-content/uploads/2019/04/2018-email-unsubscribe-report.pdf>.
- [512] The U.S. Government Printing Office. Health insurance portability and accountability act of 1996, public law 104-191, 1996. <https://www.govinfo.gov/content/pkg/PLAW-104publ191/pdf/PLAW-104publ191.pdf>.
- [513] Kurt Thomas, Devdatta Akhawe, Michael Bailey, Dan Boneh, Elie Bursztein, Sunny Consolvo, Nicola Dell, Zakir Durumeric, Patrick Gage Kelley, Deepak Kumar, et al. Sok: Hate, harassment, and the changing landscape of online abuse. In *IEEE Symposium on Security and Privacy*, pages 247–267, 2021. doi:10.1109/SP40001.2021.00028.
- [514] Kurt Thomas, Frank Li, Ali Zand, Jacob Barrett, Juri Ranieri, Luca Invernizzi, Yarik Markov, Oxana Comanescu, Vijay Eranti, Angelika Moscicki, et al. Data breaches, phishing, or malware? understanding the risks of stolen credentials. In *ACM Conference on Computer and Communications Security*, pages 1421–1434, 2017. doi:10.1145/3133956.3134067.
- [515] Kurt Thomas, Jennifer Pullman, Kevin Yeo, Ananth Raghunathan, Patrick Gage Kelley, Luca Invernizzi, Borbala Benko, Tadek Pietraszek, Sarvar Patel, Dan Boneh, et al. Protecting accounts from credential stuffing with password breach alerting. In *USENIX Security Symposium*, pages 1556–1571, 2019. <https://www.usenix.org/system/files/sec19-thomas.pdf>.
- [516] The National Network to End Domestic Violence. Power & control wheel: On technology & abuse, 2020. [https://safechatsv.org/wp-content/uploads/2016/07/NNEDV\\_TechPowerControlWheel\\_Aug08.pdf](https://safechatsv.org/wp-content/uploads/2016/07/NNEDV_TechPowerControlWheel_Aug08.pdf).
- [517] The National Network to End Domestic Violence. Technology safety, 2021. <https://www.techsafety.org/>.
- [518] Jenna Torluemke and Christine Kim. New NortonLifeLock Survey Reveals the Most Common Habits on How Exes and Partners Spy on Each Other, 2020. <https://www.businesswire.com/news/home/20200212005192/en/>.
- [519] TrustArc. EDAA rolls out pan-European consumer education campaign on OBA, 2013. <https://trustarc.com/blog/2013/06/13/edaa-rolls-out-pan-european-consumer-education-campaign-on-oba/>.
- [520] Emily Tseng, Rosanna Bellini, Nora McDonald, Matan Danos, Rachel Greenstadt, Damon McCoy, Nicola Dell, and Thomas Ristenpart. The tools and tactics used in intimate partner surveillance: An analysis of online infidelity forums. In *USENIX Security Symposium*, pages 1893–1909, 2020. <https://www.usenix.org/system/files/sec20-tseng.pdf>.



- [521] Emily Tseng, Mehrnaz Sabet, Rosanna Bellini, Harkiran Kaur Sodhi, Thomas Ristenpart, and Nicola Dell. Care infrastructures for digital security in intimate partner violence. In *ACM Conference on Human Factors in Computing Systems*, pages 123:1–123:20, 2022. doi:10.1145/3491102.3502038.
- [522] Dana Turjeman and Fred M Feinberg. When the data are out: Measuring behavioral changes following a data breach. *SSRN*, 2019. doi:10.2139/ssrn.3427254.
- [523] Joseph Turow, Michael Hennessy, and Nora Draper. The tradeoff fallacy: How marketers are misrepresenting American consumers and opening them up to exploitation. *SSRN*, 2015. doi:10.2139/ssrn.2820060.
- [524] Amos Tversky and Daniel Kahneman. Availability: A heuristic for judging frequency and probability. *Cognitive Psychology*, 5(2):207–232, 1973. doi:10.1016/0010-0285(73)90033-9.
- [525] Amos Tversky and Daniel Kahneman. Judgment under uncertainty: Heuristics and biases. *Science*, 185(4157):1124–1131, 1974. doi:10.1126/science.185.4157.1124.
- [526] Amos Tversky and Daniel Kahneman. The framing of decisions and the psychology of choice. *Science*, 211(4481):453–458, 1981. doi:10.1126/science.7455683.
- [527] United States Census Bureau. Annual Estimates of the Resident Population by Single Year of Age and Sex for the United States: April 1, 2010 to July 1, 2018, 2018. <https://www.census.gov/data/tables/time-series/demo/popest/2010s-national-detail.html>.
- [528] United States Census Bureau. Educational Attainment of the Population 18 Years and Over, by Age, Sex, Race, and Hispanic Origin: 2018, 2018. <https://www.census.gov/data/tables/2018/demo/education-attainment/cps-detailed-tables.html>.
- [529] United States Census Bureau. Households by Total Money Income, Race, and Hispanic Origin of Householder: 1967 to 2018, 2018. <https://www.census.gov/library/publications/2019/demo/p60-266.html>.
- [530] United States Census Bureau. Population by age and sex: 2018, 2018. <https://www.census.gov/data/tables/2018/demo/age-and-sex/2018-age-sex-composition.html>.
- [531] United States Census Bureau. Age and sex tables, 2022. [https://www.census.gov/topics/population/age-and-sex/data/tables.2019.List\\_897222059.html](https://www.census.gov/topics/population/age-and-sex/data/tables.2019.List_897222059.html).
- [532] United States Census Bureau. Census bureau releases new educational attainment data, 2022. <https://www.census.gov/newsroom/press-releases/2022/educational-attainment.html>.

- [533] United States Census Bureau. Selected characteristics of households by total money income, 2022. <https://www.census.gov/data/tables/time-series/demo/income-poverty/cps-hinc/hinc-01.html>.
- [534] United States Census Bureau. U.s. census bureau quickfacts, 2022. <https://www.census.gov/quickfacts/fact/table/US/LFE046220>.
- [535] United States Congress. Gramm-leach-bliley act, public law 106-102, 1999. <https://www.govinfo.gov/content/pkg/PLAW-106publ102/pdf/PLAW-106publ102.pdf>.
- [536] Blase Ur, Patrick Gage Kelley, Saranga Komanduri, Joel Lee, Michael Maass, Michelle L Mazurek, Timothy Passaro, Richard Shay, Timothy Vidas, Lujo Bauer, et al. How does your password measure up? the effect of strength meters on password creation. In *USENIX Security Symposium*, pages 65–80, 2012. <https://www.usenix.org/system/files/conference/usenixsecurity12/sec12-final209.pdf>.
- [537] Blase Ur, Pedro Giovanni Leon, Lorrie Faith Cranor, Richard Shay, and Yang Wang. Smart, useful, scary, creepy: perceptions of online behavioral advertising. In *Symposium on Usable Privacy and Security*, pages 4:1–4:15, 2012. doi:10.1145/2335356.2335362.
- [538] Blase Ur, Fumiko Noma, Jonathan Bees, Sean M Segreti, Richard Shay, Lujo Bauer, Nicolas Christin, and Lorrie Faith Cranor. “I Added ’!’ at the End to Make It Secure”: Observing Password Creation in the Lab. In *Symposium on Usable Privacy and Security*, pages 123–140, 2015. <https://www.usenix.org/system/files/conference/soups2015/soups15-paper-ur.pdf>.
- [539] Christine Utz, Martin Degeling, Sascha Fahl, Florian Schaub, and Thorsten Holz. (Un) informed Consent: Studying GDPR Consent Notices in the Field. In *ACM Conference on Computer and Communications Security*, pages 973–990, 2019. doi:10.1145/3319535.3354212.
- [540] Matthew W Vail, Julia B Earp, and Annie I Antón. An empirical study of consumer perceptions and comprehension of web site privacy policies. *IEEE Transactions on Engineering Management*, 55(3):442–454, 2008. doi:10.1109/TEM.2008.922634.
- [541] Zarina Vakhitova, Julianne Webster, Clair Alston-Knox, Danielle Reynald, and Michael Townsley. Offender–victim relationship and offender motivation in the context of indirect cyber abuse: A mixed-method exploratory analysis. *International Review of Victimology*, 24(3):347–366, 2018. doi:10.1177/0269758017743073.
- [542] René van Bavel, Nuria Rodríguez-Priego, José Vila, and Pam Briggs. Using protection motivation theory in the design of nudges to improve online security

- behavior. *International Journal of Human-Computer Studies*, 123:29–39, 2019. doi:10.1016/j.ijhcs.2018.11.003.
- [543] Ellen Van Gool, Joris Van Ouytsel, Koen Ponnet, and Michel Walrave. To share or not to share? adolescents’ self-disclosure about peer relationships on facebook: An application of the prototype willingness model. *Computers in Human Behavior*, 44:230–239, 2015. doi:10.1016/j.chb.2014.11.036.
- [544] Anthony Vance, David Eargle, Kirk Ouimet, and Detmar Straub. Enhancing password security through interactive fear appeals: A web-based field experiment. In *Hawaii International Conference on System Sciences*, pages 2988–2997, 2013. doi:10.1109/HICSS.2013.196.
- [545] Ashlee Vance. If Your Password Is 123456, Just Make It HackMe, 2010. <https://www.nytimes.com/2010/01/21/technology/21password.html>.
- [546] Kami E Vaniea, Emilee Rader, and Rick Wash. Betrayed by updates: how negative experiences affect future security. In *ACM Conference on Human Factors in Computing Systems*, pages 2671–2674, 2014. doi:10.1145/2556288.2557275.
- [547] Aditya Vashistha, Richard Anderson, and Shrirang Mare. Examining security and privacy research in developing regions. In *ACM Conference on Computing and Sustainable Societies*, pages 25:1–25:14, 2018. doi:10.1145/3209811.3209818.
- [548] Jennifer R Veltsos. An analysis of data breach notifications as negative news. *Business Communication Quarterly*, 75(2):192–207, 2012. doi:10.1177/1080569912443081.
- [549] Rafael Veras, Christopher Collins, and Julie Thorpe. On semantic patterns of passwords and their security impact. In *Network and Distributed System Security Symposium*, pages 27:1–27:16, 2014. <https://www.ndss-symposium.org/ndss2014/programme/semantic-patterns-passwords-and-their-security-impact/>.
- [550] Tony Vila, Rachel Greenstadt, and David Molnar. Why we can’t be bothered to read privacy policies: Models of privacy economics as a lemons market. In *ACM International Conference on Electronic Commerce*, pages 403–407, 2003. doi:10.1145/948005.948057.
- [551] Jessica Vitak, Yuting Liao, Mega Subramaniam, and Priya Kumar. ‘I Knew It Was Too Good to Be True’: The Challenges Economically Disadvantaged Internet Users Face in Assessing Trustworthiness, Avoiding Scams, and Developing Self-Efficacy Online. *Proceedings of the ACM on human-computer interaction*, 2(CSCW):176:1–176:25, 2018. doi:10.1145/3274445.
- [552] Silvio Waisbord. Mob censorship: Online harassment of US journalists in times of digital hate and populism. *Digital Journalism*, 8(8):1030–1046, 2020. doi:10.1080/21670811.2020.1818111.

- [553] Robin Wakefield. The influence of user affect in online information disclosure. *The Journal of Strategic Information Systems*, 22(2):157–174, 2013. doi:10.1016/j.jsis.2013.01.003.
- [554] Ari Ezra Waldman. Privacy, notice, and design. *Stanford Technology Law Review*, 21(1):74–127, 2018. <https://heinonline.org/HOL/LandingPage?handle=hein.journals/stantl1r21&div=5>.
- [555] Ari Ezra Waldman. *Industry unbound: The inside story of privacy, data, and corporate Power*. Cambridge University Press, 2021.
- [556] Kent Walker. The costs of privacy. *Harvard Journal of Law and Public Policy*, 25(1):87–127, 2001. <https://heinonline.org/HOL/LandingPage?handle=hein.journals/hjlpp25&div=17>.
- [557] Ding Wang, Zijian Zhang, Ping Wang, Jeff Yan, and Xinyi Huang. Targeted online password guessing: An underestimated threat. In *ACM Conference on Computer and Communications Security*, pages 1242–1254, 2016. doi:10.1145/2976749.2978339.
- [558] Na Wang, Bo Zhang, Bin Liu, and Hongxia Jin. Investigating effects of control and ads awareness on android users’ privacy behaviors and perceptions. In *ACM International Conference on Human-Computer Interaction with Mobile Devices and Services*, pages 373–382, 2015. doi:10.1145/2785830.2785845.
- [559] Yang Wang. Inclusive security and privacy. *IEEE Security & Privacy*, 16(4):82–87, 2018. doi:10.1109/MSP.2018.3111237.
- [560] Yang Wang, Gregory Norcie, Saranga Komanduri, Alessandro Acquisti, Pedro Giovanni Leon, and Lorrie Faith Cranor. “I regretted the minute I pressed share”: A qualitative study of regrets on Facebook. In *Symposium on Usable Privacy and Security*, pages 10:1–10:16, 2011. doi:10.1145/2078827.2078841.
- [561] Logan Warberg, Alessandro Acquisti, and Douglas Sicker. Can privacy nudges be tailored to individuals’ decision making and personality traits? In *Workshop on Privacy in the Electronic Society*, pages 175–197, 2019. doi:10.1145/3338498.3358656.
- [562] Noel Warford, Tara Matthews, Kaitlyn Yang, Omer Akgul, Sunny Consolvo, Patrick Gage Kelley, Nathan Malkin, Michelle L Mazurek, Manya Sleeper, and Kurt Thomas. Sok: A framework for unifying at-risk user research. In *IEEE Symposium on Security and Privacy*, pages 2344–2360, 2022. doi:10.1109/SP46214.2022.9833643.
- [563] Rick Wash. Folk models of home computer security. In *Symposium on Usable Privacy and Security*, pages 11:1–11:16, 2010. doi:10.1145/1837110.1837125.

- [564] Rick Wash and Molly M Cooper. Who provides phishing training?: Facts, stories, and people like me. In *ACM Conference on Human Factors in Computing Systems*, pages 492:1–492:12, 2018. doi:10.1145/3173574.3174066.
- [565] Rick Wash and Emilee Rader. Too Much Knowledge? Security Beliefs and Protective Behaviors Among United States Internet Users. In *Symposium On Usable Privacy and Security*, pages 309–325, 2015. <https://www.usenix.org/system/files/conference/soups2015/soups15-paper-wash.pdf>.
- [566] Rick Wash, Emilee Rader, and Chris Fennell. Can people self-report security accurately? agreement between self-report and behavioral measures. In *ACM Conference on Human Factors in Computing Systems*, pages 2228–2232, 2017. doi:10.1145/3025453.3025911.
- [567] Rick Wash, Emilee Rader, Kami Vaniea, and Michelle Rizor. Out of the loop: How automated software updates cause unintended security consequences. In *Symposium On Usable Privacy and Security*, pages 89–104, 2014. <https://www.usenix.org/system/files/soups14-paper-wash.pdf>.
- [568] Hue Watson, Eyitemi Moju-Igbene, Akanksha Kumari, and Sauvik Das. “We Hold Each Other Accountable”: Unpacking How Social Groups Approach Cybersecurity and Privacy Together. In *ACM Conference on Human Factors in Computing Systems*, pages 478:1–478:12, 2020. doi:10.1145/3313831.3376605.
- [569] Lynette K Watts, Jessyca Wagner, Benito Velasquez, and Phyllis I Behrens. Cyberbullying in higher education: A literature review. *Computers in Human Behavior*, 69:268–274, 2017. doi:10.1016/j.chb.2016.12.038.
- [570] Thomas L Webb and Paschal Sheeran. Does changing behavioral intentions engender behavior change? A meta-analysis of the experimental evidence. *Psychological Bulletin*, 132(2):249–268, 2006. doi:10.1037/0033-2909.132.2.249.
- [571] Ashley Weese and Dana Peiffer. Customer service: Then and now. In *ACM Conference on User services*, pages 35–38, 2013. doi:10.1145/2504776.2504804.
- [572] Miranda Wei, Pardis Emami-Naeini, Franziska Roesner, and Tadayoshi Kohno. Skilled or Gullible? Gender Stereotypes Related to Computer Security and Privacy. In *IEEE Symposium on Security and Privacy*, 2023. <https://mirandawei.com/assets/sp23.pdf>.
- [573] Neil D Weinstein. Unrealistic optimism about future life events. *Journal of Personality and Social Psychology*, 39(5):806–820, 1980. doi:10.1037/0022-3514.39.5.806.
- [574] Neil D Weinstein. Testing four competing theories of health-protective behavior. *Health psychology*, 12(4):324, 1993.

- [575] Herb Weisbaum. Dark web monitoring: Useful way to fight identity theft or marketing hype?, 2019. <https://www.nbcnews.com/better/lifestyle/dark-web-monitoring-useful-way-fight-identity-theft-or-marketing-ncna992351>.
- [576] Alan F Westin. *Privacy and Freedom*. Atheneum, New York, 1967.
- [577] Daniel Lowe Wheeler. zxcvbn: Low-budget password strength estimation. In *USENIX Security Symposium*, pages 157–173, 2016. [https://www.usenix.org/system/files/conference/usenixsecurity16/sec16\\_paper\\_wheeler.pdf](https://www.usenix.org/system/files/conference/usenixsecurity16/sec16_paper_wheeler.pdf).
- [578] Jamie White. The nightmarish experiences of an identity theft victim, 2019. <https://www.lifelock.com/learn-identity-theft-resources-nightmarish-experiences-identity-theft-victim.html>.
- [579] Kimberly A Whitler and Paul W Farris. The impact of cyber attacks on brand image: Why proactive marketing expertise is needed for managing data breaches. *Journal of Advertising Research*, 57(1):3–9, 2017. doi:10.2501/JAR-2017-005.
- [580] Susan Wiedenbeck. The use of icons and labels in an end user application program: an empirical study of learning and retention. *Behaviour & Information Technology*, 18(2):68–82, 1999. doi:10.1080/014492999119129.
- [581] Meredydd Williams, Jason RC Nurse, and Sadie Creese. (Smart) Watch Out! encouraging privacy-protective behavior through interactive games. *International Journal of Human-Computer Studies*, 132:121–137, 2019. doi:10.1016/j.ijhcs.2019.07.012.
- [582] Lauren E Willis. Why not privacy by default. *Berkeley Technology Law Journal*, 29(1):61–134, 2014. <https://heinonline.org/HOL/LandingPage?handle=hein.journals/berktech29&div=5>.
- [583] Kelce Wilson. Data breach notifications may facilitate identity theft, 2018. <https://iapp.org/news/a/data-breach-notifications-may-facilitate-identity-theft/>.
- [584] Kim Witte and Mike Allen. A meta-analysis of fear appeals: Implications for effective public health campaigns. *Health education & behavior*, 27(5):591–615, 2000. doi:10.1177/109019810002700506.
- [585] Delanie Woodlock. The abuse of technology in domestic violence and stalking. *Violence Against Women*, 23(5):584–602, 2017. doi:10.1177/1077801216646277.
- [586] Jason Michael Woodruff. Consequence and likelihood in risk estimation: A matter of balance in uk health and safety risk assessment practice. *Safety Science*, 43(5-6):345–353, 2005. doi:10.1016/j.ssci.2005.07.003.

- [587] Victoria Woollaston. Facebook and netflix reset passwords after data breaches, 2016. <https://www.wired.co.uk/article/facebook-netflix-password-reset>.
- [588] Michael Workman, William H Bommer, and Detmar Straub. Security lapses and the omission of information security measures: A threat control model and empirical test. *Computers in Human Behavior*, 24(6):2799–2816, 2008. doi:10.1016/j.chb.2008.04.005.
- [589] Justin Wu and Daniel Zappala. When is a tree really a truck? exploring mental models of encryption. In *Symposium on Usable Privacy and Security*, pages 395–409, 2018. <https://www.usenix.org/system/files/conference/soups2018/soups2018-wu.pdf>.
- [590] Yuxi Wu, W Keith Edwards, and Sauvik Das. “a reasonable thing to ask for”: Towards a unified voice in privacy collective action. In *ACM Conference on Human Factors in Computing Systems*, pages 32:1–32:17, 2022. doi:10.1145/3491102.3517467.
- [591] Yi Xie, Ke Chen, and Xiaoling Guo. Online anthropomorphism and consumers’ privacy concern: Moderating roles of need for interaction and social exclusion. *Journal of Retailing and Consumer Services*, 55:102119, 2020. doi:10.1016/j.jretconser.2020.102119.
- [592] Yaxing Yao, Davide Lo Re, and Yang Wang. Folk models of online behavioral advertising. In *ACM Conference on Computer Supported Cooperative Work & Social Computing*, pages 1957–1969, 2017. doi:10.1145/2998181.2998316.
- [593] Seounmi Youn. Determinants of online privacy concern and its influence on privacy protection behaviors among young adolescents. *Journal of Consumer affairs*, 43(3):389–418, 2009. doi:10.1111/j.1745-6606.2009.01146.x.
- [594] Alyson Leigh Young and Anabel Quan-Haase. Privacy protection strategies on facebook: The internet privacy paradox revisited. *Information, Communication & Society*, 16(4):479–500, 2013. doi:10.1080/1369118X.2013.777757.
- [595] Georgia Zara and Sarah Gino. Intimate partner violence and its escalation into femicide. *Frontiers in Psychology*, 9:1777:1–1777:11, 2018. doi:10.3389/fpsyg.2018.01777.
- [596] Brahim Zarouali, Karolien Poels, Koen Ponnet, and Michel Walrave. The influence of a descriptive norm label on adolescents’ persuasion knowledge and privacy-protective behavior on social networking sites. *Communication Monographs*, 88(1):1–21, 2020. doi:10.1080/03637751.2020.1809686.
- [597] Brahim Zarouali, Valerie Verdoodt, Michel Walrave, Karolien Poels, Koen Ponnet, and Eva Lievens. Adolescents’ advertising literacy and privacy protection strategies in the context of targeted advertising on social networking

- sites: implications for regulation. *Young Consumers*, 21(3):351–367, 2020. doi:10.1108/YC-04-2020-1122.
- [598] Achim Zeileis, Christian Kleiber, and Simon Jackman. Regression models for count data in R. *Journal of Statistical Software*, 27(8):1–25, 2008. doi:10.18637/jss.v027.i08.
- [599] Eric Zeng, Shrirang Mare, and Franziska Roesner. End user security & privacy concerns with smart homes. In *Symposium on Usable Privacy and Security*, pages 65–80, 2017. <https://www.usenix.org/system/files/conference/soups2017/soups2017-zeng.pdf>.
- [600] Yixin Zou, Shawn Danino, Kaiwen Sun, and Florian Schaub. You ‘might’ be affected: An empirical analysis of readability and usability issues in data breach notifications. In *ACM Conference on Human Factors in Computing Systems*, pages 194:1–194:14, 2019. doi:10.1145/3290605.3300424.
- [601] Yixin Zou, Allison McDonald, Julia Narakornpichit, Nicola Dell, Thomas Ristenpart, Kevin Roundy, Florian Schaub, and Acar Tamersoy. The role of computer security customer support in helping survivors of intimate partner violence. In *USENIX Security Symposium*, pages 429–446, 2021. <https://www.usenix.org/system/files/sec21-zou.pdf>.
- [602] Yixin Zou, Abraham H Mhaidli, Austin McCall, and Florian Schaub. “I’ve Got Nothing to Lose”: Consumers’ Risk Perceptions and Protective Actions after the Equifax Data Breach. In *Symposium on Usable Privacy and Security*, pages 197–216, 2018. <https://www.usenix.org/system/files/conference/soups2018/soups2018-zou.pdf>.
- [603] Yixin Zou, Kevin Roundy, Acar Tamersoy, Saurabh Shintre, Johann Rorturier, and Florian Schaub. Examining the adoption and abandonment of security, privacy, and identity theft protection practices. In *ACM Conference on Human Factors in Computing Systems*, pages 443:1–443:15, 2020. doi:10.1145/3313831.3376570.
- [604] Yixin Zou and Florian Schaub. Beyond mandatory: Making data breach notifications useful for consumers. *IEEE Security & Privacy*, 17(2):67–72, 2019. doi:10.1109/MSEC.2019.2897834.
- [605] Shoshana Zuboff. Big other: surveillance capitalism and the prospects of an information civilization. *Journal of Information Technology*, 30(1):75–89, 2015. doi:10.1057/jit.2015.5.
- [606] Shoshana Zuboff. *The age of surveillance capitalism: The fight for a human future at the new frontier of power*. Profile books, 2019.