

Understanding People's Diverse Privacy Attitudes: Notification, Control, and Regulatory Implications

CMU-LTI-23-002

May 6, 2023

Shikun Zhang

Language Technologies Institute
School of Computer Science
Carnegie Mellon University
Pittsburgh, PA, USA 15213
www.lti.cs.cmu.edu

Thesis Committee:

Norman Sadeh (Chair), Carnegie Mellon University
Alessandro Acquisti, Carnegie Mellon University
Lorrie Faith Cranor, Carnegie Mellon University
Helen Nissenbaum, Cornell Tech

*Submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy
In Language and Information Technologies*

Copyright © 2023 Shikun Zhang

Keywords: Privacy, usable privacy and security, human-computer interaction, contextual integrity, artificial intelligence, mobile app privacy, Internet of Things privacy, privacy assistants, responsible AI, ML for privacy, COVID vaccination certificates.

Abstract

With the broad adoption of smartphones, the Internet of Things (IoT) and artificial intelligence (AI) technologies, people are contributing to the generation of increasingly rich and sensitive digital footprints as they go about their daily lives. The privacy risks associated with the large and diverse amounts of data collected by these new technologies are compounded by increasingly widespread data sharing and data mining practices. In response to these developments, new privacy regulations have been introduced, such as Europe's General Data Protection Regulation (GDPR) and the California Consumer Privacy Act (CCPA). These regulations aim to increase transparency and control over the collection and use of one's personal data, yet they have also inadvertently increased user burden when it comes to managing one's privacy. In the United States, the prevailing legal framework for privacy revolves around the concept of "Notice and Choice." Notifying data subjects about all relevant data collection practices and empowering them to effectively exercise control over these practices in accordance with applicable regulations has become highly impractical. The amount of time and effort needed for a user to read all privacy policies and configure all privacy settings is unrealistically high.

This dissertation explores the diversity of people's privacy attitudes across contexts associated with the recent introduction of new technologies. Specifically, we look at (1) new data collection and use scenarios associated with the recent deployment of video analytics technologies across an increasingly broad range of contexts, (2) the privacy challenges arising from the proposed adoption of COVID-19 vaccination mandates and associated vaccination certificates, and (3) the effectiveness of mobile app privacy labels to inform mobile users about the data collection and use practices of mobile apps. Work presented herein is informed by the Contextual Integrity framework, which identifies key contextual parameters influencing people's privacy expectations and preferences. Through a collection of user studies, this thesis aims to shed light on the diversity of people's privacy attitudes in these different contexts and the challenges they give rise to. This includes looking at the complexity of informing people about the data practices associated with a representative set of video analytics scenarios, people's perception of privacy trade-offs associated with COVID-19 vaccination mandates and certificates in different contexts, and finally the challenges associated with the development of mobile app privacy labels capable of effectively addressing people's diverse privacy concerns.

This dissertation illustrates the complexity and diversity of people's privacy expectations and preferences across these different scenarios. It reveals privacy expectations that apply across broad segments of the population as well as differences in expectations among different groups of people. It shows how clustering techniques can be used to develop finer models of people's privacy expectations and preferences. It documents the challenges in reconciling privacy and user burden consideration and suggests possible solutions that range from regulations requiring APIs to communicate privacy decisions, to the use of clustering models to assist users in managing their privacy decisions.

Acknowledgments

I would like to express my sincere gratitude to my advisor Dr. Norman Sadeh and my thesis committee members Dr. Alessandro Acquisti, Dr. Lorrie Faith Cranor, and Dr. Helen Nissenbaum for their invaluable guidance and support through my Ph.D. journey. Other faculty mentors I would like to acknowledge include Dr. Lujo Bauer, Dr. Nicolas Christin, and Dr. Louis-Philippe Morency for their guidance on my research.

I extend my wholehearted appreciation to my colleagues and collaborators Yama Ahmadullah, Dr. Hazim Almuhammedi, Dr. Rebecca Balebako, Rex Chen, Dr. Jessica Colnago, Dr. Anupam Das, Dr. Martin Degeling, Dr. Yihao Feng, Dr. Yuanyuan Feng, Dr. Bin Liu, Zhengzhong Liu, Maggie Oates, Sarah Pearman, Dr. Abhilasha Ravichander, Dr. Florian Schaub, Dr. Yan Shvartzshnaider, Dr. Daniel Smullen, Dr. Peter Story, Chelse Swoopes, Dr. Yaxing Yao, Hongliang Yu, for their collaboration with me, their assistance, and their camaraderie during my years as a grad student. I would also like to thank our lab manager, Ms. Linda Moreci, for her unfailing support and Ms. Stacey Young for all her help.

I am immensely grateful to my parents and my grandparents for their unwavering support and unconditional love that they have shown me throughout my life. I feel incredibly fortunate to have had my best friend and husband by my side to share this arduous journey of pursuing a Ph.D. I am also very lucky to have so many friends in LTI, HCII, S3D (formerly ISR), and CyLab whose companionships have filled my time at CMU with support, joy, and love.

I would also like to acknowledge the funding that supported in part the research reported in this dissertation. This includes funding through a Carnegie Mellon University CyLab Presidential Fellowship. This also includes funding for the “Personalized Privacy Assistant” Project (DARPA/AFRL Brandeis program grant FA8750-15-2-0277 and NSF CNS-1513957 grant), the “Usable Privacy Policy” Project (NSF SaTC CNS-1330596 grant), the “Contextual Integrity: From Theory to Practice” Project (NSF SaTC CNS-1801316 grant), and the “Automatically Answering People’s Privacy Questions” project (NSF SaTC CNS-1914486 grant) and unrestricted grants from Google and Mozilla.

Contents

1	Introduction	5
1.1	Background and Privacy Challenges	5
1.2	Thesis Contributions	7
1.3	Thesis Outline	8
2	Related Work	9
2.1	Usable Privacy Notices and Control Mechanisms	9
2.1.1	Usable and Effective Privacy Notice	9
2.1.2	Usable Privacy Control Mechanisms	10
2.1.3	Designing and Implementing Privacy Assistants	11
2.2	Sampling and Modeling Privacy Preferences	11
2.3	Studying Privacy through Contextual Integrity	12
2.4	Privacy Challenges of Video Analytics	13
3	Understanding Privacy Expectations and Preferences of Video Analytics Technology	15
3.1	Overview	15
3.2	Study Design	16
3.2.1	Experience Sampling Method	16
3.2.2	Selecting Realistic Scenarios	16
3.2.3	Factorial Design	17
3.2.4	Study Protocol and Procedures	17
3.2.5	Ensuring Study Validity	19
3.2.6	Recruitment and Ethics	20
3.3	Participants and Responses	20
3.3.1	Qualitative Data Set and Analysis	23
3.4	Privacy Preferences	24
3.4.1	Study Validity and Benefits of ESM	25
3.4.2	Factors Impacting Privacy Attitudes	26
3.4.3	Attitude Change Between Start and End of the Study	31
3.4.4	Correlation Between Privacy Expectations and Allow/Deny Preferences	35
3.5	Privacy Concerns and Attitudes	35
3.5.1	Impressions of Facial Recognition	35
3.5.2	Beneficial and Concerning Contexts	36

3.5.3	Beneficial and Concerning Entities	38
3.5.4	Concerns About Facial Recognition	39
3.5.5	Perceived Privacy Risks of Facial Recognition	41
3.5.6	Proposed Actions and Responses	44
3.6	Exploring the Development of Predictive Models	46
3.6.1	Feature Selection and Clustering	46
3.6.2	Predictive Power of Cluster Profiles	47
3.6.3	Example of Cluster Profiles	48
3.6.4	Possible Application in the Context of Privacy Assistants	49
3.7	Discussion	50
3.7.1	Limitations	50
3.7.2	Lack of Awareness and Desire for Greater Transparency	50
3.7.3	Privacy Preferences Are Complex and Context-Dependent	51
3.7.4	Implications for the Design of Privacy Assistants	51
3.7.5	Evolving Notification Preferences	52
3.7.6	Combating Inaccuracy and Bias	53
3.7.7	Contextualizing Perceived Privacy Risks	53
3.7.8	Designing Effective Notice and Choice	53
3.8	Summary of Main Contributions	54
4	A Contextual Integrity Analysis of Vaccination Certificates	57
4.1	Overview	57
4.2	Study Methodology	58
4.2.1	CI-Based Vignette Survey	58
4.2.2	Survey Deployment	60
4.2.3	Data Analysis	61
4.2.4	Limitations	62
4.3	Results	63
4.3.1	VCs as <i>de facto</i> passports	64
4.3.2	Examining VC mandate vignettes	65
4.3.3	Examining Scenarios on Re-sharing VC Information	66
4.3.4	Different Views on VCs: Qualitative Analysis	68
4.4	Clustering Analysis	70
4.5	Discussion	71
4.6	Results from an Additional Survey with a Large Sample	72
4.6.1	The COVID States Project	72
4.6.2	Participants and Demographics	73
4.6.3	Acceptance Levels toward Various Usage Scenarios	73
4.6.4	Clustering	77
4.7	Summary of Main Contributions	78

5	Evaluating the Effectiveness of Mobile App Privacy Labels: How Usable are Today’s Mobile App Privacy Labels?	79
5.1	Overview	79
5.2	Method	82
5.2.1	Recruitment and Screening	82
5.2.2	Interview Protocol	82
5.2.3	Interview Design and Piloting	83
5.2.4	Data Analysis	83
5.2.5	Demographics	84
5.2.6	Limitations	85
5.3	Results	85
5.3.1	Perceptions about App Privacy	85
5.3.2	Perceptions of Privacy Labels	87
5.3.3	Misunderstandings of Privacy Labels	89
5.3.4	Suggested Improvements	95
5.4	Discussion	97
5.4.1	Helping Users Comprehend Complex App Privacy Practices	97
5.4.2	Improving Privacy Labels’ Salience	98
5.4.3	Promoting Privacy Labels’ Role in App Privacy Management	98
5.4.4	Reducing User Burden in App Privacy Management	99
5.5	Conclusion	100
6	Evaluating the Effectiveness of Mobile App Privacy Labels: To What Extent do Apple and Google Privacy Labels Address People’s Privacy Questions?	101
6.1	Overview	101
6.2	A Dataset of Privacy Questions Users Have About Mobile Apps	102
6.3	Methodology	102
6.4	Results	104
6.4.1	Question Themes	104
6.4.2	Question Themes Mostly Answered by Labels	105
6.4.3	Implicit Answers	106
6.4.4	Question Themes Not Addressed by Labels	106
6.4.5	A Comparative Summary of iOS and Google Labels	108
6.5	Discussions	109
6.5.1	Limitations	109
6.5.2	Missing Key Information	109
6.5.3	Implicit Answers and Mismatching Mental Models	110
6.5.4	Privacy Question Answering Functionality	111
6.6	Conclusion	112
6.7	Summary of Main Contributions for Chapter 5 and 6	112

7	Conclusion	115
7.1	Summary of Contributions	115
7.2	Ongoing Work for the Label Study	116
7.3	Challenges and Future Work	117
A	Understanding Privacy Expectations and Preferences of Video Analytics Technology	121
A.1	Scenarios	121
A.2	Evening Review	125
A.3	Post Study Survey	125
A.4	Interview Scripts	126
B	A Contextual Integrity Analysis of Vaccination Certificates	129
B.1	Full Survey Text	129
B.2	Full Survey Text	132
C	Usability of iOS App Privacy Labels	141
C.1	Screening Questionnaire	141
C.2	Interview Scripts	142
C.3	Post-Interview Questions	147
C.4	Codebook	148
	Bibliography	151

List of Figures

3.1	Screenshots of the study app and the web survey used for the evening review	21
3.2	Questions shown to participants in situ	24
3.3	Summary of collected responses organized around 16 different purposes	25
3.4	Percent of participants/notifications reporting specific reasons for discomfort. Participants only selected reasons for notifications that they indicated discomfort (N=1,369). N is the used as the denominator to calculate the percent of notifications.	27
3.5	A Sankey diagram shows the change of participants' reported notification preferences before and after the study	31
3.6	Participants' desire to be notified decreases as the study progresses	33
3.7	The notification and allow/deny preferences of all participants in chronological order over the course of the study.	34
3.8	Accuracy and efficiency of models plotted against the number of clusters used to build them.	47
3.9	Privacy profiles associated with a 6-cluster model	48
4.1	Example of first-hand sharing (top) and re-sharing (bottom) of VC information vignette questions with marked CI parameters. Note that, as per CI theory, in the re-sharing template, the sender value does not match the subject, indicating that the sender is not sharing their own information.	59
4.2	CI parameters used for vignettes involving re-sharing VC information	59
4.3	Reported acceptance levels for VC passport vignettes organized by recipients. . . .	63
4.4	Participants' acceptance levels for nine vignettes. The top row displays the averaged response across nine vignettes. The right graph shows a box plot of the ordinal data with the mean marked in orange.	65
4.5	A heat map of the average of all participants' responses under a combination of four CI parameters	69
4.6	Profiles associated with a 7-cluster model.	71
4.7	U.S. map labeled with the number of participants recruited in each state	74
4.8	Darker red shading indicates oversampling and lighter yellow and white shading indicates the number of participants from this state is under-sampled.	74
4.9	Reported acceptance levels for 22 vignettes.	75
4.10	Dendrogram of agglomerative clustering with a cut-off at 9 clusters	76
4.11	Profiles associated with a 9-cluster model.	77

5.1	Screenshots of compact privacy labels of DoorDash and Chipotle in the iOS App Store	80
5.3	Frequency Terminology Illustration	84
7.1	Examples of our privacy label prototypes	117
B.1	An example vaccination certificate shown to survey participants.	130

List of Tables

3.1	Contextual attribute evaluated in the study	18
3.2	Survey participant demographics and respective %	20
3.3	Occupations of survey participants and respective %	22
3.4	Comparison of IUIPC scores of our participants with an MTurk sample	22
3.5	Example quotes from participants' evening reviews explaining their in-situ answers.	28
3.6	Generalized Linear Mixed Model Regression with Logit Link. A positive coefficient(estimate) shows likeliness of participants' to deny a data collection	30
3.7	Correlation matrix of privacy preference variables	35
3.8	Codes from Content Analysis and the Percentages of Participants Who Mentioned Them	38
4.1	CI parameters used for all vignettes involving first-hand VC information sharing . .	60
4.2	Demographics of our study participants $N = 890$	61
4.3	Cumulative Linear Mixed Model Regression. A positive coefficient (estimate) shows participants' decreased acceptance	67
4.4	Demographics of our study participants $N = 10,631$	75
5.1	Demographics of Participants in the Privacy Label Study	84
5.2	Apple's definitions of data linked to you and data not linked to you [12]	91
5.3	Apple's definitions of analytics, app functionality, and product personalization purposes [12]	92
5.4	Apple's definitions of advertising-related purposes [12]	92
5.5	Apple's definition of product interaction [12]	93
5.6	Apple's definitions of browsing and search history [12]	94
5.7	Apple's definitions of terms under "User Content" [12]	95
6.1	Themes identified, types of questions under each theme, and the number of questions under this theme	103
6.2	Sample user questions and corresponding privacy label entry in the iOS and Google Play Stores.	107
6.3	Questions can be answered by Google Play or iOS privacy labels	108
A.1	Scenario text shown to participants.	121

Chapter 1

Introduction

1.1 Background and Privacy Challenges

With the broad adoption of technologies such as smartphones, cameras, and Internet of Things (IoT) systems, an increasing amount of information is collected about us. Smartphone usage continues to grow in the U.S. and around the world [3, 222, 223, 226, 227]. These devices are capable of continuously collecting a broad range of data such as location, audio, video, fitness data, and much more [111]. According to Pew Research, 54% of mobile app users have refrained from using an app, 30% have declined to install an app, and 19% have disabled location tracking on their devices due to privacy concerns [27]. These findings indicate that a significant number of individuals are concerned about the collection and use of their information by the apps on their mobile devices. In the meanwhile, the number of connected IoT devices is expected to reach 15.9 billion in 2030 [224]. These devices and systems themselves increasingly produce data streams that are fed into machine learning algorithms. For instance, video footage is increasingly processed by video analytics functionality, whether it is for face recognition, facial expression recognition, scene recognition, or some other purpose [110, 123, 144, 259]. Data privacy has been a central area of concern surrounding the deployment of IoT technologies [177] especially when these systems rely on the collection and use of personally identifiable information [37, 176, 178]. Much of the data collection and processing in mobile and IoT is taking place without users' knowledge, let alone their consent. Another trend that has accelerated over the past few years revolves around the use of digital technologies and the demand for disclosing information for public safety and public health purposes. This ranges from requests by public authorities to hand over Ring doorbell video footage, to mandates by the government to show proof of vaccination or sharing contact tracing data on smartphones [109, 160, 168, 171]. For instance, vaccination certificates have become a prime example of this phenomenon as the prolonged and devastating COVID-19 pandemic has affected every aspect of people's lives. The collection and use of the information contained in vaccination certificates, such as an individual's ID number, full name, date of birth, gender, nationality, and vaccination records, may not be restricted to its intended context or purposes, especially when without proper policies and technology-backed measures. The digitization and re-purposing of this information pose significant risks, including privacy violations, with the potential of widening

inequalities, and discrimination [23, 41, 99, 140, 232].

In response to some of these developments and challenges, new privacy regulations have been introduced, such as Europe’s General Data Protection Regulation (GDPR), the California Consumer Privacy Act (CCPA), and the California Privacy Rights Act (CPRA), aimed at safeguarding individuals’ privacy and data protection rights. However, despite these efforts, regulations often lag behind technological advancements, making it challenging to keep up with the ever-evolving data collection and processing practices. Sometimes these regulations can also be aspirational and may be difficult to implement using existing technologies. This gap between regulations and technology is further exacerbated by the emergence of new technologies such as artificial intelligence, machine learning, and especially deep learning, which present novel challenges to privacy and data protection. Though existing regulations aim to increase transparency and control over the collection and use of personal data, they have inadvertently also increased user burden associated with managing one’s privacy. More and more details about the data collection, such as the purposes for which data is collected and used or whether the collected data is shared with third parties, need to be communicated to users to help them make informed privacy decisions. This is a positive development, as research has consistently shown that people’s privacy expectations and preferences vary with the purpose for which data is collected and with whom that data might be shared. Yet, providing people with this additional information further increases the amount of time and effort they would have to devote to learn about the data practices of technologies with which they interact. The same is true for privacy controls (e.g., opt-in/opt-out) as well as data subject rights. While regulatory requirements to offer these to users are beneficial to consumers, it is unclear that people actually have the time and motivation to engage with these options and really take advantage of them. Whether it is while browsing the web or interacting with smartphones, users are expected to manage an unrealistically large number of privacy decision [136, 138]. For instance, a typical smartphone user can easily have well over 100 permission settings to configure on their smartphone [111, 137, 225]. In IoT environments, users are often unaware of the presence of multiple sensors and lack interfaces to restrict the collection and use of their data [51, 104, 259].

Information privacy is about informing people about the collection and use of their data and about empowering them to exercise adequate control over these processes [245]. In the United States, the prevailing legal framework for privacy revolves around the concept of “Notice and Choice.” Notice is typically addressed through the publication of a privacy policy. In practice, users seldom read these privacy policies, which are not just long and difficult to read but also tend to be ambiguous or silent about important issues [149]. Choice is typically offered in the form of opt-in or opt-out decisions such as the recently introduced “Do Not Sell” opt-outs required by CCPA/CPRA, and the opt-in choices mandated by GDPR. Choices are also supplemented with additional data subject rights such as the right to erasure or the right to obtain a copy of one’s data. Notifying data subjects about all relevant data collection practices and expecting them to take advantage of all the choices made available to them thanks to new regulations is impractical [139, 259]. The user burden needed to take advantage of this information and these controls is unrealistically high, often leading people to a state of resignation, where they effectively give up on the idea of trying to control the collection and use of their data. Examples of resignation abound, from people’s attitudes toward cookie consent interfaces [237] to how people feel about managing privacy on social media [148].

1.2 Thesis Contributions

This dissertation explores the diversity of people’s privacy attitudes across contexts that are representative of recent developments in society, including the broad adoption of mobile and IoT technologies as well as the introduction of vaccination requirements that have emerged as a result of the COVID pandemic. Through a collection of user studies, this thesis also sheds light on the challenges associated with empowering people to exercise their right to be informed about and exercise control over the collection and use of their data, given the diversity and complexity of their privacy attitudes. Specifically, we look at new data collection and use scenarios associated with the recent deployment of video analytics technologies across an increasingly broad range of contexts, the privacy challenges arising from the proposed adoption of COVID-19 vaccination mandates and associated vaccination certificates, and the effectiveness of mobile app privacy labels to inform mobile users about the data collection and use practices of mobile apps. Work presented herein is informed by the Contextual Integrity framework, which identifies key contextual parameters influencing people’s privacy expectations and preferences [162]. As emerging technologies continue to proliferate, privacy norms are also changing in response to people’s growing knowledge and awareness of these technologies as well as their experience interacting with these technologies. Findings from this dissertation can provide valuable insights into the evolving privacy norms surrounding these emerging technologies, and help inform the design of public policies and regulations. It is crucial for regulators to recognize that users’ privacy preferences are diverse and context-dependent. Instead of expecting users to repeatedly engage with tedious privacy management tasks, regulators should be open to and in fact promote the adoption of technologies that can empower users to exercise their privacy rights without overwhelming them with repetitive manual tasks. Usable mechanisms should be developed, tested, and made available to assist users in managing their privacy choices in alignment with their unique preferences as they pertain to the context at hand. The main contributions of this thesis include:

- The development of detailed models of people’s privacy expectations and preferences across a broad cross-section of realistic data collection and use practices associated with video analytics deployments, COVID-19 vaccination certificate deployment, and mobile app privacy notice. This includes the identification of key contextual parameters influencing people’s privacy expectations and preferences across these scenarios.
- Beyond the identification and modeling of privacy expectations and preferences that reflect attitudes of broad cross-sections of the population, this dissertation also offers a finer-grained analysis of how some privacy expectations and preferences also vary from one individual to another, and how subgroups of like-minded individuals can often be identified and modeled using clustering techniques.
- We show how user models resulting from our analysis, including the identification of clusters of like-minded individuals, and the introduction of new (APIs) could also serve as a basis for reducing user burden when it comes to managing notice and choice functionality designed to empower users to regain control over their data.
- Our findings shed new light on the unrealistic burden currently placed on users when it comes to managing their privacy across common mobile app and video analytics deployment scenar-

ios. We argue that these findings provide strong support for the introduction of additional regulation that would require the availability of mechanisms, APIs, and protocols designed to reduce user burden. We detail some of such mechanisms in the context of IoT, mobile app, and web browsing scenarios.

1.3 Thesis Outline

The organization of this thesis is as follows: Chapter 2 provides a summary of prior research on usable privacy notices and control mechanisms. It reviews prior work on modeling users' privacy expectations and preferences and the Contextual Integrity framework for studying privacy. Chapter 3 presents an experience sampling study that explores users' privacy expectations and preferences in the context of realistic video analytics deployment scenarios and uses machine learning to model individuals' privacy preferences. Chapter 4 describes a study that focuses on addressing privacy expectations towards the use of COVID-19 vaccination certificates and mandates. It also includes a large-scale analysis of privacy norms. Chapter 5 and 6 present a sequence of two studies focused on the limitations of current mobile app privacy labels, namely succinct and standardized labels intended to inform people about particularly salient data collection and use practices. Chapter 5 describes an interview study that explores lay users' experiences, understanding, and perceptions of iOS app privacy labels in the iOS App Store. Chapter 6 details an analysis of a crowd-sourced corpus of privacy questions collected from mobile app users. The analysis suggests that people's privacy questions are diverse and that an important percentage of these questions are not answered or only partially addressed in today's labels. Finally, the last chapter of this dissertation provides a more detailed discussion of this dissertation's contributions, implications, and future possible work.

Chapter 2

Related Work

2.1 Usable Privacy Notices and Control Mechanisms

The prevailing legal framework for privacy in the U.S. is built upon the concept of “Notice and Choice” derived from the Fair Information Practices Principles (FIPPs) [209]. Privacy notices are declarations of how entities collect, process, retain, and share personal data. We first summarize the privacy literature on four key criteria for usable and effective privacy notices. Then we review prior research on privacy choice.

2.1.1 Usable and Effective Privacy Notice

First, the **readability** of privacy notices is crucial for conveying information. Research has repeatedly shown privacy policies are too long and often require unrealistic education levels to read [68, 153, 215], discouraging people from reading them [68, 149, 153, 215, 242]. Research also indicates that concise privacy notices written in plain language tend to be more effective than lengthy privacy policies [62, 88].

Second, effective privacy notices should promote **comprehension** by the intended audience. Privacy policies often use legal jargon and vague language to allow potential future uses of collected data [194], making it difficult for an average person to comprehend the disclosed data practices [7, 35, 193, 239]. Vu and colleagues’ eye-tracking study found that participants poorly comprehended privacy policies even if they were written at their level of education [243]. Researchers have proposed non-textual privacy notices in addition to privacy policies to convey privacy concepts, such as various indicators [195] and icons [155], but user comprehension of these notices remains a challenge [96].

Third, **salience** determines the likelihood that people will actually find and pay attention to privacy notices. Effective privacy notices should be prominently displayed and easy to access both initially and when users want to revisit them. An eye-tracking experiment found that participants were more likely to read and understand privacy policy information when it was displayed by default rather than accessible only by following a link [228]. Another study found that a prototype Android app privacy label was more likely to be noticed and remembered by users when displayed after they

downloaded an app than when displayed only in the app store [20]. A recent study also indicates that concise privacy notices displayed in a salient way significantly increased user awareness of potentially risky data practices [62].

Forth, **relevance** also impacts the effectiveness of privacy notices. Frequent exposure to lengthy privacy policies containing too much irrelevant information may cause privacy fatigue [36]. Therefore, privacy notices should highlight the most relevant information to their audience, particularly about unexpected, risky data practices [74, 187]. Also, contextually relevant privacy notices tend to be more effective [63, 205]. “Just-in-time” notices like mobile app permissions can provide users contextual information when a specific data practice is about to happen, allowing them to make informed privacy decisions when choices are also provided [73, 205].

2.1.2 Usable Privacy Control Mechanisms

Even though privacy notices are necessary to inform data subjects, usable privacy controls are also imperative to empowering users to exercise the necessary actions they desire to make [7, 49]. Prior research has shown users who were notified about data practices but lack control can resort to privacy resignation [43]. Actually, actionable information about **control** makes privacy notices more useful. This typically means integrating privacy notices with privacy choices (e.g., consent, control options), allowing users to take actions about their privacy based on the disclosures in the notices [49, 75, 205]. One prevalent example of presenting privacy information and obtaining consent is the “ask on first use” approach used by app permissions management systems on Android and iOS [26, 152].

Recent years have also seen a growing amount of user-centered research on usable and effective privacy choice, especially since the introduction of privacy regulations, which mandated a proliferation of privacy settings. For example, since the introduction of GDPR, cookie consent interfaces mandated by GDPR have received considerable attention in the research community [55, 97, 142, 165]. Similarly, Pearman et al. investigated the usability of different health data disclosure authorization designs for a healthcare chatbot in compliance with HIPAA and argued the need for research on alternate approaches to obtain meaningful consent [173]. Moreover, there has been a growing number of user studies that aim to effectively communicate privacy choices to users with the help of tools such as icons, pop-ups, labels, dashboards, and nudges [6, 10, 96, 113, 114, 155, 165, 195, 201].

One barrier to usable privacy controls is dark patterns, namely practices designed to influence users to make privacy choices that ostensibly are not in their best interest [90]. Dark patterns have been found in cookie consent interfaces [91], in emails to unsubscribe from marketing communications [92]. Recently, Habib and Cranor synthesizes the approaches used in prior usability evaluations and introduced a comprehensive framework for systematically conducting evaluations of privacy control mechanisms [94]. They defined usability for privacy choices in terms of seven aspects, and their framework can help provide design recommendations that would improve the usability of these choice mechanisms.

2.1.3 Designing and Implementing Privacy Assistants

The past ten years have seen a proliferation of privacy settings, whether to enable users to block web trackers or to deny mobile apps access to their location. In practice however, users often struggle to configure privacy settings to match their privacy preferences, whether it is because these settings are unintelligible [204], or because the number of available settings is unmanageable [5, 136, 138, 212], or both.

To overcome these usability challenges, recent research has advocated the introduction of “privacy assistants” to (1) notify people about sensitive data collection and use practices and motivate them to manage associated privacy settings [10], and to (2) also help them configure privacy settings [139, 188]. Privacy assistants can be enhanced by incorporating machine learning models of individuals’ privacy preferences to further reduce user burden [136, 138, 139, 217, 249]. For example, Liu et al. successfully demonstrated an Android privacy assistant app that relied on machine learning to generate personalized recommendations about which permission to grant or deny to different apps based on a small number of personalized questions answered by each user [139]. Users could review the recommendations and decide whether or not to accept them. The authors report on a pilot of this technology in the wild, with users indicating they saw value in the way in which this technology made it easier for them to manage a large number of privacy decisions without taking away control over their privacy decisions.

There is a growing body of research focusing on helping people manage their privacy in IoT contexts [51, 65]. This work ranges from the delivery of machine-readable privacy notices to users who are responsible for manually making all privacy decisions [105] to functionality that leverages models of individuals’ privacy preferences to help them manage their privacy. The latter includes the use of machine learning to generate privacy setting recommendations that users can review and accept (or reject) [139] as well as functionality that attempts to automate some privacy decisions on behalf of users [65]. Recent work generally indicates that people appreciate privacy assistant technology that helps them manage privacy decisions, while it also reveals that not everyone feels the same way about how much control they are willing to give up in return for a lighter user burden [43]. The work reported herein is intended to supplement this prior research by providing a more in-depth understanding of individuals’ privacy expectations and preferences in the context of a diverse set of video analytics scenarios. By understanding how rich and diverse people’s expectations and preferences actually are across these scenarios, we aim to build a better understanding of the complexity involved in notifying people about the presence of video analytics deployments and in enabling them to effectively manage associated privacy choices.

2.2 Sampling and Modeling Privacy Preferences

We review previous research on privacy preference modeling. Prior work has shown that individual privacy preferences vary greatly from one person to another and across different data collection and use scenarios [128, 136, 212]. One-size-fits-all models are often unable to capture individuals’ diverse privacy preferences when it comes to the collection and use of their data by mobile and IoT technologies. Research on mobile app permission preferences has shown that it is often

possible to identify common patterns among the privacy preferences of different subgroups of users [108, 135, 136, 138, 147, 191]. Some of this work has also demonstrated the use of machine learning models to predict individuals' privacy preferences [136, 138, 247] and help them manage their privacy decisions [139, 249].

Prior research has also successfully implemented the method using real users with their own devices [139]. For example, Liu et al. demonstrated an Android privacy assistant app that relied on machine learning to generate personalized recommendations about which permission to grant or deny to different apps based on a small number of personalized questions answered by each user [139]. Users could review the recommendations and decide whether or not to accept them. The authors report on a pilot of this technology in the wild, with users indicating they saw value in the way in which this technology made it easier for them to manage a large number of privacy decisions without taking away control over their privacy decisions.

Researchers also have made initial progress in discovering privacy norms with IoT technologies in general by sampling people's privacy expectations and preferences through vignette scenarios using large-scale online surveys [15, 158]. However, vignette studies are limited because participants have to imagine themselves in hypothetical scenarios that are not immediately relevant [4]. The experience sampling method (ESM), where both the context and content of individuals' daily life are collected as research data, better examine links between external context and the contents of the mind [100]. Particularly, mobile-based ESM can prompt participants with the actual context they are in, enabling the collection of higher quality, more valid responses [21, 47]. This motivates us to use ESM to elicit people's privacy expectations and preferences.

2.3 Studying Privacy through Contextual Integrity

The theory of Contextual Integrity (CI) [162, 163] provides a practical way to study privacy and assess the implications of data handling practices. CI defines privacy in terms of the appropriate and legitimate flow of information. Appropriate flow, generally, is a function of conformance with established contextual norms, which are expressible in terms of five CI parameters: three actor parameters (`sender`, `recipient`, `information subject`), an `attribute` parameter, specifying the type of information, and the `transmission principle` parameter, constraining the conditions under which information flows. Being able to specify the values for all 5 parameters is imperative to evaluating the privacy implication of any practice involving information flows. CI posits that a potential privacy violation occurs when one, or more of the information flow parameters, deviates from an established norm. For example, it might be considered appropriate for a doctor to collect the patient's date of birth and prescription drug use for diagnosis purposes. However, if the doctor were to collect this information for advertising purposes or sharing with a pharmaceutical company, the resulting flow—with a different transmission principle and recipient—would deviate from the established expectation.

Nissenbaum's privacy as contextual integrity framework [162] is a theory well suited to evaluate the appropriateness of data practices of new technologies by considering important contextual factors, such as in the case of video analytics deployments and vaccination certificate (VC) deployments. There are increasing privacy concerns about pandemic mitigation technologies re-sharing

people’s personal information, such as controversies related to contact tracing data being shared with law enforcement [61, 151, 197]. Building on the insights from prior studies structured by CI [14, 16, 146, 213], our work on VCs focuses on assessments of appropriateness that explicitly distinguish between initial information flows (i.e., when the data subject is the sender) and the subsequent re-distribution practices (when sender is a different party from subject.) Our study draws on CI to uncover the factors that are likely to affect people’s attitudes and acceptance of re-sharing of information associated with VCs. Accordingly, our study draws on CI to compare reactions both to the initial information flows as well as to the subsequent re-sharing of VC information. The outcome we seek is a comprehensive understanding of people’s attitudes towards the complicated information sharing practices associated with VCs.

However, privacy norms can vary across societies/cultures and may change over time. For example, Gerdon et al. [84] conducted a CI-based longitudinal study in Germany in 2019, before the pandemic, examining people’s acceptance of using individual health data during a pandemic, for public health or for private purposes. In 2020, in the wake of the pandemic, they were able to perform another such (opportunistic) study. Through the lens of CI their findings revealed that the COVID-19 pandemic altered German individuals’ perspective on sharing health data with a public agency, from least acceptable before the pandemic to acceptable in the wake of the COVID-19 pandemic. Open questions remain on whether the perception will swing back after the pandemic subsides. In another CI-based study, Utz et al. [238] examined how these applications handle health information and people’s willingness to adopt them in Germany, the US, and China. They found that participants from Germany and the US perceived sharing “corona app” data with law enforcement agencies as inappropriate. Nevertheless, a restrictive transmission principle (e.g., limited purpose or use) increases the overall appropriateness of information flows. Additionally, compared to Germans and Americans, Chinese respondents considered sharing unique IDs with government servers and digital health certificates overall as more acceptable, highlighting the cultural differences in social norms and privacy expectations.

This thesis uses Contextual Integrity as an organizing framework to explore privacy norms associated with emerging technologies.

2.4 Privacy Challenges of Video Analytics

Video analytics, often equipped with facial recognition, is increasingly being integrated with the Internet of Things (IoT) systems [110, 123, 144]. Data privacy has been a central discussion in IoT [177] because IoT systems rely on the collection and use of contextual information (e.g., people, time, location, activity) in environments that often contains identifiable personal data [37, 176, 178]. Researchers have explored technical solutions to safeguard user data in IoT [58, 60, 203], including algorithms to avoid being tracked by video analytics or facial recognition [210, 211, 250], and systems to enable real-time opt-out of facial recognition systems [50, 51, 202]. However, transparency around IoT data privacy remains an unsolved issue [37, 185, 186]. People often have no way to know the existence of video analytics deployments in their daily environments, what personal data is being collected, what purpose the footage is used for, and how long the footage will be retained. Moreover, video analytics has unique data privacy challenges. First, it can be used to

capture a variety of sensitive information about people, from biometric data (e.g., facial features and body pose) [72, 179, 206, 233] to information about people's activities (e.g., where they are, whom they are with, and what they do) [72, 251] all the way to their emotions (e.g., attentive, depressed, and surprised) [131]. Such information is generally considered more sensitive than people's digital footprints. Second, video analytics can be applied later to video footage already collected by existing cameras for a myriad of purposes (e.g., security, operation optimization, targeted advertising).

These challenges indicate that the privacy implications of video analytics differ greatly in real-world scenarios, and should be evaluated case by case. Nissenbaum's privacy as contextual integrity framework [162] is a theory best suited to evaluate the appropriateness of data practices of new technologies by considering important contextual factors. Under the framework, data practices can be evaluated against certain privacy norms in five information flow parameters — the sender, the recipient, the attribute, the subject, and the acceptable transmission principle. Changes to these parameters are likely to cause a privacy norm violation and must be examined closely [163]. However, privacy norms can vary across societies/cultures and may change over time, so existing privacy norms may not be suitable for new technologies like facial recognition in video analytics. Therefore, the first step to address data privacy challenges of video analytics is to establish a baseline of privacy norms by understanding people's opinions and attitudes towards the technology.

Chapter 3

Understanding Privacy Expectations and Preferences of Video Analytics Technology

3.1 Overview

In recent years, video analytics has been widely deployed in public places, such as airports for security and surveillance purposes, department stores for automatic detection of known shoplifters, and rental car companies for self-checkout [70, 85, 161, 230]. While video analytics can contribute to security, productivity, convenience, and more, its broad deployment also gives rise to serious privacy concerns [216]. These concerns have prompted increased scrutiny from both privacy advocates and regulators [46, 53, 117]. Despite the diverse applications and growing prevalence of video analytics, little is known about how people actually feel about the many different contexts where this technology is being deployed. In this chapter, we report on the findings of an experience sampling study that aims to better understand how people feel about video analytics deployments in different contexts, looking both at the extent to which they expect to encounter them at venues they visit as part of their everyday activities and at how comfortable they are with the presence of such technologies across a range of realistic scenarios.

Our study is organized around two broad sets of questions. The first set focuses on understanding individuals' privacy expectations and preferences. This includes looking for possible social norms that apply to a large fraction of the population [162], or alternatively identifying differences in how people respond to the same deployment scenarios. The second set of questions is motivated by recent technical advances introduced by Das et al.[50], namely (1) the development of real-time face denaturing functionality that enables video analytics software to only be applied to people who provide consent, and (2) the development of a privacy infrastructure for the Internet of Things (IoT) [202]that enables entities deploying video analytics software to publicize their data practices and allow data subjects to opt in or out of data collection, analysis, and sharing practices. Using this functionality, it becomes possible to notify people in real-time as they approach areas where video analytics technologies are deployed and allow them to selectively opt in or out —as might be required in some contexts by regulations such as GDPR or CCPA. Because expecting people to manually opt in or out of video analytics each time they come within range of video analytics

functionality could entail an unrealistically high number of privacy decisions, we use our data to explore the feasibility of developing predictive models that could assist users with their privacy decisions—with users able to review, adopt, adjust, or reject recommendations from the predictive models.

Our in situ study, which spanned 10 days, reveals that participants have rather diverse privacy attitudes towards video analytics deployments and also shows how challenging it could be to empower data subjects to effectively learn about and control data practices associated with these deployments. We find that individuals’ privacy preferences and expectations are complicated and vary with a number of factors such as the purpose for which footage is captured and analyzed as well as the particular venues where it is captured. To alleviate user burden when it comes to managing the many privacy decisions people could be presented with as they come across a variety of video analytics deployments during the course of their daily lives, we explore the feasibility of developing privacy assistants that could possibly be configured to help people more effectively manage these decisions. We discuss how such assistants would require the adoption of standardized APIs and taxonomies of data practices, and how they could also benefit from the use of machine learning techniques.

This work was published at PoPETS 2021 and SOUPS 2021 [255, 259].

3.2 Study Design

3.2.1 Experience Sampling Method

Context has been shown to play an important role in influencing people’s privacy attitudes and decisions [163]. Studying people’s privacy attitudes through online surveys is often limited because participants answer questions about hypothetical scenarios and often lack context to provide meaningful answers. Accordingly, we conducted an experience sampling study to collect people’s responses to a variety of video analytics deployments (or “scenarios”) in the context of their regular everyday activities. The experience sampling method [100] has been repeatedly used in clinical trials [124, 241], psychological experiments [32, 103], and human-computer interaction (HCI) studies [76, 192], yielding “a more accurate representation of the participants’ natural behaviour” [240]. This enables us to engage and survey participants in a timely and ecologically valid manner as they go about their normal daily lives [175]. Participants are prompted to answer questions about plausible video analytics scenarios that could occur at the location in which they are actually situated.

3.2.2 Selecting Realistic Scenarios

Previous research mainly surveyed participants’ privacy attitudes in the context of generic IoT scenarios, including some facial recognition scenarios [129, 158]. By systematically exploring more concrete scenarios in actual settings associated with people’s day-to-day activities, we are able to elicit significantly richer reactions from participants and develop more nuanced models of their awareness, comfort level, and notification preferences pertaining to different deployment

scenarios. The scenarios considered in our in-situ study were informed by an extensive survey of news articles about real-world deployments of video analytics in a variety of different contexts (e.g., surveillance [198], marketing [200], authentication [17], employee performance evaluation [52], and church attendance tracking [18]). These scenarios provided the basis for the identification of a set of relevant contextual attributes which were randomly manipulated and matched against the different types of venues our subjects visited.

Our baseline scenario described the use of generic surveillance cameras with no video analytics. All other scenarios in our study involved the use of some type of video analytics. *Security-related* scenarios included automatic detection of petty crime [198], and identification of known shoplifters and criminals in public places [2, 45, 79, 107]. Scenarios for *commercial* purposes included helping businesses to optimize operations [156, 172, 200], displaying personalized advertisements based on the detection of demographic features [67, 79, 181, 219], collecting patrons' facial reaction to merchandise [25, 30, 40, 208], and detecting users' engagement at entertainment facilities [130, 141, 246]. Other significant use case scenarios revolve around *identification* and *authentication*. Here, we considered two broad categories of scenarios: (1) replacing ID cards with facial authentication in schools, gyms, libraries and places with loyalty programs [17, 64, 159, 214], and (2) attendance tracking in the workplace, at churches, and at gyms [17, 18, 81]. Lastly, we included a small number of plausible, yet hypothetical, scenarios inspired by emerging practices as discussed in news articles or as contemplated in research. This includes health insurance providers using facial recognition and emotion analysis to make health-related predictions [8, 133, 184]; employers using emotion analysis to evaluate employee performance [52, 125, 132]; and hospitals using emotion recognition to make health-related predictions [1, 69, 93].

In total, we identified 16 purposes, as shown in Table 3.1, representative of a diverse set of video analytics scenarios. A representative list of the scenarios as well as the corresponding text shown to participants to elicit their reactions can be found in the Appendix (Table A.1). The scenario text was crafted through multiple iterations to sound plausible without deceiving participants.

3.2.3 Factorial Design

We employed a factorial study design and developed a taxonomy that captured a representative set of attributes one might expect to influence individuals' privacy attitudes. These attributes are shown in Table 3.1. We specified a discrete set of possible values for each attribute, taking into account our desire to cover a broad spectrum of scenarios while also ensuring that we would be able to collect a sufficiently large number of data points for each scenario. Here, we differentiate between the retention time of raw footage and of video analytics results because raw video data, containing biometrics, can be very sensitive, and possibly be exploited for additional analyses subsequently.

3.2.4 Study Protocol and Procedures

The 10-day study comprised the following five stages.

Stage 1: Eligible participants completed the consent forms for this study and downloaded the study app from the Google Play Store. Upon installing the app, participants completed a pre-study

Attribute Name	Values
Purpose	Generic Surveillance Petty crime detection Known criminal detection (Anonymous) people counting (Individualized) jump the line offers (Anonymized) demographic ad targeting (Individualized) ad targeting (Anonymized) sentiment-based ad targeting (Individualized) sentiment-based ad targeting (Anonymous) sentiment-based customer service evaluation (Individualized) customer engagement detection Attendance tracking Using face as IDs Work productivity predictions Health predictions - eatery visits Health predictions - medical visits
Anonymity level	No video analytics Anonymous face detection Facial recognition
Retention of raw footage	ephemeral, 30 days, unspecified
Retention of analysis results	ephemeral, 30 days, unspecified
Sharing specified	Yes, No
Detection of whom people are with	Yes, No
Type of places	store, eatery, workplace, education, hospital, service, alcohol, entertainment, fitness, gas, large public places, transportation, worship, library, mall, airport, finance

Table 3.1: Contextual attributes: Among all the possible combinations of these attributes, our study focused on a subset of 65 scenarios representative of common and emerging deployments of video analytics technology.

survey about their perceived knowledge level, comfort level, and notification preference with regard to facial recognition.

Stage 2: Participants were instructed to go about their regular daily activities. The study app collected participants' GPS locations via their smartphones. As they visited points of interest, namely places for which we had one or more plausible deployment scenarios, the app would send them a push notification, prompting them to complete a short survey on a facial recognition scenario pertaining to their location, as illustrated in the app screenshots in Figure 3.1a–Figure 3.1d. The protocol limited the number of scenarios presented to each participant to six per day, though most of the time participants' whereabouts would trigger a smaller number of scenarios—closer to three

per day.

Stage 3: On the days participants received push notifications via the app, they also received an email in the evening to answer a daily summary web survey (“evening review”). This web survey showed participants the places they visited when they received notifications, probed reasons for their in-situ answers, and asked a few additional questions. See Figure 3.1e for an example of the evening review.

Stage 4: After completing 10 days of evening reviews, participants concluded the study by filling out a post-study survey administrated via Qualtrics. This survey contained free-response questions about their attitudes on facial recognition, the 10-item IUIPC scale on privacy concerns [143], as well as additional demographic questions like income, education level, and marital status.

Stage 5 (Optional): Participants who indicated they were willing to be interviewed in their post-study survey may be invited to an online semi-structured interview. The interview contained questions about study validity, perceptions of scenarios, and clarifications with regard to their earlier responses. The full text of the post-survey and the interview scripts can be found in the Appendix (Appendix A.3 and Appendix A.4).

To maximize the contextual benefits provided by the experience sampling method [39], we designed a sophisticated payment scheme to incentivize prompt responses to in-situ notifications. Participants were compensated \$2 per day for each day of the study. They received an additional 25 cents per notification they responded to within 15 minutes, or 10 cents if they responded to the notification between 15 and 60 minutes. We also compensated them \$2 for the time spent on answering pre-study and post-study surveys. An additional \$15 was awarded when they finished the study. In total, participants could earn between \$37 and \$52 and were compensated with Amazon gift cards. Participants who completed the online interviews were awarded \$10.

3.2.5 Ensuring Study Validity

Due to the complexity and the number of components of the study framework, we conducted several pilot rounds, with initial rounds involving members of our research team and later rounds involving a small number (N=9) of external participants. Each pilot round helped identify issues that needed to be addressed, whether in the form of small refinements of our protocol or adjustments to technical components of our system (e.g., study app, web survey app, study server). Below, we briefly discuss the two most important refinements that were made as a result of this process.

Because of the limitations of location tracking functionality, we determined that we could not automatically pinpoint the location of our subjects and use that location to automatically identify a relevant video analytics scenario. Instead, we opted to use location tracking to automatically generate a drop-down list of venues near our subject. We then asked them to select the actual venue where they were. The drop-down list of venues always included three additional options: “I was somewhere else in the area,” “I was passing by,” and “I was not there.” This ensured that our protocols also accounted for missing venues, situations where our subjects were passing by a given location (e.g., being stuck in traffic), as well as situations where location tracking was potentially inaccurate. Participants still received payments for each scenario when they selected one of these three additional choices. In other words, they had no incentive to select a place that they did not visit.

During the first pilot, we found that some participants did not seem to pay close attention to some of the scenario attributes (Table 3.1). This was remedied by introducing two multiple-choice attention check questions (see Figure 3.1b). These questions required participants to correctly identify two different and randomly selected contextual attributes assumed in the scenario (attributes in Table 3.1, excluding type of places). Participants were only allowed to proceed with the remaining in-situ questions once they had passed the two attention checks. These attention checks proved rather effective, as discussed in Section 3.4.1.

3.2.6 Recruitment and Ethics

We recruited participants using four methods: posts on local online forums for the Pittsburgh area (e.g., Craigslist, Reddit), posts in a university-based research participant pool, promotional ads on Facebook, and physical flyers posted on local community bulletin boards and at bus stops. Potential participants were asked to take a short screening survey to determine eligibility (age 18 or older, able to speak English, using an Android smartphone with data plan). The screening survey also displayed the consent form for the study and collected basic demographic information such as age, gender, and occupation. Recruitment materials, the consent form, and the screening survey did not mention or refer to privacy. We tried to avoid convenience samples of undergraduate college students, and purposely looked for participants with a variety of occupations.

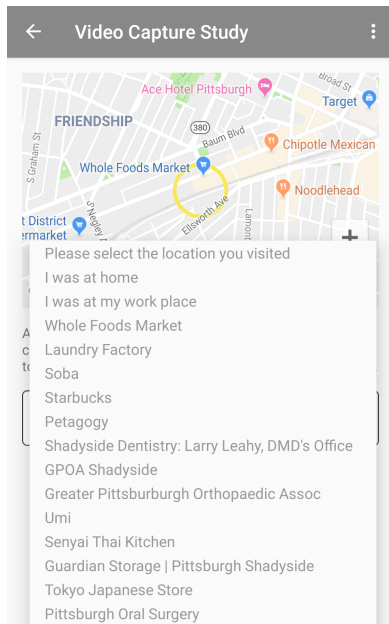
This research was approved by our university’s institutional review board (IRB) as well as the funding agency’s human research protection office. As location data collected over a period of time can be particularly sensitive, we refrained from using off-the-shelf experience sampling software and developed our own system and location-aware Android app.

3.3 Participants and Responses

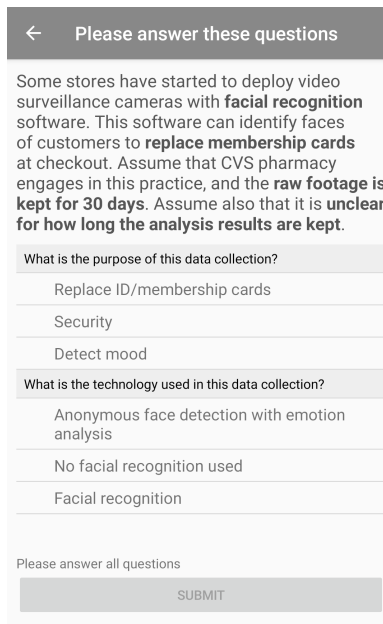
Gender	%	Age	%	Education	%	Income	%	Marital Status	%
Female	57.7	18-24 years old	8.1	Some high school	.8	Less than \$25,000	14.6	Single, never married	50.4
Male	40.7	25-34 years old	54.5	High School	4.1	\$25,000 to \$34,999	14.6	Married	41.5
Other	1.6	35-44 years old	23.6	Some college	13.8	\$35,000 to \$49,999	9.8	Separated	1.6
		45-54 years old	8.1	Associate’s degree	7.3	\$50,000 to \$74,999	22.0	Divorced	3.3
		55-64 years old	3.3	Bachelor’s Degree	35.0	\$75,000 to \$99,999	14.6	Widowed	0.8
		65-74 years old	2.4	Master’s Degree	23.6	\$100,000 to \$149,999	14.6	I prefer not to answer	2.4
				More than Master’s Degree	12.8	\$150,000 to \$249,999	2.4		
				Other	1.6	I prefer not to answer	7.3		

Table 3.2: Survey participant demographics and respective %

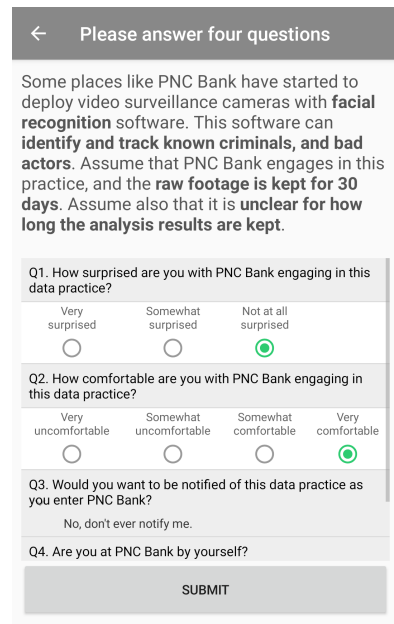
A total of 164 individuals (excluding 9 pilot participants) took part in the study and downloaded our study app from the Google Play Store between May and November 2019. Of these, 124 completed the 10-day study. One participant was removed due to poor response quality as that person selected “I was somewhere else” for all the notifications received. Among the remaining 123 participants, 10 (8%) were 18-24 years old, 67 (54.5%) were 25-34, 29 (23.6%) were 35-44, 10 (8%) were 45-54, 4 (3%) were 55-64, and 3 (2%) were between 65 and 74. In our sample,



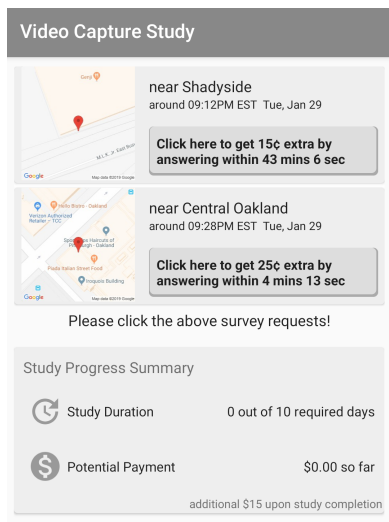
(a) Prompting users to clarify their location



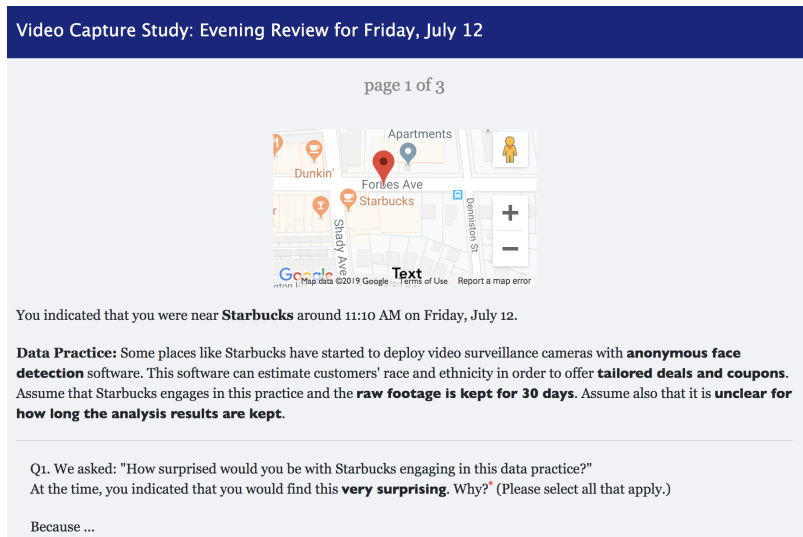
(b) Two attention check questions designed to ensure participants read about relevant attributes



(c) Four in-situ questions



(d) Dashboard showing prompts to complete two in-situ surveys, including monetary incentives to respond as quickly as possible



(e) Partial screenshot of evening survey associated with a given scenario encountered earlier during the day

Figure 3.1: Screenshots of the study app and the web survey used for the evening review

58% identified as female, 41% as male, and 2% as other. Most participants were highly educated: 43 (35%) had bachelor's degrees, and 46 (37%) had graduate degrees. Half of the participants

Occupation	%	Occupation	%
Business, or sales	12.2	Legal	3.3
Administrative support	9.8	Other	3.3
Scientist	8.9	Graduate student	2.4
Service	8.1	Skilled labor	2.4
Education	8.1	Homemaker	2.4
Computer engineer or IT	7.3	Retired	2.4
Other salaried contractor	7.3	Government	1.6
Engineer in other fields	6.5	Prefer not to say	1.6
Medical	6.5	Art or writing	.8
Unemployed	4.1	College student	.8

Table 3.3: Occupations of survey participants and respective %

were single and never married, and 42% were married or in a domestic partnership. The majority of our participants (82%) reported having no children under 18 living with them. Participants reported diverse occupations (see Table 3.3). The average IUIPC factor scores of our participants are shown in Table 3.4. Comparing our results with those of a large MTurk sample from another study (N=1007) [158] using Mann-Whitney U tests, we found no difference in the collection and the awareness factors, and a significant difference in the control factor with a small effect size ($r = 0.1, p < 0.01$).

	Ours Mean [SD]	MTurk Mean [SD]	Reject H0
IUIPC-Collection	5.90 [1.04]	5.79 [1.11]	No
IUIPC-Control	6.21 [0.78]	5.95 [0.90]	Yes
IUIPC-Awareness	6.53 [0.66]	6.44 [0.82]	No

Table 3.4: Comparison of IUIPC scores of our participants (N=123) with an MTurk sample (N=1007). H0 stipulates that two samples come from the same population. Cannot reject H0 means that 2 groups are not significantly different.

We recruited interviewees about halfway through the study. Participants were selected based on their demographics. We sent out 17 invitations and conducted online interviews with 10 participants who followed up.

In total, participants were sent 3,589 notifications prompting them to identify their specific location (Figure 3.1a). In the majority of cases (65%), our system was able to retrieve a scenario relevant to the location reported by the participant, such as the two scenarios shown in Figure 3.1b and Figure 3.1c. For the remaining 35%, the system did not have a pre-identified scenario that matched the response provided by the participant, in which case we were unable to elicit any

additional information from the participant for that particular location. Based on answers provided by participants, common examples of such situations included the participant being at home or visiting a partner, friend, or relative. Other situations included the participant waiting for a bus or passing by a location. In some instances, participants reported that they did not see the location at which they were in the drop-down menu shown to them (Figure 3.1a). This seemed to most commonly occur when participants were in parks, parking lots, farmers' markets, new establishments, or small local stores.

When the system was able to retrieve a plausible scenario relevant to the participant's location, the participant was presented with the scenario and prompted to answer a few quick questions related to that scenario (e.g., see Figure 3.1b and Figure 3.1c). In addition to these in-situ responses, they were also requested to answer a more complete set of questions about the scenario in the evening. As a result, we were able to collect in-situ and evening responses for a total of 2,328 scenarios. Each participant on average provided in-situ and evening responses to 19 scenarios over a 10-day period, and received an average compensation of \$41.

3.3.1 Qualitative Data Set and Analysis

We also analyzed the qualitative data set collected from the 10-day experience sampling study. The data set includes 2,562 entries of text responses from participants' daily summaries, 1,230 entries of text responses in the post-survey, and 10 interview transcripts. The interviews ranged from 26 to 40 minutes (mean=33) and were fully transcribed. A total of 326 minutes of transcripts were analyzed. I read and familiarized myself with all the transcripts and then applied thematic analysis [28] to open code the transcripts. A second research met with me regularly to iterate on the themes.

In order to answer the research questions, it is crucial that the qualitative data collected reflects participants' attitudes towards facial recognition. Since we adopted an experience sampling method presenting realistic scenarios of facial recognition to participants over 10 days, we believe the data collected following these contextual cues would capture participants' perceptions and attitudes.

From the 10-day study, we collected 2,562 entries of text responses from participants' daily summaries and 1,230 entries from the post-survey. In the post-survey, there were 10 open-ended questions. The first question was "What is the first thing that comes to your mind when you think about facial recognition technology?" We coded the sentiment (i.e., positive, negative, neutral, mixed) in each response.

We included two questions in the post-survey asking participants' perceived beneficial and concerning contexts to use facial recognition technology. We also asked questions eliciting participants' privacy concerns about facial recognition deployment scenarios. After reading the survey responses, we realized many participants shared their attitudes and experiences with facial recognition deployment scenarios regardless of to which question they were responding. Since the daily summaries were also addressing similar issues, in our analysis, we broke down the boundaries between the data sources and conducted a content analysis [229] of all the participants' 3792 textual responses.

Two authors started from inductive coding [28] to extract codes that show participants' perceived benefits or concerns about facial recognition technology and developed a codebook. In total, we summarized 13 main codes with 32 subcodes about the benefits of facial recognition and 19 main codes with 40 subcodes about the concerns. In the end, we used a deductive coding approach,

applying the codebook to the entire dataset. Two authors independently coded all data and met to resolve any discrepancies.

3.4 Privacy Preferences

When surveying participants’ responses to facial recognition scenarios, we focused on four related questions: how surprised they were by the scenario presented to them (**surprise level**), how comfortable they were with the collection and use of their data as assumed in that scenario (**comfort level**), to what extent they would want to be notified about the deployment scenario at the location they visited (**notification preference**), and whether, if given a choice they would have **allowed** or **denied** the data practices described in that scenario at that particular location at the time they visited that location (**allow/deny preference**). These questions are shown in Figure 3.2.

How surprised are you with *Controller* engaging in this data practice?

Very surprised Somewhat surprised Not at all surprised

How comfortable are you with *Controller* engaging in this data practice?

Very uncomfortable Somewhat uncomfortable Somewhat comfortable Very comfortable

Would you want to be notified of this data practice as you enter *Controller*?

Yes, notify me every time it happens.
Yes, but only once in a while to refresh my memory.
Yes, but only the first time I enter this location.
I don't care whether I am notified or not.
No, don't ever notify me.

If you had the choice, would you allow or deny this data practice?

Allow Deny

Figure 3.2: *Controller* being a variable that would be instantiated with the name of the venue participants were visiting

Figure 3.3 provides a summary of collected responses organized around the 16 categories of scenarios (or “purposes”) introduced in Table 3.1. As can be seen, people’s responses vary for each scenario. In other words, “one size fits all” would fail to capture individuals’ diverse preferences when presented with these scenarios. At the same time, some scenarios elicit more consistent responses from participants than others. For instance, generic surveillance scenarios appear to surprise participants the least and to elicit acceptance by the most (close to 70% would agree to such scenarios, if given a choice and fewer than 10% reported feeling “very uncomfortable” with such scenarios). Yet, even in the presence of such scenarios, 60% of participants reported they would want to be notified at least the first time they encounter these scenarios at a given

venue and over 35% indicated they would want to be notified each time. At the other end of the spectrum, scenarios involving facial recognition for the purpose of evaluating employee productivity or tracking attendance at venues elicited the greatest level of surprise and lowest level of comfort among our participants, with barely 20% reporting that, if given a chance, they would consent to the use of these technologies for the purpose of evaluating employee productivity. Similarly, participants expressed significant levels of surprise and discomfort with scenarios involving the use of facial recognition to make health and medical predictions or to track the attendance of individuals.

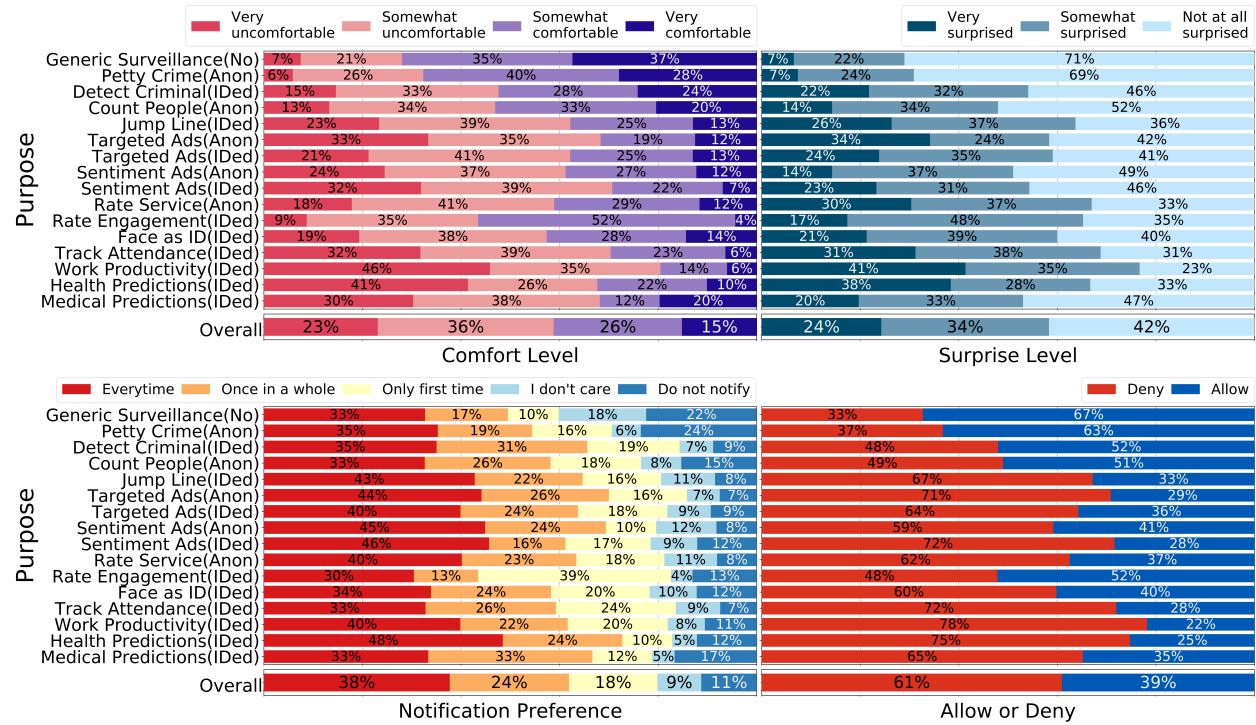


Figure 3.3: Summary of collected responses organized around 16 different purposes. The bottom row shows the aggregated preferences across different purposes.

3.4.1 Study Validity and Benefits of ESM

Below we report results on study validity, focusing on three aspects: whether participants carefully read the scenarios, whether they thought the scenarios could happen, and how the ESM helped anchor their responses to their everyday life experience.

Overall, 81% of the time participants successfully completed both attention check questions associated with the scenarios assigned to them within two attempts. Attention questions were found to be useful by 8 out of the 10 interviewees. For instance, one participant (P107) stated, “I think you definitely had to read them [scenarios]. I think there was one or two that I saw the bold words, and thought that they were the same as older questions, so I picked the same answer, and

it was a different one. So once I re-read it, I saw that it was a little different.” Five interviewees reported attention questions helping them discern between retention for raw footage, and retention for analysis results, as P55 said, *“But the first couple of times, I mixed up the raw footage with the analysis results, but after that [the attention checks] I remembered to look for the distinction.”* These comments suggest that the attention checks contributed to participants noticing the contextual attributes associated with each scenario and that the responses we collected most likely reflect privacy attitudes that take these contextual attributes into account.

As 68% of in-situ questions were answered within 15 minutes and 87% within 1 hour, the actual location visited by the participant and the context associated with the scenario were likely still fresh in their mind (e.g., what the participant was doing at that particular location, or whom they might have been with). When asked about whether the scenarios matched actual video collection practices at the places participants were visiting in the exit interviews, most ($N = 7$) stated that they found the scenarios to be realistic, and *“it is entirely possible that it is happening in those places”*(P55). P107 explained, *“I don’t know if they actually use any of the strategies right now, but they did seem to fit pretty well with the places like grocery stores offering coupons, or targeting some ads towards you.”*

Furthermore, the experience sampling method provided context to participants’ responses, with participants reporting that context played an important role in influencing their attitudes towards different video analytics deployments. When the participants selected in situ that they felt somewhat or very uncomfortable about a scenario, in daily the evening reviews they can select multiple-choice options and provide additional free responses to further explain their discomfort. Figure 3.4 plots the reasons participants selected, many of which are directly related to the in-situ context. The figure also shows the percentages of participants who ever reported considering each reason: many reasons were taken into account by the majority of 123 study participants. Our qualitative analyses of free responses in evening reviews also revealed that study participants had context in mind when they explained their in-situ comfort level. Their responses also reflected various aspects of data flows as by Nissenbaum’s framework of CI [162]. Example quotes listed by purpose are shown in Table 3.5.

3.4.2 Factors Impacting Privacy Attitudes

The responses collected as part of this in-situ study provide rich insight into people’s awareness of the many different ways in which facial recognition is deployed, how comfortable they are with these deployments, and to what extent they would want to be notified about them. Our analysis is organized around the different contextual factors already identified in Table 3.1.

On average each participant responded to a total of about 19 deployment scenarios. These 19 different scenarios covered an average of 9.9 different “purposes,” as defined in Table 3.1, and 5.9 different types of venues, thereby offering rich insight into how people feel about facial recognition deployments across a range of different situations.

Allow/Deny Decisions

We first investigate whether people’s decisions to allow or deny data collection have a relationship with the contextual attributes in Table 3.1. We constructed our model using generalized linear

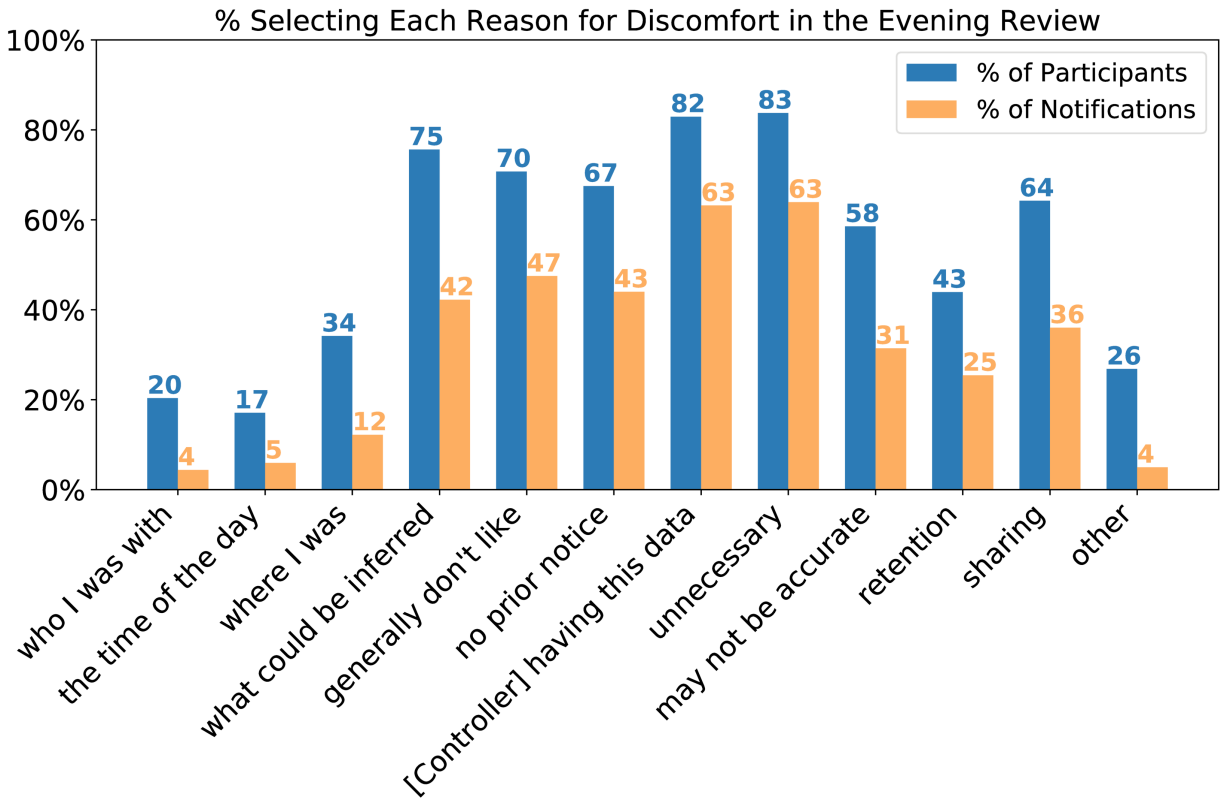


Figure 3.4: Percent of participants/notifications reporting specific reasons for discomfort. Participants only selected reasons for notifications that they indicated discomfort (N=1,369). N is the used as the denominator to calculate the percent of notifications.

Purpose	Example Quotes	Values
Generic Surveillance (No)	I'm fine with it to keep banks more safe. — P27	A,TP
Petty Crime (Anon)	When it comes to law enforcement and public safety I am more ok with giving up privacy. But there is an expectation that the data is protected. But any data collected for one reason is expected to stay within that original use. — P12	R,TP
Detect Criminal (IDed)	Because this is a bar, I feel like I would be more willing to acquiesce to a certain degree of surveillance for my own safety. — P59	A,TP
Count People (Anon)	It's anonymous and seems like a good use of the technology. — P68	TP
Jump Line (IDed)	The cafe is never super crowded when I go, and the space is small. I am surprised they would need something like that due to area and logistics. — P16	R
Targeted Ads (Anon)	I've heard that Target has the most advanced security, so it's kind of unsettling because I don't know exactly what they're doing. — P7	R,TP
Targeted Ads (IDed)	It's the facial recognition of it and keeping of derived data that bothers me. — P13	A,TP
Sentiment Ads (Anon)	It's anonymous so I don't care as much. Also I have pretty good brand loyalty to Target and trust them more than I probably should. — P40	TP,R
Sentiment Ads (IDed)	The errands I do there are acceptable for all audiences. — P9	A
Rate Service (Anon)	I would expect this practice from larger chains rather than a small, local store, so it weirded me out a little to think the surveillance technology was there. — P27	R
Rate Engagement (IDed)	It might help improve the experience. — P110	TP
Face as ID (IDed)	I trust this location with footage as it is my local gym, and it actually would be convenient in this case. — P106	R,TP
Track Attendance (IDed)	It's a military base with 100% ID check at the gate, so I know about it and basically trust them. — P25	R
Work Productivity (IDed)	Big Brother is watching. I did not consent. — P104	TP
Health Predictions (IDed)	I don't like sharing data with health insurance companies. — P13	TP
Medical Predictions (IDed)	Emotion analysis combined with facial recognition makes me more uneasy than other ways this tech is implemented, especially coming from a healthcare provider. — P58	TP,R

Table 3.5: Example quotes from participants' evening reviews explaining their in-situ answers. Their responses were coded by relevant parameter values of contextual integrity. A—Attribute: Any description of information type. R—Recipient: Any entity (person, company, etc.) that receives the information. TP—Transmission Principle: The conditions under which information may be used or collected [163].

mixed model (GLMM) regression [22], which is particularly useful for data analysis with repeated measures from each participant. Our GLMM model was fit by maximum likelihood (Laplace approximation) treating the user identifier as a random effect, using a **logistic link** function for the binary response (allow/deny).

Among all the attributes introduced in Table 3.1, we find that “purpose” exhibits the strongest correlation with the decision to allow or deny data practices associated with our scenarios. In particular, when compared against “generic surveillance” scenarios, 12 out of 15 other purposes came out as being significantly more likely to result in a “deny” decision. Participants were respectively 23.5 ($=e^{3.16}$) times and 29 ($=e^{3.37}$) times more likely to respond with a “deny” to deployment scenarios for predicting work productivity, and for predicting health, compared to generic surveillance scenarios with no facial recognition. The odds of participants denying purposes for targeted advertising were at least 6 ($=e^{1.87}$) times and up to 16 ($=e^{3.16}$) times greater than the odds for generic surveillance. Even for the purpose of using faces for authentication and identification, participants were still more likely to deny data collection (odds ratio $= e^{1.70} = 5.5$). Three purposes turned out not to be significant: detecting petty crime, using anonymous facial detection to count the number of people in the facility, and using facial emotion detection to rate engagement. The last of the three purposes, despite being relatively intrusive in comparison with the previous two, did not seem to have an important impact. We suspect that this might be partially due to the low number of occurrences ($N = 23$) of this purpose as this scenario was only associated with visits to places like movie theaters, museums, and amusement parks.

Contrary to our expectations, we found that whether targeted ads relied on identifying individuals or treating them anonymously did not elicit substantially different responses from our participants. In fact, participants reported being more likely to respond with a “deny” to facial recognition scenarios used in targeted ads based on demographic features like race or ethnicity than to scenarios which involved individually targeted ads. The interview data revealed that some participants (3 out of 10) were viewing advertising based on demographics (e.g., race and age) as a form of profiling. For example, P106 stated, *“I do think it will divide us more if they are targeting specifically based on what you look like, not even necessarily your profile and who you are ... I think it just gives an overall weird and gross feeling, especially in today’s society where it comes up a lot.”*

Some of the place type attributes were also found to have an influence on participants’ allow or deny decisions. When we compare different place types to the baseline of large public places (e.g., sports stadiums, parking garages, city hall buildings), we find that participants were more likely to deny data practices at eateries (odds ratio $= e^{1.09} = 3$), at libraries (odds ratio $= e^{1.71} = 5.5$), and at gas stations (odds ratio $= e^{1.36} = 3.9$). Participants were significantly less likely to respond with a “deny” to deployment scenarios at transportation locations (buses stops, train stations, metro stations) than at the more generic baseline (odds ratio $= e^{-1.87} = 0.23$). The number of days participants had been in the study also seemed to influence their allow/deny decisions. Participants proved more likely to respond with a “deny” as the study progressed. None of the other attributes were statistically significant ($p < 0.05$). We present the complete results from the regression in the Table 3.6.

Factors	Est.	Std. Err	Z	p
Intercept	-1.79965	0.60789	-2.96	0.003072**
purpose:baseline = Generic Surveillance				
Petty Crime(Anon)	0.57922	0.52134	1.111	0.266563
Criminal Detection(IDed)	1.08567	0.43613	2.489	0.012799*
Count People(Anon)	0.54011	0.56511	0.956	0.339187
Jump Line(IDed)	2.12133	0.53749	3.947	7.92E-05***
Targeted Ads(Anon)	2.77327	0.56614	4.899	9.66E-07***
Targeted Ads(IDed)	1.87295	0.5265	3.557	0.000375***
Sentiment Ads(Anon)	2.03323	0.70039	2.903	0.003696**
Sentiment Ads(IDed)	2.7837	0.59923	4.645	3.39E-06***
Rate Service(Anon)	1.92574	0.55494	3.47	0.00052***
Rate Engagement(IDed)	0.9621	0.92536	1.04	0.298478
Face as ID(IDed)	1.70491	0.51797	3.292	0.000997***
Track Attendance(IDed)	2.56281	0.60284	4.251	2.13E-05***
Work Productivity(IDed)	3.15627	0.63879	4.941	7.77E-07***
Health Predictions(IDed)	3.37146	0.58706	5.743	9.30E-09***
Medical Predictions(IDed)	1.92103	0.7824	2.455	0.014077*
Raw retention:baseline=30 days				
Ephemeral	0.10859	0.3799	0.286	0.775005
Unspecified	0.23487	0.4079	0.576	0.564742
Analytics retention:baseline=unspecified				
Ephemeral	-0.02068	0.81819	-0.025	0.979836
30 days	-0.22812	0.30495	-0.748	0.454423
Association: baseline=No				
associationID	0.27251	0.18042	1.51	0.130937
Shared: baseline=No				
sharedID	-0.09074	0.26258	-0.346	0.729666
dayIndex	0.79628	0.27167	2.931	0.003378**
placeType:baseline=large public places				
store	0.73456	0.42748	1.718	0.085732
eatery	1.09194	0.41956	2.603	0.009252**
work	0.46835	0.50123	0.934	0.350094
education	-0.48813	0.50161	-0.973	0.330493
hospital	1.11144	0.65184	1.705	0.088178
service	0.67614	0.52179	1.296	0.195037
alcohol	0.81001	0.4635	1.748	0.08053
entertainment	0.80385	0.61804	1.301	0.193377
fitness	1.06873	0.66162	1.615	0.10624
gas	1.36253	0.58379	2.334	0.019598*
transportation	-1.48697	0.5998	-2.479	0.013171*
worship	-0.27275	0.81689	-0.334	0.738463
library	1.71228	0.71968	2.379	0.01735*
mall	1.19774	0.89793	1.334	0.182241
airport	0.08364	0.96362	0.087	0.930832
finance	-1.13355	1.16506	-0.973	0.33058

Table 3.6: Generalized Linear Mixed Model Regression with Logit Link. A positive coefficient(estimate) shows likeliness of participants' to deny a data collection

Comfort Level, Surprise Level, and Notification Preference

Here we explore how the different contextual attributes considered in our study seem to influence participants’ comfort level, surprise level, and notification preferences. As those responses are not binary or linear, GLMM is not suitable due to its inability to model ordinal dependent variables. Instead, we opted for cumulative link mixed models (CLMM) fitted with the adaptive Gauss-Hermite quadrature approximation with 10 quadrature points using the R package `ordinal` [38]. We constructed one CLMM model for each dependent variable, adopting the same set of independent variables and random effect, as is the case with allow/deny decisions described in Section 3.4.2.

Similarly to the case with allow/deny decisions, purpose remains the attribute with the strongest influence on participants’ comfort level, surprise level, and notification preferences. Participants are more likely to feel uncomfortable, surprised, and are more likely to want to be notified when confronted with scenarios involving facial recognition than with our baseline “generic surveillance” scenario with no facial recognition. Data sharing with other entities seems to also contribute to a significant reduction in comfort among participants. As is the case with allow/deny decisions, we also found that the number of days in the study was significantly correlated with participants’ surprise level and notification preferences. Participants reported being less surprised over time, likely because they had already encountered similar scenarios earlier in the study. Over time, participants became slightly more inclined to deny scenarios, while their notification preferences became somewhat more selective. These results are further explored in Section 3.4.3 and Section 3.4.3.

3.4.3 Attitude Change Between Start and End of the Study

In our pre-study and post-study surveys, we asked participants the same questions about their understanding of, comfort level with, and notification preferences for facial recognition. In the post-study, we also asked them to provide open-ended responses to why their level of concern may have (not) changed. We analyzed these responses using inductive coding. Two researchers iteratively improved the codebook and independently coded all responses. Coding discrepancies were discussed and reconciled. We reported results from comparing both surveys and qualitative coding.




Figure 3.5: A Sankey diagram shows the change of participants’ reported notification preferences before and after the study

Increased Awareness

By the end of the study, 60% of participants ($N = 74$) reported increased awareness resulting from participation in the study. They did not realize facial recognition could be used for so many different purposes, at such a diverse set of venues, and with this level of sophistication. For instance, P68 wrote, *“Some of the scenarios and growth of the technology you mentioned, I had never considered. Freaked me out.”* 11% of the above group reported learning the benefits of facial recognition. P106 explained, *“In the beginning I was very uncomfortable with the fact that this tech could be abused or that law enforcement could use it. However, as the scenarios came up in the study, I realized it could be helpful in my life as long as there are safeguards in place to prevent abuse.”* At the end of the study, when rating how much they thought they knew about facial recognition, one third of participants rated their knowledge of facial recognition lower than what they had reported at the start. This situation could be explained by the Dunning-Kruger effect, a cognitive bias wherein people tend to overestimate their knowledge in areas which they have little or no experience [126]. As participants grew more aware of possible video analytics deployments, they gained a more grounded estimate of their knowledge level. In interviews, 5 out of 10 interviewees indicated their awareness had increased. For instance, P50 mentioned *“I didn’t know when I started there were so many different potential uses. I only thought that it could be used for tracking someone who committed a crime, so I was really surprised that there are so many different things being developed. And I definitely do think there are good uses and some that are more invasive.”* Three interviewees described their deliberation on facial recognition usages as the study progressed. For example, P56 recounted *“I feel like I might’ve started to get more negative about the use of cameras... I could easily how see all of this information would go to very bad places... In some ways now that I am more aware of it, I’ve certainly put more thought into it and became more negative about it.”*, and P107 gave an account of his thought process: *“I think it’s just thinking about it more, being asked a couple of different times, and then you get asked once you just kind of answer it, but then twice and the third, I really think about it. It’s been in my mind already, so then the answer is probably more close to what I think... by the end, maybe I am not so sure about them having that information. But I think by the last 3 or 4 days, they were more consistent, consistently no for certain ones.”* This could possibly explain why the number of days in the study was a significant predictor of participants’ allow and deny preferences and why they tended to deny more as the study progressed as reported at the end of Section 3.4.2.

Evolution of Notification Preferences

Before the study, 95.9% of all participants claimed that they wanted to be notified about facial recognition deployment scenarios, including 51.2% who indicated they wanted to be notified every time they came within range of facial recognition. As shown in Figure 3.5, between the beginning and end of the study 55.3% of participants changed their preferences regarding whether and how often they wanted to be notified about facial recognition deployments. Among participants who originally wanted to be notified every time, 44% of them opted for less frequent notifications. This is also supported by the positive coefficient associated with the number of days predictor of the CLMM regression model for notification preferences, as stated in Section 3.4.2, as well as the

descending line in Figure 3.6, which plots the percentage of notifications where participants want to be notified every time or once in a while against the number of days in the study.

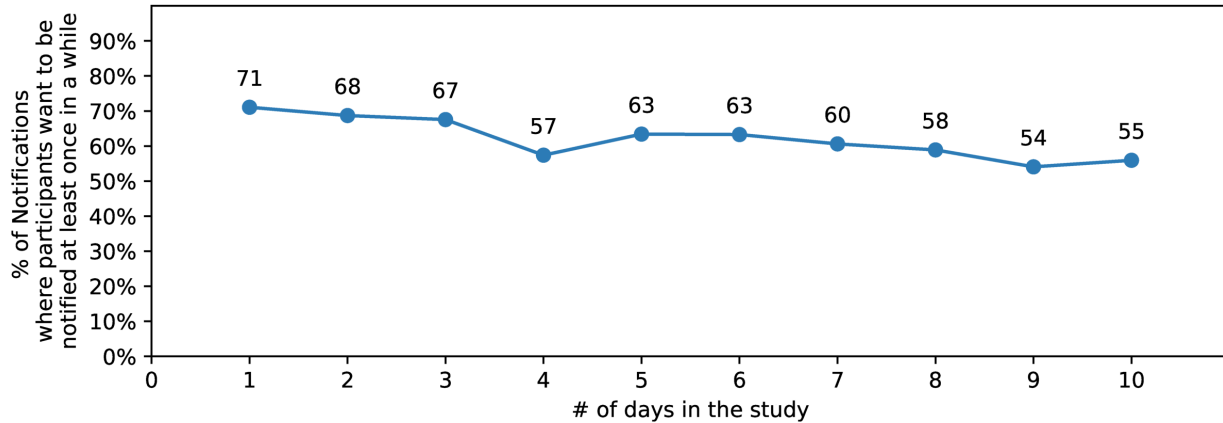


Figure 3.6: Participants’ desire to be notified decreases as the study progresses

One possible explanation is that people gradually developed a better appreciation for the broad deployment of these scenarios, and the possibility of receiving a large number of notifications, as P53 described, “*I think at first when I first started, I was saying once in a while and then I realized that would be really annoying to get multiple notifications.*” Some participants also expressed resignation. For instance, P89 said, “*The whole concept has become normal to me. I’ve definitely been reminded, through the app, that cameras with facial recognition are used in many, many places. I’ve become desensitized to the practice, and in fact, what I had considered in some ways[sic] to be negative because I want my privacy.*” It is also worth noting that, as can be seen in Figure 3.5, a simple “Ask on First Use” approach would not accommodate most users. If anything, changes identified in participants’ responses before and after the study indicate that people seem to become more sophisticated over time in their notification preferences with a substantially smaller fraction of participants requesting to be notified every time by the end of the study. The majority are looking for some type of selective notification solution.

On the other hand, we also noticed that a sizable minority of participants (shown in bottom of Figure 3.7) stayed relatively consistent throughout the study with regards to their notification preferences, as they wanted to be notified every time facial recognition is in use. Results from interviews revealed that some participants would always want to be notified, like P56 noted “*The fact that I wanted everything to be always reminding me... I think it is worth letting people know upfront, and every time so you don’t get used to it and complacent.*” P52 also explained why he would always want to be notified at his workplace: “*At my work, if I didn’t think it was necessary or appropriate, then it wouldn’t register in my head that I was being watched. I would have to be reminded every time.*”




Figure 3.7: The graph displays the notification and allow/deny preferences of all participants in chronological order over the course of the study. Each participant is represented by a row. The left graph arranges participants in order of increasing desire to receive notifications, while the corresponding allow/deny preferences are shown on the right. On the left graph, participants' desire to get notified less over time is illustrated by the gradually brightening color from left to right, especially in the top right corner.

3.4.4 Correlation Between Privacy Expectations and Allow/Deny Preferences

Prior research has shown that comfort is often correlated with the degree of surprise people express towards different data collection and use practices [135]. We compiled pairwise correlations between the four types of responses collected from our participants across the 2,328 scenarios evaluated in our study (Table 3.7). Correlations were calculated using the Spearman rank correlation with Bonferroni-corrected p -values. Not too surprisingly, we find a significant correlation with a large effect size between people’s comfort level and whether they would allow or deny a given scenario. As reported in prior research [135], we also find a moderate correlation between surprise about some deployment scenarios and comfort with these scenarios. On the other hand, correlation between allow/deny decisions and desire to be notified seems nearly non-existent, suggesting people’s notification preferences do not simply correspond to their allow/deny preferences across different scenarios. An example of this case was mentioned in the previous section: only 30% of participants would deny data practices for generic surveillance purposes, but 60% reported that they would like to be notified. Our qualitative results in Section 3.4.3 and Figure 3.7 also seemed to suggest that individuals’ notification preferences are rather distinct from their allow/deny preferences, and serve different needs.

	comfort	surprise	notification	allow/deny
comfort	1			
surprise	0.442***	1		
notification	0.183***	0.214***	1	
allow/deny	0.604***	0.350***	0.046	1

Table 3.7: Correlation matrix where *** indicates $p < 0.001$

3.5 Privacy Concerns and Attitudes

In this section, we present findings from qualitative analysis of interview and textual response data collected from evening reviews of in-situ scenarios participants received. We first present findings on participants’ attitudes towards facial recognition technology and the reasons behind their attitudes. We then show the perceived beneficial and concerning contexts of facial recognition usage. We also unveil participants’ concerns about the use of facial recognition, with a particular focus on privacy-specific concerns, as they are among the most prominent themes. Finally, we flesh out participants’ proposed actions in responses to these deployment scenarios.

3.5.1 Impressions of Facial Recognition

We first present findings on participants’ sentiment towards facial recognition technology. This is based on our coding of sentiment in participants’ responses to the first question in the post-survey: “What is the first thing that comes to your mind when you think about facial recognition technology?”

Participants tend to be more negative towards FR

We observed that participants tended to be more negative towards the use of facial recognition: 51 (42%) participants displayed negative impressions while only 13 (11%) expressed positive sentiments. The negative connotation mostly revolves around problems of the technology, like the infringement on their right to privacy. Those negative first impressions also echo entrenched perceptions on problematic usages and privacy risks of facial recognition that are revealed in our subsequent analysis in Section 3.5.4 and Section 3.5.5.

Among the 13 participants with positive impressions, most praised facial recognition's usefulness, like its ability to increase public safety and catch criminals. A few also mentioned the "*advancement in technology*" (P36, positive). We also noted a mixed perspective of facial recognition from 11 (9%) respondents: "*It's invasive and big brother esque. It can provide good information for law enforcement but is easily abusable*" (P83, mixed). 48 participants (39%) indicated their neutral impressions typically by describing main use cases or depicting how facial recognition works: "*the ability of computers to see normal people in plain view and identify their identity. This can then be passed to another decision-making system for a distinct purpose: law enforcement, advertising, efficiency, etc.*" (P12, neutral).

Participant views may be influenced by media portrayals

A few concepts also emerged from these responses, mostly related to media portrayals of facial recognition. Some participants were reminded of what they have watched in the movies or crime shows relating to facial recognition: "*I think of face scanners and searches people do when looking for criminals in crime tv shows and movies*" (P42, neutral). Other respondents made references to a dystopian world, with many citing the concept of Big Brother from the book 1984 — "*Cyberpunk dystopias, "Big Brother," and similar instances in fiction, satire, and socio-political discussion about invasion of privacy on the part of powerful political and economic entities*" (P39, negative). China was brought up 7 times as the example of a surveillance state, which was associated with more negative sentiments (5 out of 7) than neutral tones (2 out of 7). For example, P80 alluded to a negative use case, "*China and the way they micromanage their citizens lives,*" and P5 expressed a more neutral impression: "*I think of China because the only times I've seen it on the news, it was being used in China.*"

In summary, respondents expressed more negative views about facial recognition than positive ones. Many were wary about potential problems linked to the technology. Around a quarter of participants' views were influenced by the media portrayal of facial recognition (e.g., news, movies, and books).

3.5.2 Beneficial and Concerning Contexts

We present findings on users' perceived beneficial and concerning use of facial recognition. This is based on the deductive coding of textual responses to the questions asking participants to identify up to 5 contexts each where they found the use of facial recognition technology to be beneficial and concerning. On average, each participant identified 2.7 ± 1.4 beneficial contexts, and

3.0 ± 1.4 concerning contexts. A paired Wilcoxon signed-rank test showed that participants recorded significantly more concerning contexts than beneficial contexts ($Z = 2.65, p < 0.01, r = 0.24$).

The findings are organized based on the major codes in the codebook, as shown in Table 3.8. These codes were further categorized into two groups: purposes for using facial recognition and entities that use facial recognition. We first report beneficial and concerning purposes in this subsection.

Beneficial purposes: security, authentication, and commerce

The majority (104 out of 123) of participants reported that security is a beneficial context for facial recognition. Among those, 42% thought that facial recognition could increase public security in general, and 32% thought that it is beneficial to use facial recognition to identify and catch criminals. Another important context for security, raised by 20%, is to find missing individuals. For example, P26 mentioned that facial recognition could be helpful in *“locating missing/abducted children and adults.”* 13% of them also mentioned that facial recognition could be beneficial to deter crime, as expressed by P27 *“in public, especially in isolated places like parking garages, to help preserve women’s safety.”* Another context for facial recognition that 51 participants (42%) identified as beneficial is authentication. About half of them (24 out of 51) stated that facial recognition could be used to replace IDs and confirm identity. 31% mentioned that it could be used to log in devices and/or replace passwords. A quarter maintained that facial recognition could be useful to grant access in secured locations, which P46 described as *“helping identify people in high-security areas.”* 14% considered authentication in stores via facial recognition as a way to replace membership or reward cards to be beneficial as well. A sizable minority (27 out of 123 — 22%) of participants also saw merits in leveraging facial recognition in commercial settings; using facial recognition to improve services and tailor customer experiences was deemed beneficial by about half of those 27 participants, for example, in contexts like *“relocating people between the crowded check-out areas”* (P63) and *“customization of service based on who you are and known preferences”* (P55). Others considered marketing and tailored advertisement of potential benefit, like in *“retail scenarios (catered advertising)”* (P46) and *“providing information to retail companies about their customers”* (P111).

Concerning purposes: advertisement, profiling, and prediction

Most participants (64%) raised concerns about various purposes for which facial recognition is used. Specifically, 36 out of 123 (29%) participants found using facial recognition for advertisement troubling: P117 said, *“It can be used for marketing and branding purposes that are generally antagonistic.”* 18 participants were concerned about facial recognition used for profiling — *“using it to profile someone based on race or gender”* (P21). 17 respondents found *“when emotion recognition is in use”* to be concerning. 12 participants (10%) were specifically against their data being sold for profit *“random companies selling and profiting off of it”* (P40). 11 were worried about use cases of facial recognition that involves predicting or estimating intentions or behaviors — *“Any assessments that are psychologically based since there is a lot that could be wrongly inferred by only taking into account visual data”* (P12).

Purpose				Entity			
Beneficial		Concerning		Beneficial		Concerning	
Code	%	Code	%	Code	%	Code	%
Security	84.6	Ads	29.3	Law/Gov	14.6	Law/Gov	18.7
Authentication	41.5	Profiling	14.6	Public	11.4	Employer	17.1
Commercial	22.0	Emotion	13.8	Health	8.1	Business	15.4
Personal	9.8	Profit	9.6	Employer	5.7	Insurer	14.6
Other	8.1	Predictive	8.9	Myself	5.7	Health	7.3
		Security	5.7	Business	4.9		

Table 3.8: Codes from Content Analysis and the Percentages of Participants Who Mentioned Them

3.5.3 Beneficial and Concerning Entities

The right-hand side of Table 3.8 shows the percentages of participants who identified different entities (law/government, employers, etc.) as beneficial and/or concerning when they deploy facial recognition.

Weighing between beneficial versus concerning

It is interesting to observe that law enforcement/the government were deemed concerning and beneficial both by a sizable number of respondents, which is also similar in the case of health-related entities (e.g., hospitals and clinics). The neck-and-neck numbers seem to suggest that those entities entail both rather apparent pros and cons of using facial recognition. For example, “*law enforcement falsely accusing someone*” (P83) is rather concerning, while facial recognition aids “*law enforcement to track and apprehend criminals*” (P42) is clearly beneficial. On the other hand, significantly more participants considered businesses, employers, or health insurers’ use of facial recognition more concerning than beneficial. More participants see harm than benefit brought by facial recognition usages by these entities, as elaborated by P59, “*The data collected seems worth more to the company than any coupons could possibly be for me.*”

Attributes influencing attitudes towards entities

The interview data revealed in-depth deliberations participants had while weighing various entities obtaining their facial recognition information. **Trust** was one of the factors that can erase participants’ doubts about potentially questionable facial recognition usages. Two interviewees explained why they trust their employer or the government/law enforcement, therefore trusting their use of facial recognition. P55 explained, “*I trust my manager personally to have my own interests in heart...Right now, personally, I have a good relationship with my manager and with the company. So I am pretty comfortable with what they do, decide to do, and feel like that they are not going to use it against me.*” Believing in the democratic government, P57 maintained, “*The government supposedly is "by the people, for the people" as supposed to private corporations...So if it’s used by law enforcement, I am a bit more comfortable with that.*”

More evidence on trust being an influential factor also emerged in the answers from evening surveys: “Because law enforcement and the government have a history of using data for purposes other than what they were intended for or what we were told it was for”(P26), “I don’t trust insurance companies to make fair decisions”(P116), “I trust the library mostly not to do anything bad with the video” (P97), “This is a large entity that I trust”(P51), and etc.

Besides trust, whether entities that deploy facial recognition have **control** over data subjects is another important attribute. Three interviewees were reserved about their employer or the government using this technology as those entities intrinsically have more **control** over them. In their views, facial recognition can be used against them by powerful entities, such as governments, employers, and big corporations, as expressed in the following quotes.

“I am used to people that advertise to me, trying to sell me something...I have more control over that relationship because I can always turn down buying something, even with coercive tactics that are manipulative. But with my boss or the government, I don’t have the power in that relationship at all. So it’s more information for them that they can use against me basically.” — P50

“I mean whoever’s behind it [facial recognition] has more data and information, what people need, what individual person wants, and how to best serve the people around, like get their product to the people. And also they have more control...over their customers.” — P52

Three interviewees were worried about advertisers’ or corporations’ usage that could decrease their sense of **autonomy**. Thanks to facial recognition technologies, businesses would leverage highly fine-grained and even real-time data to improve their marketing techniques. For example, P56 expressed her concern, “With the ability to read your reactions and then be able to market responses specifically to you, you might be losing some free choice. Because they are able to pinpoint and push harder things they think are important to you, because you are reacting to them, they can get real-time reactions to products...They can start using terms that look like something and trick you into buying something.” Such practices can be manipulative and encroach on people’s freedom.

3.5.4 Concerns About Facial Recognition

Participants were concerned about facial recognition even for anonymous demographic detection

Current facial recognition software enables different levels of identification: some can recognize the shape of faces and humans; some can detect specific demographic features; others can match faces to images of people stored in databases. Demographic detection has been used in contexts like targeted advertising and marketing [67, 79, 181, 219].

When designing the study, we initially conjectured that people would be more comfortable with anonymous demographic detection than personally identifiable detection. Nonetheless, 9% of participants expressed reservations about using anonymous demographic detection for advertising as they saw it as a form of profiling. P50 explicitly pointed out, “I was also pretty concerned when the notifications popped up about predicting purchases based on racial classifications because that just seemed very racist to me. Just because someone is African American or Hispanic, you can’t predict what they are going to want to buy based on their race; that seems a really not very good policy.”

Others were really against gender-based advertising. For example, P50 mentioned, *“And gender, there is such a spectrum, just because you’re female, that doesn’t mean you are going to wanna wear makeup or buy pretty dresses. Same thing for guys. I just think lumping every person into a classification is over-generalized; you are going to miss people.”* Some participants questioned the efficacy of advertising based on gender and race, *“ I wouldn’t think it will be very accurate, you could target something to me being white that would not at all relate to me still based on that one factor. But it may relate to a non-white person. I think it wouldn’t even be accurate. I think you need a lot more than race and gender to advertise to someone effectively”* (P106). This type of practices, even though beneficial at times, can also reinforce existing gender and cultural stereotypes — *“I understand that some ethnic groups might benefit from this (for instance, African American women need specific hair care products that aren’t always easy to find.) But I am concerned about the potential for misuse of this technology to discriminate. Also, people don’t always “look like” the racial or ethnic background with which they identify”* (P27).

Some participants, including some parents, were leery of age-based advertising, especially worrying about kids being susceptible to those practices. *“Things are marketed to kids nowadays, and kids can buy things on apps without their parents even knowing...I don’t think they should be marketed towards kids necessarily”* (P50). We also observed reservations from participants who were afraid of being labeled as a specific demographic group, such as religious groups. P53 said, *“I think it is kind of dangerous to pinpoint one person as part of a group vs. just the individual. So I think the times I was most concerned during the research was when I would go to someplace that was religious[ly] affiliated or like a non-profit organization. If there was a video of me and my friends maybe at a church or at a Jewish organization. Does that put us more in danger if we are associated with that group? I feel like there is this danger of having a label placed on you, and if the wrong person gets that information, and that could be a catalyst for violence.”* P89 summarized her feelings towards demographic-based facial recognition, *“I do think it will divide us more if they are targeting specifically based on what you look like, not even necessarily your profile and who you are...I think it just gives an overall weird and gross feeling, especially in today’s society where it comes up a lot.”*

Participants were worried about incorrect detection and interpretation

About a third of the participants reported their concerns about the accuracy of facial recognition during the study. Some were worried about the technology not accurate enough and could make *“mistakes in the face recognition (twins, relatives)”* (P65). One interviewee P107 shared his firsthand experience with inaccurate facial recognition in details, *“I don’t know how accurate they would be based on stuff that I have tested out before. Like even with having a beard, it throws off a lot of things that try to guess things. Actually, at work, just for fun, one of the guys had it. It is for visually impaired people who are blind. It scans anything and tells you what it is. It scans faces and got a lot of people like “39 male,” and it would be really close, but when it comes to me, it would say 40 where I am 25. It would say frowning even though I am smiling because of it tracking the mustache...if they are trying to pick up people with negative emotions for security purposes, maybe it could be pretty wrong.”* Others also echoed their doubts about the accuracy of emotion detection, like P68 *“I don’t see how it (emotion analysis) could be that accurate unless you are monitoring*

what I am saying too. Like I said, I went through a breakup that week, and sometimes I was not in a good mood no matter where I was, no matter how good the food was. How are they supposed to know? It just seemed like it was an unnecessary addition that wouldn't end up being very accurate."

In addition to questioning how accurate facial recognition can be, some participants also argued that seemingly suspicious behavior, when viewed out of context, can be misinterpreted by those systems, potentially resulting in grave consequences. For example, P53 described one such scenario in her friend's life that could be misconstrued, "*I think a lot of the times like my friend she locked herself out of her apartment this past weekend, so she tried to jump in through her window. So if a recognition program saw that, they might think that it is a thief or criminal or whatever. And that is not the case. She is not breaking into her own house. It needs to be able to interpret scenarios correctly. It needs to be able to have a context for them. Not just to assume that something looks like a criminal act is a criminal act."* Similarly, P68 gave another example, "*I think it could misinterpret scenarios, it could misinterpret the guy trying to break into his own car to get his keys out, or the boyfriend putting his hand in the girlfriend's pocket."* An interviewee P57 was worried about such inaccuracies leading to deadly consequences — "*because if someone was marked for shoplifting and they didn't do, that could cost a lot of trouble, in some scenarios that could cost someone's life."*

Participants were concerned about racial and other biases introduced by facial recognition

One-tenth of our participants reported being concerned about potential bias in the facial recognition systems, especially about the deep implications it might have on minority groups. Many were worried that racial bias in these algorithms could exacerbate the entrenched bias and infringe upon the rights of those impacted groups. Two interviewees' elaborated accounts provide us with more insights: P68 stated, "*Any system I've seen has inevitably been used only to profile people of color and the LGBTQ+ community. I think even if we have this surveillance, somebody is like, "Oh, it is just gonna automatically detect petty crimes." The reality is that it will still be looking harder at a black person and their actions to see if that is a petty crime than it could with a white person. I still think at the end of the day, a human is gonna analyze the data. I think you still have a lot of misidentification where people of color and LGBTQ+ community members are going to be scrutinized more strongly, not given the benefit of the doubt that white people are."* Similarly, P53 noted, "*I wouldn't want a program like that to decide that for example, a black man equals thief or even to give a warning sign to a program to flag that because that is not the case. So I think that is the danger of having that type of use for facial recognition. I think it can too easily be biased, intentionally or unintentionally. The person programming it might think that they might have statistics to back up the demographics of thieves or demographics of criminals, but I don't think that is a good way of deciding who is or who is not a criminal."*

3.5.5 Perceived Privacy Risks of Facial Recognition

Privacy is repeatedly brought up as a key concern by our study participants. Around 70% of participants voiced privacy concerns during the study. In this section, we summarize the major themes around perceived privacy risks of facial recognition, in light of concepts from established

privacy frameworks (i.e., Solove’s “Taxonomy of Privacy” [220] and Westin’s states of privacy from *Privacy and Freedom* [245]).

Violation of solitude

The feeling of surveillance prevails A third of our respondents found surveillance through facial recognition to be concerning. Surveillance can exert adverse psychological effects like discomfort and anxiety on subjects. For example, P68 pointed out that *“I had this paranoia that I would be judged based on every action I took at work without the full context.”* Similarly, P29 stated that *“always being watched and analyzed which in itself is scary.”* Moreover, surveillance is also harmful due to its infringement on people’s freedom to act. P89 contextualized this concern — *“There is a feeling of freedom as I enter the library where I participate in a Spanish speaking group on Wednesday morning...in the small classroom where we speak, I would feel rather self-conscious if I were videoed.”* This infringement upon freedom can also possibly lead to inhibition and behavior alteration, as P84 noted *“I’d always have to be concerned about how my actions might be perceived on camera,”* and in P20’s view, *“I want to know where all of the cameras are, so I can always be aware and I can always be on guard and vigilant. So if something happens, I can be ready to defend myself or defend the findings.”* Surveillance can also have a chilling effect on civil and political engagement. For instance, P117 pointed out that facial recognition *“is used to identify anti-fascists and peaceful protesters”*, and P39 found *“any and all efforts at using such technology against political dissenters”* to be concerning.

Deprived of the right to be let alone Warren and Brandeis first articulated privacy as the “right to be let alone” [244]. Privacy risks also lie in the probing action itself which perturbs this right, making “the person being questioned feel uncomfortable” as noted by Solove [220]. Two-fifths of our participants regarded some deployment scenarios of facial recognition as unwarranted and prying. For instance, P68 manifested their concern, *“It is the idea of somebody being able to surveil my life and know my business...Even though on sight it’s something different through a camera, that knowing somebody is interested in the data, and wants it, and is just getting it for free. Something about it really bothers me.”* Some participants responded to data collection of facial recognition rather abruptly, *“It’s none of anyone’s business, as long as I’m obeying the law, where I am and what I’m doing”*(P114). Some participants reported that facial recognition is intrusive into one’s life, and they cannot be let alone under the presence of facial recognition. For example, P83 mentioned that they are *“unable to hide from people”*, and P104 noted, *“I feel like I’m being stalked by the man, the powers that be, wealthy corporations.”* Others regarded facial recognition as disruptions to their daily activities: P69 mentioned, *“Don’t want to be filmed eating,”* and P62 commented on their experience in stores, *“It’s like being stared at in the face by someone while I’m just trying to shop.”*

Unwanted exposure and violation of anonymity

Not able to stay anonymous 17% of participants stressed the importance of anonymity and scrutinized how facial recognition enabled the identification of normal people in plain view. P63 gave examples of circumstances when people may want to stay anonymous, *“Probably if you go to*

some kind of clinics, like sexual health clinics, or food pantry.” P12 voiced their concerns about facial recognition used for advertising, “If it is generating tailored advertising then it implies it is tracking my shopping habits and linking it to my face.” P55 elaborated a situation when he wants to remain anonymous, “I don’t do any sort of very secretive things. The only possible scenarios are if I was trying to...plan a surprise birthday party for my wife, some notification got sent to both of us of where I was, and then she figures that out...There is a mixed scenario of people who are doing slightly illegitimate things but are legal to do, like having affairs with people on their partners, they would definitely not like stuff like that.” Identification, a method to connect people to collected data, is hard to avoid as the deployment of facial recognition technologies becomes widespread.

Unwanted exposure to others This issue involves “exposing to others of certain physical and emotional attributes about a person,” which often “creates embarrassment and humiliation” as defined by Solove [220]. 22% of our participants pointed out that it is easy to reveal emotions under contexts of facial recognition involving emotion recognition. For example, P68 described her personal experience, *“I went through a breakup that week. I was really emotional a lot of the time. I do not want my health insurance, my employer, my parents getting updates like “hey, she’s trying to get through the pain while she is working today.””* P50 commented on a facial recognition scenario that occurred at the vet they went to, *“People experience deep personal emotions at the vet.”* Some respondents were cautious about carrying out private actions. For example, P89 elaborated, *“I might be caught at the gym entering and adjusting a bra strap, etc.,”* and *“doing something like picking your nose, something like that, not doing something against the law, but something you don’t want others to see.”* Such unwanted exposure in public spaces might not have been feasible without facial recognition technologies.

Non-consensual and insecure disclosure

Secondary use without consent This refers to the privacy issue of data collected for additional purposes without data subjects’ knowledge or consent. In the context of facial recognition, this problem is exacerbated because of the lack of ways to properly convey data practices to subjects other than using signs that say “face recognition security cameras in use.” Given the sensitive nature of facial recognition data, around a quarter of our participants reported concerns about unauthorized secondary use. Many respondents questioned whether companies would retain data for intended use only, as P12 described, *“As I’m doing this study more, I think it’s my trust in their ability to safe keep the data and only for that use. I would doubt their compliance even if I do want them to get the competitive advantage by the use of video surveillance.”* P89 also hoped for regulations to prevent secondary use, *“If there were laws in place that they could never ever use it for anything else like they couldn’t sell it to marketing companies.”* P106 provided a concrete example of secondary use with regards to workplaces using facial recognition to track attendance, *“I think if used to replace a time card is fine, but I could see it being abused by overbearing managers.”* A few participants expressed concerns about their data being sold, which can also be regarded as a secondary use.

Fear of data leakage and abuse About one-third of our participants expressed their concerns about their facial recognition data being hacked or abused. Because it is almost impractical to relinquish biometric data when compromised, the security of facial recognition data is ever more pressing. Many of the participants reported that they do not trust data collectors' ability to safeguard their data. For example, P122 noted, *"I don't think data security is a strong priority for these companies, and when they do have data leaks, they don't care because it doesn't affect them, and the punishment is not enough to incentivize them to change their practices,"* which parallels the concerns of P54 about identified frivolous activities being leaked, *"Frivolities that end up being insecure, like entertainment or stores."* Also, the fear of insecurity can induce privacy risks by placing people to whom it pertains in a vulnerable state, as corroborated by P122, *"It's very troubling to think of how this info could be used by bad actors."*

Inaccurate dissemination and violation of reserve

Dissemination of inaccurate or misleading information Around one-third of our participants were concerned about the dissemination of inaccurate or misleading information [220]. This issue is also mostly linked to the inaccuracies of facial recognition as presented in Section 3.5.4. Our participants were concerned about being falsely identified or judged out of context. For example, P46 noted, *"Bad luck or timing could lead law enforcement to be suspicious of an innocent citizen."* P11 referred to their experiences when shopping in stores, *"I would really not like supposedly meaningful data to be recorded if I happened to smile remembering something while walking down the condom aisle."* Distortion can be detrimental, as illustrated by P59, *"Reputational damage could occur if someone is falsely accused of a crime."*

Decisional interference Solove defined this as the intrusion on private decisional making, especially by the government [220]. In our study, participants mostly focused on the unwarranted influence on their purchasing autonomy by private companies with the help of facial recognition. This is also discussed in Section 3.5.3. In addition, P89 lamented, *"It's machines taking over and my freedom circumvented."* P122 echoed this thought, *"I do not want to have this information used against me or used to try and subvert my thinking."*

3.5.6 Proposed Actions and Responses

Our qualitative data also reveals participants' reported desire to take action when encountering facial recognition in their everyday life. They also express a desire for transparency and indicate they would like to be notified about nearby deployments of facial recognition technology. At the same time, their notification preferences vary with some participants expressing concerns about potentially overly disruptive notifications.

Participants want transparency and control over the collection of their data

About 30% of participants expressed strong views about the need for entities collecting sensitive facial recognition data to notify them and to actively obtain consent from them before data collection.

For example, P50 commented, *“I think if they are going to record our image, they should have to notify you before they do anything with it like if they are going to use it for a specific purpose, we should be able to know what they are using it for, and we should be able to say “yes, that’s fine,” or “no, it’s not. Delete my stuff from your system.”*” While most participants agreed about the need to obtain consent, they did not provide consistent answers with regard to the frequencies of such notifications. Some participants wanted to be notified every time when such data collection is taking place, as illustrated by the quote from P56, *“I think it is important to know when you are in areas where data is being collected, passive consent really disturbs me. I know it happens all the time when I am on my phone or computer, and it is really hard to know what data is being collected, what it is being used for, etc....So, if I have my preference, I would want to know every time someone is engaging in this practice,”* whereas others were wary of repeated reminders and preferred less frequent notices, as P17 elaborated, *“I frequent this establishment pretty often, so a constant reminder would annoy me. It would be nice to be reminded every now and then in case I simply forget.”* These results suggest a need for customizable notification functionality where different individuals can select from a number of notification options.

Participants find existing notice mechanisms inadequate

While the majority of participants wanted to be informed about facial recognition in use, our follow-up interviews disclosed the specific ways how some participants found the existing notice mechanisms inadequate. For instance, P68 described how they missed the existing signs in physical spaces that were supposed to notify them about the presence of cameras, *“There will be places where I would want to be notified every time, and then I look over, and see a sign that I have just passed by a dozen times, and realize I am being notified.”* When probed about what is a good way to give them notice or obtain their consent, some interviewees reported that no existing mechanisms would achieve the goal, as P53 said, *“I think that [obtaining consent] is hard...It is hard because you cannot pass a form when you walk into a restaurant or a store, it cannot be formal...I guess trying to do it remotely like through the Internet or your phone would be the easiest.”* Specifically, P50 expressed their desire to provide consent based on different purposes of facial recognition, *“It would depend on what they were using it for. If it was just like someone committed a crime, and they needed FR for that, then that’s fine. Maybe if it’s to replace a swipe card or a membership cards, that would be okay, but if it’s for tracking my purchases, or tracking my attendance, emotions.”* The information on the purposes for which facial recognition is deployed is not available to data subjects in the majority of current deployments. Also, it is also hard to design notice mechanisms with the desired level of intrusiveness, as P89 elaborated, *“I would not want to think about it at all times, so I want it to be subtle whatever the notification is, but also not so subtle that you don’t know that it is happening ever,”* which highlights the problem of privacy as a secondary goal.

Some participants fear being overwhelmed by frequent notifications

While most participants report that they want to be notified, more than half are also weary of too frequent notifications. In particular, some participants realized during the course of the study that the number of notifications they would receive might become a nuisance if they request to be notified

each time they get within range of facial recognition technology. For instance, P53 described her thought process, *“When I first started, I was saying once in a while, and then I realized that would be really annoying to get multiple notifications.”* About half (55 out of 123) of participants reported that they were unlikely to avoid places that deploy facial recognition technology, even if they indicated being concerned about these deployments, revealing a general sense of resignation. For instance, P11 underscored, *“There is nothing I can do about it, and this is the only accessible grocery around my workplace, so I don’t have an alternative.”* A similar sense of helplessness and resignation was expressed by P67: *“I give up. Spy on me. What can I do about it? I’m old. I’ll be dead soon.”*

At the same time, not all participants reported concern. We also observed a small number of participants who did not care about the usage of facial recognition in general, referencing the “nothing to hide” argument. For instance, P55 elaborated, *“I am not likely to be so concerned about it, because I don’t do any sort of very secretive things... There are more legitimate reasons why people would want to value their privacy more than I do, but I am not sure how much of the population that would really affect.”*

3.6 Exploring the Development of Predictive Models

Under regulations such as GDPR data subjects are supposed to be notified and agree to having their footage captured by video analytics software at or before the point of collection. Because of the increasingly widespread deployment of video analytics software, if data subjects are asked to manually opt in or out of video analytics each time they encounter such functionality, they will not only face an unrealistically high number of decisions, but can quickly become annoyed or desensitized as we have observed in the 10-day study. Recent technical advances introduced in prior work by Das et. al [50] open the door to scenarios where a user, with a “privacy assistant” app running on their smartphone, would be alerted to the presence of video analytics software and would be given the choice to opt in or out of such processing. In this section, we use our data to explore the feasibility of developing predictive models to assist users in managing these privacy decisions and discuss different possible deployment strategies for such models. Specifically, we focus on the development of models to predict people’s allow/deny decisions across the different types of scenarios using data collected as part of our in-situ study.

3.6.1 Feature Selection and Clustering

As discussed in Section 3.4.2, purpose appears to be the most significant attribute when modeling people’s allow and deny decisions. Accordingly, we develop models that use purpose as feature—it is likely that more complex models could be developed with possibly even better results. As prior work showed promising results of clustering like-minded users in the mobile app permission space [136, 139], we adopted a similar approach and applied agglomerative clustering with ward linkage on the feature vectors to cluster participants. After we obtained the resulting clusters of users, we calculated the privacy profiles of each cluster using two-thirds majority vote. If more than two thirds of participants in a given cluster allow (deny) a given data practice for a particular

purpose, then the cluster profile recommends allowing (denying) that practice for that particular purpose. If there is no majority decision, or the number of data points in the cluster for the particular practice and purpose is too small, the cluster profile does not recommend allowing/denying the practice for the given purpose (i.e., no recommendation).

3.6.2 Predictive Power of Cluster Profiles

We want to evaluate how well the cluster profiles generated could help predict people’s allow/deny decisions for incoming users not present in the clusters. We first randomly select 90% of the participants to build clusters as described in the previous section, and use the remaining 10% of participants to evaluate the predictive power of the clusters by calculating the following two metrics *accuracy* and *efficiency*. *Accuracy* is defined as the percentage of time the prediction of a cluster profile (when such prediction is available) matches the actual allow/deny decisions made by users assigned to that profile. We define *efficiency* as the percentage of allow/deny decisions made by a user for which the assigned cluster of the user offers a prediction (or recommendation). In other words, if for every allow/deny decision a user needs to make, the cluster to which the user is assigned offers a prediction, efficiency is 100 percent—theoretically the user does not need to manually make any decision, though the accuracy of the predictions could be less than 100 percent, as some predictions could be erroneous.




Figure 3.8: Accuracy and efficiency of models plotted against the number of clusters used to build them.

We repeated 10 times the process of generating clusters from randomly drawing 90% of participants, and of evaluating the predictive power of these clusters using allow/deny decisions of the remaining 10% of participants. Average *accuracy* and *efficiency* results are shown in Figure 3.8. As can be seen, there is a substantial increase in both accuracy and efficiency when we move from a global one-size-fits-all profile (single cluster) to models with two or more clusters. We can observe the trade-off between efficacy and accuracy as the number of clusters grows. Accuracy increases with the number of clusters, as these clusters become more targeted. Yet, efficiency decreases given that, as the number of clusters increases, the size (or population) of each cluster decreases,




Figure 3.9: Profiles associated with a 6-cluster model. Each cluster profile contains 3 columns: the left one displays the average mean value (deny=-1, allow=1), and the right column represents the cluster profile, where the blue color represents an allow decision, red means a deny, and white means no decision, either because not enough data points are available or for lack of a two-thirds majority. The middle column shows the variances, ranging from 0 to 1. The 3 numbers (D/A/T) in each entry in the the right column represent the distribution of deny ("D") and allow ("A") collected for members of the cluster for the corresponding purpose, with T=D+A representing the total number of decisions collected for the given purpose from members of the cluster.

eventually making it more difficult to generate predictions as some entries have too few data points to obtain majority voting. The results for six clusters seem to provide the highest harmonic mean of accuracy and efficiency. It is worth noting that a model with 6 clusters achieves an efficiency of 93.9%, namely the clusters are able to predict 93.9% of the allow/deny decisions our participants had to make with an accuracy of 88.9%. It is likely that with additional data, more complex models, taking into account additional features beyond just purpose, could achieve even greater predictive power.

3.6.3 Example of Cluster Profiles

As shown in Figure 3.8, one-size-fits-all models based on lumping all users in a single cluster fail to capture the rich and diverse responses towards facial recognition deployments captured in our study. However, models obtained by organizing participants in a small number of clusters seem to achieve much higher predictive power. Here we look at the profiles associated with a 6-cluster model, (see Figure 3.9), namely the model that yielded the highest harmonic mean in the previous section, and discuss what these profiles tell us about how people report feeling towards different deployment scenarios.

As can readily be seen, participants in Cluster 1 and Cluster 5 represent polar extremes, with participants in Cluster 5 indicating they would largely respond with an “Allow” to all the deployment scenarios covered in our study, whereas participants in Cluster 1 would largely respond with a “Deny” to all these scenarios. It is worth also noting the low variances found in these two clusters for most deployment scenarios, indicating that people’s responses in these clusters tend to be particularly homogeneous. All other clusters also exhibit low variances for many scenarios, though each of these other 4 clusters has a few scenarios for which responses are less homogeneous, with each of these other 4 clusters having one or more deployment scenarios where the model is unable to

make a prediction (e.g., “Rate Service (Anon)” in the case of Cluster 4). Comparing Cluster 3 with Cluster 5, we see that like in Cluster 5, participants in Cluster 3 tend to respond with an “Allow” to scenarios associated with a variety of different purposes, except when it comes to sensitive purposes like tracking attendance or evaluating work productivity. They tend to also be more reticent in the presence of facial recognition scenarios designed to support health predictions. Members of Cluster 2 exhibit significantly more conservative responses and are generally uncomfortable with a much larger set of deployment scenarios than members of Cluster 3, though they appear to be fine with the use of facial recognition to capture demographic information in support of anonymous targeted advertising scenarios (e.g., adjusting the ad shown in a store window based on demographic features of the person looking at the window [67, 79, 181, 219]). In comparison, members of Cluster 4 seem to exhibit somewhat different sensitivities. While they too object to many deployment scenarios, they appear to be fine with the use of facial recognition to fight crime and to also anonymously count people.

3.6.4 Possible Application in the Context of Privacy Assistants

The above analysis sheds some light on how different groups of people share many privacy preferences when it comes to opting in or out of different video analytics scenarios and how these preferences vary across different groups. The privacy profiles can also function as meaningful default settings in privacy assistants. Users would just be asked a few questions, which would be used to suggest a particular profile to them. This profile could include a hundred or more privacy settings, which would accurately capture many of their preferences. Users could then review these settings and edit them as they see fit. This approach proved quite effective in the context of a privacy assistant developed to recommend mobile app privacy permission decisions, with most users indicating that they liked the functionality and adopting most of the recommendations made by their privacy assistant [137, 139]. In the context of mobile apps, recommendations can be organized by categories of apps. In the context of video analytics scenarios, recommendations could be organized based on a taxonomy of relevant contextual attributes such as the 16 different purposes considered in our analysis. As can be seen in Figure 3.9, most profiles have recommendations for the vast majority of the 16 different purposes considered in the study. This means that, depending on the particular profile to which a user would be assigned, that user would have recommendations for most of the privacy decisions they would encounter. They would obviously be able to review and adjust these recommendations, as they see fit. Assuming video analytics deployments with standardized APIs to communicate relevant contextual attributes and to process opt-in/opt-out decisions selected by a user, a privacy assistant would be able to automatically communicate the user’s opt-in/opt-out choices for the particular video analytics deployment a user encounters, saving the user the burden of communicating the same decisions over and over again as they keep on running into similar video analytics deployments.

3.7 Discussion

3.7.1 Limitations

We do not claim that our results are representative of the general population. Our sample population skews young and more educated, which could have induced bias in our results. In addition, participants were recruited only from Pittsburgh, Pennsylvania, a mid-sized city in the United States.

Our study protocol determined the type and frequencies of scenarios participants saw, which in turn likely impacted their attitudes over time and in particular their notification preferences. We strived to keep the study realistic by presenting each participant with scenarios representative of the venues they visit in their everyday life. The actual frequency and types of video analytics participants would encounter could, however, be different from those in our study, and are likely to evolve over time. Our analyses were conducted using data provided by participants when presented with plausible deployment scenarios, rather than based on observations in the presence of actual deployments. While our use of an in-situ methodology was intended to mitigate this issue, it is possible that some of the data collected is not fully representative of participants' actual behaviors.

While describing study scenarios, we strived to maintain a balanced narrative without overly emphasizing benefits or potential risks associated with different deployments, but rather leaving it to participants to decide how they felt about them. This being said, we acknowledge that the phrasing of these types of scenarios is an art and that on occasions our phrasing might have primed participants in one direction or the other.

Our participants generally expressed somewhat negative views of various facial recognition deployment scenarios. This could, in part, be a reflection of the fact that they did not actually experience true interactions with these deployment scenarios and, as a result, may not have had a chance to appreciate what they consider as benefits associated with some of these scenarios (e.g., marketing scenarios).

Finally, we also acknowledge that more sophisticated predictive models could be built with even better performance. We purposefully limited ourselves to simple clustering solutions to emphasize that even simple models can produce strong predictive power in this domain.

3.7.2 Lack of Awareness and Desire for Greater Transparency

Our results clearly indicate that many people were taken by surprise when encountering a variety of video analytics scenarios considered in our study. While many expect surveillance cameras to be widely deployed, few are aware of other types of deployments such as deployments for targeted advertising, attendance, productivity, and more. These less expected scenarios are also those that generally seem to generate the greatest discomfort among participants and those for which, if given a chance, they would often opt out (or not opt in). These results make a strong case for the adoption of more effective notification mechanisms than today's typical "this area under camera surveillance" signs. Not only are people likely to miss these signs, but even if they do not, these signs fail to disclose whether video analytics is being used, for what purpose, who has access to the footage and results, and more. Our study shows that many of these attributes have a significant impact on

people’s desire to be notified about deployments of video analytics. And obviously, these signs do not provide people with the ability to opt in or out of these practices.

Our findings support new disclosure requirements under regulations like GDPR, which mandates the disclosure of this information at or before the point of collection. Our findings also demonstrate the urgent need to provide people with choices to decide whether or not to allow the collection and processing of their data, as our participants expressed diverse levels of comfort with these scenarios with many not feeling comfortable with at least some of them. Regulatory disclosure requirements help improve transparency of video analytics deployments. While some study participants grew more concerned about facial recognition, we observed others becoming more accepting of it as they learned about potential benefits of some deployments. These findings suggest that increased transparency and awareness would help data subjects make informed decisions.

3.7.3 Privacy Preferences Are Complex and Context-Dependent

Our findings show that people’s privacy preferences are both diverse and complex. They depend on a number of contextual attributes such as the purpose for using video analytics, who has access to the results, where the user is at the time of collection, and other factors. As such, our findings are another illustration of contextual integrity principles introduced by Nissenbaum [162]. The importance of purpose information identified in our study (i.e., for what purpose video analytics is being applied) is largely consistent with results reported in earlier publications. This includes earlier work conducted by Lin et al. [136] and Smullen et al. [217] in their studies of privacy preferences when it comes to configuring mobile app permission settings. This also includes prior work by Emami-Naeini et al. [158] looking at privacy preferences across generic IoT scenarios. In contrast to these earlier studies, our work took a more systematic approach to exploring the nuances in video analytics scenarios, including the type of analysis, the purpose for which the analysis is conducted, whether information is being shared with other entities, and the venue where video analytics is deployed; those factors all have an impact on individuals’ privacy attitudes.

3.7.4 Implications for the Design of Privacy Assistants

Our findings can also inform the design of privacy assistants that help users manage privacy decisions related to the deployment of video analytics and other Internet of Things (IoT) technologies. Das et al. have introduced a privacy infrastructure for IoT, where users rely on “privacy assistant” mobile apps to be notified about the presence of nearby IoT resources such as cameras running video analytics software [51]. Using these privacy assistants, users can access opt-in or opt-out functionality made available by IoT resources to indicate whether they agree or not to the collection and processing of their data. However, given the growing deployment of cameras, taking advantage of such functionality would still be hampered by the number of notifications and decisions a typical person would be confronted to each day when passing within range of cameras.

A more practical approach would involve allowing users to configure privacy assistants to only notify them about those deployments they care to be notified about, and to possibly also configure any available opt-in/opt-out settings in accordance with their individual preferences. Based on our findings, it is easy to see that different users would likely select different configurations of their

settings, namely different notification settings and different combinations of opt-ins/opt-outs. To keep user burden manageable, one would likely include settings that allow users to automatically opt in or out of scenarios for which they have pretty definite preferences (e.g., “I want to opt out of any video analytics deployment that shares my data with insurance companies”). For other scenarios, they would be notified and prompted to make manual opt-in or -out decisions. Given how rich and diverse people’s privacy preferences are, enabling users to accurately specify their notification and opt-in/opt-out preferences would require a large number of privacy settings (e.g., differentiating between a variety of different video analytics deployments, different notification preferences). Recent work on privacy assistants has shown that it is possible to use machine learning to reduce user burden when it comes to configuring such complex privacy settings. For instance, Liu et al. have demonstrated the use of machine learning techniques to help users configure mobile app privacy settings [139]. Similar results have been observed by the authors using data collected as part of the present video analytics study, where models of privacy preferences were built to predict participants’ allow/deny decisions [258]. The idea is that these models are used to recommend settings to users, who can review the recommendations and decide whether or not to accept them.

Our results showing that individual’s preferences for notification of video analytics deployments are quite diverse suggest that different people would select different setting configurations, with some people preferring to be systematically informed about each deployment and being prompted to manually decide whether to opt in or out, and other people preferring more selective notification settings and greater delegation of opt-in opt-out decisions. This is also consistent with results from a recent study by Colnago et al. [43] It goes without saying that effective implementation of notification functionality and opt-in/opt-out settings such as those we just discussed, settings that our findings seem to call for, would substantially benefit from the development of standardized APIs. Ideally such APIs would enable privacy assistant functionality to (1) discover video analytics deployments in the vicinity of their users, (2) selectively notify their users, and (3) transmit opt-in or opt-out requests on their behalf (whether these requests are made manually or derived from settings selected by users).

3.7.5 Evolving Notification Preferences

In our study, we observed that participants’ notification preferences evolved over time with many people opting for less frequent notifications as time passes. This change in preferences is attributed to some level of fatigue as people got a better appreciation for the number of times they were likely to be notified about the same or similar scenarios, and as their level of surprise in the face of some of these scenarios also diminished over time. Even taking into account this general trend in receiving less frequent notifications over time, it is clear that people’s notification preferences are not adequately met if one relies on a simple “Ask on First Use” approach—as is typically the case today when dealing with mobile app permissions, for instance. Individuals’ notification preferences are more complex and also more diverse, ultimately requiring a more sophisticated set of configurations that users could choose from and also modify over time, as their preferences evolve. Here again we see opportunities for the use of AI-based privacy assistant functionality [42, 139] that would adapt to their user’s preferences over time, possibly through a combination of nudges designed to motivate users to think about options available to them [6, 10] and dialogues designed to capture

people’s evolving preferences. Our study also uncovers how individuals’ allow/deny preferences are distinct from their notification preferences. However, how to properly notify people without overwhelming them remains an understudied direction as the majority of work on modeling privacy preferences focused on allow/deny “choice” rather than “notice.”

3.7.6 Combating Inaccuracy and Bias

While most of participants reported seeing benefits in facial recognition deployments such as security and authentication scenarios, their reported attitude towards many other scenarios was generally more negative. Part of their willingness to embrace the technology was dampened by concerns over accuracy and bias of facial recognition systems, echoing concerns voiced by marginalized interviewees in a prior study [98]. Our data suggest that these concerns extend to the more general population. Recent reports of people wrongly arrested due to faulty facial recognition algorithms likely contributed to reservations captured in our study [101] and also illustrate the severe consequences that deployment of this technology can have if deployed and relied upon without adequate safeguards. Minimally, technology should be evaluated for potential biases and minimal levels of accuracy, especially when deployed in support of particularly sensitive activities such as law enforcement. Their performance and limitations should be clearly communicated and taken into account. And decisions based on these algorithms should be meticulously cross-checked and manually vetted if we are to avoid more of these nightmarish scenarios.

3.7.7 Contextualizing Perceived Privacy Risks

Our analysis organized perceived privacy risks associated with facial recognition deployments around key dimensions identified in well-established privacy frameworks [220, 244, 245]. We were able to elicit more nuanced and contextualized privacy concerns than prior work [33, 216, 221] as shown in Section 3.5.5. While legal arguments support people’s reasonable expectations of privacy in public places [102], our study provides strong evidence that these expectations are real and widespread and that some facial recognition deployment scenarios are perceived as overstepping the boundaries of personal solitude, making people feel deprived of “the(ir) right to be let alone” [244]. These concerns are further exacerbated by the sensitive nature of biometric data, the information that can be inferred from facial recognition data (e.g., location, activity, and mood), as well as risks of secondary use of this data and its security. These findings underscore the need for more transparency in notifying people about not just the deployment of facial recognition technology but also sufficient details for individuals to gauge their perceived privacy risks.

3.7.8 Designing Effective Notice and Choice

Our study confirms that privacy concerns are a major obstacle to acceptance of a variety of facial recognition scenarios [33, 37, 186], although these deployments are becoming increasingly widespread. Responses from our participants indicate a strong desire to be notified about different deployment scenarios and to have some control over the collection and analysis of their data. Current deployments generally fall short when it comes to effectively notifying people about the presence of

facial recognition technologies, including details about the type of analysis they rely on and how results are being used and possibly shared. Also, current deployments generally fail to provide people with opt-in or opt-out choices.

How to effectively notify people and offer them adequate controls is not trivial. Entities deploying facial recognition should inform data subjects in a clear and noticeable manner. Today’s “this area under camera surveillance” signs do not provide them with enough information, such as type of analysis, the purpose for collection and analysis, sharing, etc. Privacy controls (e.g., opt-in and opt-out choices) should obviously include mechanisms to authenticate data subjects (to make sure they are whom they claim to be when they request to opt in or out of some practices), giving rise to privacy issues. With the possible exception of security-related deployments, which many view as generally beneficial, people should be offered some control over the collection and use of their footage — preferably in the form of opt-ins.

One solution involves requiring people to opt in by providing training data about their face [50, 51]. In this system, a privacy-aware infrastructure is used to notify people about the presence of nearby facial recognition deployments, including who has deployed the technology, what analysis is performed, and for how long the footage is retained. Users who do not opt in for facial recognition by default have their face (or possibly their entire body) obfuscated in real-time in the captured footage. Notifications about nearby facial recognition deployment are provided via a “Privacy Assistant” mobile app that users install on their smartphones. This infrastructure has been deployed to support notice and choice for a variety of Internet of Things data collection processes—not just facial recognition [51, 202].

Our data highlight individuals’ diverse notification preferences, with some preferring to be systematically notified about FR deployments, while others only would prefer just occasional notices and reminders. The Internet of Things Privacy Infrastructure introduced by Das et al. offers users of its “Privacy Assistant” mobile app different settings they can configure to specify the types of data collection processes they want to be notified about as well as the frequency of these notifications (e.g., “only the first time,” “every time,” or “never”). These settings are consistent with results discussed in Section 3.5.6, which indicate that different participants have different notification preferences and that these preferences can also evolve. Further research is needed to determine what personalized settings are likely to work best and how to alleviate the user burden that might be entailed by opt-in or opt-out settings associated with a potentially large number of facial recognition deployments.

Finally, our study indicates that participants fear losing their autonomy when commercial entities can assemble and leverage near real-time facial recognition data, including their emotions, to tailor advertisements presented to them. Our participants also expressed reservations about the power this technology can bestow on already powerful entities such as their employers or law enforcement authorities. These results further emphasize the need for more effective notice and choice mechanisms if people are to become less fearful about the deployment of facial recognition.

3.8 Summary of Main Contributions

In summary, research reported in this chapter yielded the following contributions:

- We conducted a first 10-day in-situ study of individuals' privacy expectations and preferences across a wide range of realistic video analytics deployment scenarios. We offered an in-depth analysis of the data collected as part of a study involving 123 participants who provided us with detailed insight into their degree of awareness and comfort across a total of 2,328 deployment scenarios.
- Our analysis revealed that many people have little awareness of many of the contexts where video analytics can be deployed and also showed diverse levels of comfort with different types of deployment scenarios.
- Notification preferences were also shown to be diverse and complex, and seemed to evolve over time, as people became more sophisticated in their expectations as well as in their realization of the number of notifications they were likely to receive if they were not selective in their notification preferences.
- Our qualitative analysis contextualized participants' perception of privacy risks associated with video analytics and explored their concerns about the limitations and bias found in some of these systems.
- We used the data collected as part of our study to explore the feasibility of developing predictive models to help people cope with the large number of allow/deny decisions they would otherwise have to make each time they encountered video analytics deployments. We showed that even using simple clustering techniques, it was possible to accurately predict many privacy decisions a user would want to make when encountering diverse video analytics deployments.
- We discussed the implications of people's rich and diverse privacy preferences when it comes to notifying them about different video analytics scenarios and to supporting opt-in or opt-out choices associated with the collection and use of their data under these scenarios. This included discussing different possible configurations of privacy assistant functionality to accommodate individuals' diverse privacy attitudes and preferences. In particular, we discussed how standardized APIs could be used to help recognize different types of video analytics scenarios and automatically communicate a user's privacy decision for that particular type of deployment, saving the user the burden of repeatedly communicating the same privacy decisions over and over again as they kept running into similar deployments. In addition, we also discussed how results from our study suggest the development of selective notification functionality, which users would be able to configure to match their particular preferences.

This work was published at PoPETS 2021 and SOUPS 2021, presented in PrivacyCon 2021, and received the 12th Annual Privacy Papers for Policymakers Award [255, 258, 259].

Chapter 4

A Contextual Integrity Analysis of Vaccination Certificates

4.1 Overview

In Chapter 3, we investigate the diverse privacy attitudes towards video analytics technology in depth. In this chapter, we turn to another recent development in society that has emerged in response to the COVID-19 pandemic: vaccination requirements and their potential profound privacy implications. The prolonged and devastating COVID-19 pandemic has affected every aspect of people’s lives as well as the global economy. In an attempt to curb the spread of highly contagious variants, governments around the world have contemplated or adopted vaccination mandates (VMs) and vaccination certificates (or passports) (VCs) in schools, hospitals, public transportation, and other social contexts [44, 87, 154, 157, 168, 171, 180]. Against the background of the global pandemic and increasing adoptions of VMs and VCs, it is important to ensure that the information flow rules embodied in the new technologies adhere to prevailing societal norms for each given context.

In our work, we want to explore how privacy influences the acceptance of vaccination certificate (VC) deployments across different realistic usage scenarios. Analysis of results collected as part of this study is used to derive general normative observations about different possible VC practices and to provide guidance for the possible deployments of VCs in different contexts. We use the Contextual Integrity (CI) [163] framework to evaluate the privacy implication of VC technologies in terms of appropriateness and legitimacy of information flows they generate. We aim to answer the following two research questions: 1) In what contexts are VC deployments and mandates perceived appropriate? 2) How does the practice of re-sharing VC information affect the perceived appropriateness?

Our study provides a comprehensive and nuanced understanding of people’s diverse attitudes and expectations towards VC deployment scenarios as well as the complex privacy decisions faced by society such as those we face in today’s fight against COVID-19. Beyond the blunt approach one often hears—that privacy must be traded off against public health—our findings open the door to more informed and nuanced alternatives that allow the pursuit of public health even as we reinforce appropriate information-flow practices that conform with the wide attitudes of the general public.

Our study shows that contextual parameters have a significant impact on people’s attitudes towards VC usage when it comes to deciding whether they view such usage as acceptable or not. Our work also provides further insight into the diverse attitudes towards COVID-19 vaccination in the form of a cluster analysis. This evidence-based research can provide policymakers with timely insight that balances public health and privacy concerns amid an unprecedented pandemic.

This study was published at FAccT 2022 [257].

4.2 Study Methodology

Our study explores the privacy and societal implications of information flows resulting from the use of vaccination certificates (VCs) in enforcing vaccination mandates (VMs). We survey a demographically-stratified US sample on Prolific [182] to investigate how various VC information sharing practices affect people’s perceptions of norms.

4.2.1 CI-Based Vignette Survey

We use a CI-based vignette survey method [16, 213] to gauge the effects of contextual factors on the perceived appropriateness of information sharing practices associated with common VC usage scenarios. We generated vignettes using the five CI parameters (see Table 4.1 and Figure 4.2), based on a review of existing VC proposals [44, 78, 169, 171] and related news articles [44, 48, 61, 87, 106, 118, 121, 145, 166, 171, 197]. Our study included vignettes describing two types of VC information sharing practices: (1) “first-hand” VC information sharing, where the sender shares their own VC information, and (2) VC information re-sharing, where the sender shares someone else’s VC information. These hypothetical vignettes reflect a wide range of real-world scenarios regarding the use of VCs.

First-hand information sharing vignettes

Using the following template with the CI parameters in Table 4.1, we generated 21 vignettes describing “first-hand” information sharing when people present their VCs, as *de facto* passports, to gain access or use services potentially on a regular basis:

[Recipient] ask [Sender] to show their (Subject) vaccination certificates (Attribute) to [Transmission Principle]. Would such a practice be acceptable?

To avoid potential respondent fatigue [82, 127] and limit survey completion time, we presented each participant with three randomly selected vignettes out of the 21. In addition, we curated another nine “first-hand” *VC mandate vignettes* pertaining to in-person education, employment, international travel, and apartment rental. These nine vignettes, shown at the bottom half of Table 4.1, are based on relevant and/or debated contexts where people comply with a VM by sharing their VCs [24, 34, 54, 57, 71, 112, 166, 170]. We showed these nine VC mandate vignettes to all participants in randomized order with an attention check.

VC information re-sharing vignettes

To analyze the perceptions towards possible VC information re-sharing outside the context of the original collection, we used the following question template:

Would it be acceptable for [Sender] to share [Subject] [Attribute] with [Recipient] for [Transmission Principle]?

For the sender values in the above question template, we used the recipient values from the first-hand VC information sharing vignettes, listed in Table 4.1, alongside additional CI parameter values in Figure 4.2. Figure 4.1 shows an example of the two types of vignette questions presented to participants.

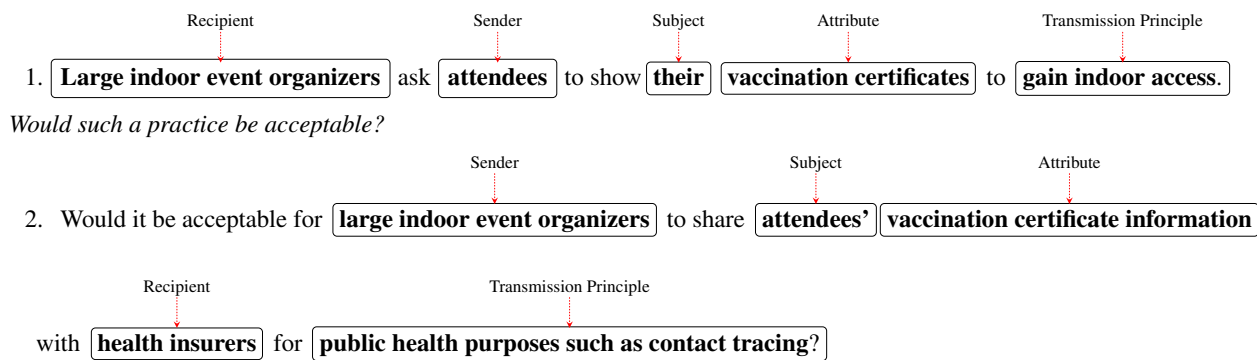


Figure 4.1: Example of first-hand sharing (top) and re-sharing (bottom) of VC information vignette questions with marked CI parameters. Note that, as per CI theory, in the re-sharing template, the sender value does not match the subject, indicating that the sender is not sharing their own information.

Clustering on Acceptance Levels

Since every participant responded to all nine occasional vignettes, we used their acceptance levels to form a 9-dimensional feature vector representing each participant. We applied agglomerative

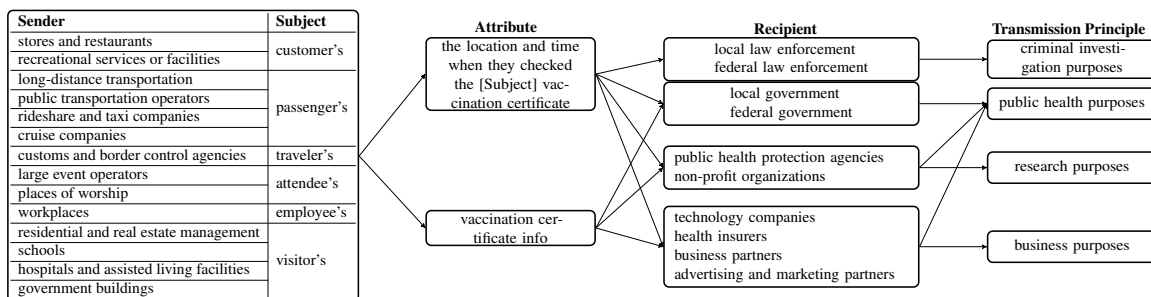


Figure 4.2: CI parameters used for vignettes involving re-sharing VC information

VC Passport Vignettes		
Sender	Recipient	Transmission Principle
customers	restaurants and cafes stores, malls, and supermarkets gyms entertainment establishments (e.g., movie theatres, museums, theatre halls) personal care businesses (e.g., nail salons, barber shops) hotels and short-term rentals (e.g., airbnb)	gain indoor access
visitors	government facilities (e.g., DMVs, courthouses) assisted living facilities hospitals places of worship apartment building management schools (K-12 and higher education)	
employees	workplaces	
attendees	large indoor event organizers large outdoor event organizers	gain access
passengers	public transportation operators	board
	long-distance bus or train companies (e.g., Megabus and Amtrak) cruise companies airline companies	
	taxi drivers, or ridesharing drivers (e.g., Uber drivers) ridesharing companies (e.g., Uber, Lyft)	use the service
VC Mandate Vignettes		
passenger	airlines	take an international flight
foreign travelers US nationals	customs and border controls	enter the United States enter a foreign country
students teachers	schools (K-12 and higher education)	return to in-person learning
job applicants	employers	be considered for a job be considered for a job in hospitals be considered for a job in assisted living facilities
potential renters	building management	rent an apartment

Table 4.1: CI parameters used for all vignettes involving first-hand VC information sharing

clustering with ward linkage on the feature vectors to identify groups of like-minded participants. We selected the number of clusters based on analyzing the dendrogram [231]. This analysis enables us to obtain a snapshot of the different types (or privacy profiles) of our survey participants on a high level.

Free-text Questions

We asked participants additional questions about their attitudes related to COVID-19 and their vaccination status, given the divided public opinion on COVID-19 vaccines and VCs in the US [80]. To contextualize participants' responses to the vignettes, we included optional free-text questions to allow participants to explain their choices.

4.2.2 Survey Deployment

We administered our survey using Qualtrics [183] and ran two pilot surveys with 75 participants each in June 2021 on the Prolific platform [182]. We chose Prolific because prior findings show that their participants provide high-quality data and are relatively diverse [174]. We used the results

from the pilots only to improve the survey questions.

For our study, we used Prolific’s “representative sample” option to recruit a demographically-stratified sample of 1,000 participants based on the age, gender, and ethnicity of the 2015 US Census data [234]. The data collection took approximately four days to complete in July 2021, and the median time spent on the survey was 13 minutes. Out of the 1,006 respondents recruited, we rejected six low-quality submissions and compensated the remaining participants \$2.00 for completing the survey. To further ensure data quality, we excluded results from the 110 respondents who failed one of the attention questions. In total, we analyzed valid responses from 890 participants. Their reported demographics is shown in Table 4.2. The survey study protocol was approved by the Institutional Review Board at Carnegie Mellon.

Gender		Age		Ethnicity	
Female	51.0%	18–27	19.6%	Asian	6.6%
Male	47.4%	28–37	18.8%	African American	12.5%
Other	1.2%	38–47	16.6%	Caucasian	71.6%
Decline to answer	0.3%	48–57	16.3%	Hispanic	4.7%
		58+	28.8%	Other	3.6%
				Decline to answer	1.0%

Table 4.2: Demographics of our study participants $N = 890$

Timing of the Survey

We conducted our survey in July 2021. At the time of the study, vaccines were widely available to all adults aged over 16, and 48.3% of the US population was fully vaccinated (55.9% had received at least one dose) [199]. By early July 2021, the relaxed COVID-19 measures and the Delta variant had led to a resurgence of positive cases and hospitalizations. At the time of the survey, states across the US had adopted or were about to adopt widely diverging policies regarding VCs. States such as California, New York, Louisiana, and Hawaii started to use digital vaccination records [168, 171], whereas states like Florida and Georgia had passed a state-wide ban on digital vaccination records [56, 167]. The debate over the use of VCs or similar vaccination verification systems remains a timely and controversial topic in public discourse [252]. This study should be viewed within this particular context.

4.2.3 Data Analysis

In our study, we measured people’s acceptance levels towards CI-based VC usage scenarios using the 5-point Likert scale and performed a qualitative analysis of the free texts about respondents’ attitudes related to COVID and VCs.

Acceptance Levels for VC Information Sharing Practices

We first compiled and graphed participants' acceptance levels towards various VC usage scenarios, which provided an overall picture of the survey responses. Then, we ran Wilcoxon and Mann-Whitney U tests, which do not assume normal distributions, to compare ordinal distributions means of first-hand sharing and re-sharing vignettes.

Regression Analysis on Acceptance Levels

We constructed a regression model of the five essential CI parameters (i.e., sender, attribute, subject, recipient, transmission principle) to measure their effects on the perceived acceptance of the respective information flow these parameters define. For the re-sharing vignettes, we set up a cumulative link mixed model [38] (CLMM), treating perceived acceptance levels as ordinal dependent variables. The CI parameters are independent variables, and every participant is treated as a random effect. The model was fitted with the adaptive Gauss-Hermite quadrature approximation with five quadrature points. The resulting model is well defined with a condition number of the Hessian less than 10^4 [38]. The re-sharing vignettes are more suitable for a regression analysis than first-hand vignettes because the CI parameter values in the re-sharing vignettes are relatively independent of each other.

Analysis of the Free-text Responses

For the qualitative analysis of free-text responses, we conducted a streamlined thematic analysis [28] of 6,230 responses. The first author open coded all free-text responses and discussed the coded data with two other authors. We discuss the resulting themes in Section 4.3.4.

4.2.4 Limitations

Our study has several limitations. First, similar to previous efforts in CI-based surveys [16, 213, 254], our study is limited to the information flow space defined by the CI parameter values. As we discussed in Section 4.2.1, we purposefully elicited the CI parameter values from relevant news on COVID and VC deployments. These values are not comprehensive and might change as the real-world situation evolves. Future work can examine these changes. Second, our results may not be generalizable to the US population, as crowd workers recruited from Prolific may differ from the general public. We tried to mitigate this issue by recruiting a large demographically-stratified sample based on the US census data. Our sample has a vaccination rate of 75% compared with the national rate of 56% at the time of the survey [199], which might induce bias in our results. Also, we only surveyed US participants, which means the results may not apply to other nations, as information norms may vary across cultures. Finally, as with all survey work, we rely on participants' self-reported data, which may be prone to biases such as social desirability bias.

4.3 Results

This section details our analysis of vignettes as discussed in Section 4.2.1 where individuals are asked to share their VC information with a range of entities for various purposes and under different constraints. As discussed in Section 4.2.1, we also examine vignettes that describe the possible re-sharing of one’s VC information by the receiving entity beyond the context of the original data collection. By varying the different contextual parameters across vignettes (see Section 4.2.1), we can better understand the privacy expectations and converging norms regarding VC information sharing around the following research questions:

- **In what contexts are VC deployments and mandates perceived appropriate?** In Section 4.3.1 and Section 4.3.2, we report and compare the levels of acceptance towards VC deployment and mandate under different contexts.
- **How does the practice of re-sharing VC information affect the perceived appropriateness?** In Section 4.3.3, we compare the levels of acceptance of first-hand VC information sharing (when the sender is also the subject of the information) to the re-sharing of VC information (when the sender shares someone else’s information).




Figure 4.3: Reported acceptance levels for VC passport vignettes organized by recipients. The box plots on the right indicate the variances of the acceptability scores. Recall that the survey only showed each participant three randomly selected vignettes. The denominator of the percentages is the number of responses for each vignette. The top row shows an overall acceptance level across all vignettes.

4.3.1 VCs as *de facto* passports

The 21 first-hand VC information sharing vignettes reflected the scenarios in which people show their VCs, as *de facto* passports, to gain access to a service, venue, or facility. Figure 4.3 summarizes the acceptance levels of providing VC information to 21 different CI recipients in this particular context. A majority of respondents viewed “VC as passport” scenarios as acceptable or somewhat acceptable. For scenarios involving gaining access to assisted living facilities, cruises, and airlines, respondents expressed on average high levels of acceptance, where over 75% of participants considered those at least somewhat acceptable. The least acceptable scenarios involve asking visitors to show their VCs to enter apartment buildings or visit worship places. Fewer than 50% of responses indicate requiring VCs in apartment buildings as acceptable (27%) or somewhat acceptable (21%). Worship places elicit a similar reaction with only 34% and 16% of responses suggesting requiring VCs is “acceptable” and “somewhat acceptable”. Furthermore, asking to show VCs in public transportation, government buildings, shops, worship places, and apartments elicited more diverse reactions.

Variances in perceptions We analyze variances in perceptions across 21 vignettes, which is an indicator of norm formation. A low variance is a sign of a relative agreement within the scenario. Figure 4.3 shows the variances of appropriateness scores among the scenarios. We observe low variances in perceptions for scenarios in assisted living facilities, cruises, airlines, and indoor events. This is in contrast to the high overall variances in perceptions associated with hospitals, workplaces, shops, worship places, and apartments.

Essential services and basic facilities

Our results showed that asking for VCs in a non-essential facility is considered significantly more appropriate than asking for VCs in an essential facility. For example, 68% of responses indicate it is “acceptable” (47%) or “somewhat acceptable” (21%) to show VCs in eateries compared with the lower 57% (37% “acceptable” and 20% “somewhat acceptable”) for showing VCs in stores. Several accompanying free-text comments potentially explained the discrepancy. P213 commented on restaurants requiring VCs: *“The spaces are just too small, and the ambient air is not efficiently exchanged. This is the number one place for requiring people to be vaccinated. People are voluntarily choosing to go, so should have to show a pass.”* Yet, P156 noted: *“Freedom to access a source of food such as a supermarket should be effortless. Having to show vaccination certificates to enter would cause mayhem.”*

Noticeably, asking visitors to show their VCs in hospitals is significantly less acceptable than that in assisted living facilities (Mann–Whitney $U = 7924.5$, $n_1 = 117$, $n_2 = 116$, $p < 0.01$, $Cohen'd = 0.293$, two-tailed), although both are places with COVID-19 vulnerable populations. P72 provided a possible line of reasoning: *“Even though some people may not feel comfortable getting the vaccine, they should still be granted access to hospital resources indoors. If an individual lacks certification, they should be wearing a mask.”*

The results also reveal a similar contrast between public transportation and other forms of transportation such as airlines, trains, and taxis or ride-sharing services. Only about 60% of respondents found it “acceptable”(39%) or “somewhat acceptable”(22%) to show their VCs to use

public transportation. P209’s open-ended comment provides some context to the reported contrast: “Safety is important here too, but unlike flying, public transportation is more of a necessity and shouldn’t be hindered by this.”

In summary, the results highlight the relationship between the nature of the context—whether it is deemed essential or non-essential—and the perceived appropriateness. This suggests a need for nuanced policy making with regard to using VCs as passports.

4.3.2 Examining VC mandate vignettes

The nine VC mandate vignettes reflected publicly debated scenarios in the context for which governments around the world are seeking mandates to require VCs, such as for international travel, returning to in-person learning, applying for a job, and renting an apartment, as mentioned in Section 4.2.1. Figure 4.4 summarizes the acceptance levels for each of the nine vignettes from all participants. Overall, 74% of participants found the selected vignettes to be “acceptable” (58%) or “somewhat acceptable” (16%).

A VC mandate for international travel is perceived appropriate to take a flight or use at the border. Our results show that requesting VCs for international travel is largely perceived as appropriate: 82% of all participants stated that it is acceptable (68%) or somewhat acceptable (14%) for passengers to show VCs to take an international flight. Similarly, 85% of respondents perceived showing VCs to customs and border control agencies as acceptable (70%) or somewhat acceptable (15%), both for entering the US or a foreign country.




Figure 4.4: Participants’ acceptance levels for nine vignettes. The top row displays the averaged response across nine vignettes. The right graph shows a box plot of the ordinal data with the mean marked in orange.

A VC mandate for employment: Perceived appropriate to apply for a job at assisted living facilities or hospitals. 81% of respondents expressed similar levels of acceptability for sharing

vaccination certificate information with employers to be considered for a job in assisted living facilities and hospitals, with 14% stating it was somewhat acceptable and 67% viewing it as acceptable. In comparison, when it comes to applying for a general position, only 60% participants considered showing vaccination certificates to potential employers for a job as acceptable (37%) or somewhat acceptable (23%). A Wilcoxon Signed-Rank test shows that levels of acceptance for the general case were statistically significantly lower than the levels of acceptance for the cases involving hospitals and assisted living facilities ($Z = 5.42, p < 10^{-26}$).

A VC mandate for education: Perceived appropriate for teachers, less so for students When asked whether it is appropriate to share VC information with schools for returning to in-person learning, the acceptance levels depended on whether the sender is the students or teachers. These two vignettes involved the same CI parameters except for the sender. Using a Wilcoxon Signed-Rank test ($Z = 9.89, p < 0.000001$), we noted the perceived levels of acceptability for students were statistically significantly lower than those for teachers. This means that even though the majority of our survey respondents considered the VC mandate in schools acceptable, they regarded asking students to share their VC information with the school as less acceptable than asking teachers to do so.

A VC mandate in residential settings: Perceived as inappropriate overall Respondents viewed showing VCs to building management to rent an apartment as the least acceptable. With only 17% stating that it was somewhat acceptable, a slightly higher percentage of the respondents (26%) saw it as acceptable. Such low acceptance is also consistent Section 4.3.1 where respondents considered showing VCs to visit an apartment as the least acceptable.

4.3.3 Examining Scenarios on Re-sharing VC Information

We examined respondents' perceived appropriateness regarding VC information re-sharing practices: a situation in which a VC shown in a given context is being shared by the original recipient with a different entity for a new purpose or under a new condition. For example, when businesses share their customers' VC information with the health protection agency for public health purposes such as contact tracing. For a full list of vignettes and CI parameters, see Figure 4.2.

Figure 4.5 shows a heat map of average acceptance levels of vignettes describing VC information re-sharing. Overall, the practice of re-sharing and re-purposing of VC information is perceived as less appropriate compared with the first-hand VC information exchange in the original context. We found a statistically significant difference between the two types of information flows using a Wilcoxon matched-pair signed rank test ($Z = 4.80, p < 10^{-7}$).

Regression analysis of vignettes' CI parameters

A closer examination of CI parameters in the re-sharing vignettes reveals varied levels of perceived appropriateness. Table 4.3 shows the results of the CLMM regression analysis (see Section 4.2.3) of factors affecting participants' perceived acceptance levels of re-sharing VC information. We

Factors	Est.	Std. Err	Z	p-value
Sender: baseline=customs and border control				
government buildings	0.0448	0.1042	0.4299	0.6672
hospitals and assisted living facilities	0.1404	0.1057	1.3280	0.1842
long-distance transportation	0.1713	0.1022	1.6758	0.0938
cruise companies	0.1921	0.1052	1.8256	0.0679
workplaces	0.2633	0.1079	2.4404	0.0147*
large event organizers	0.3089	0.1080	2.8603	0.0042**
schools	0.3545	0.1072	3.3076	0.0009***
stores and restaurants	0.3804	0.1053	3.6127	0.0003***
recreational services or facilities	0.4255	0.1062	4.0054	6.2e-5***
public transportation operators	0.4280	0.1081	3.9588	7.5e-5***
places of worship	0.4600	0.1090	4.2195	2.4e-5***
residential and real estate management	0.5122	0.1090	4.7005	2.6e-6***
rideshare and taxi companies	0.6041	0.1046	5.7761	7.7e-9***
Recipient: baseline=public health protection agencies				
local government	0.9994	0.0839	11.9142	2.2e-16***
federal government	0.9788	0.0838	11.6736	2.2e-16***
non profit organization	2.2955	0.0813	28.2482	2.2e-16***
health insurer	2.5199	0.0821	30.7003	2.2e-16***
business partners	3.9754	0.0872	45.5826	2.2e-16***
technology company	4.2075	0.0886	47.4724	2.2e-16***
advertising and marketing partners	4.7668	0.0926	51.4647	2.2e-16***
Attribute: baseline=vaccination certificate information				
location and time	-0.0472	0.0415	-1.1355	0.2561
Transmission Principle: baseline=public health purposes				
criminal investigation	0.7099	0.0819	8.6722	2.2e-16***
research	0.6427	0.0927	6.9334	4.1e-12***
business	0.2994	0.0475	6.3036	2.9e-10***

Table 4.3: Cumulative Linear Mixed Model Regression. A positive coefficient (estimate) shows participants' decreased acceptance

found that values of three CI parameters—sender, recipient, and transmission principle—have a statistically significant effect on participants’ perceived appropriateness.

Sender We used “customs and border control agencies” as the baseline in our regression analysis of the sender parameter as such a sender is the most accepted among all senders. Out of all 14 senders of the vignettes, rideshare drivers/companies and residential management were perceived as the most unacceptable sender, with respective odds ratios of 1.6 ($=e^{0.6041}$) and 1.4 ($=e^{0.5122}$) compared to the baseline value. In other words, VC information sharing by rideshare companies is 1.6 times less acceptable than customs and border controls, holding constant all other variables.

Recipient Our results show that sharing VC information with public health protection agencies (which we used as the baseline) is significantly more acceptable than sharing VC information with other receiving entities. The least unacceptable recipients included advertising and marketing partners, followed by technology companies, and business partners, as indicated by the decreasing coefficients.

Transmission Principle Our results indicate that sharing VC-related information for public health purposes (the baseline) is significantly more acceptable than for other purposes or conditions.

Attribute In addition to the information in the VC itself, we looked at the meta-data associated with VC information sharing, such as the location and the timestamp. This information, when shared with other entities, could be further used to surveil or track individuals. Our analysis, however, shows no statistically significant difference in perceptions of re-sharing the VC information or the residual meta-data (location and time) associated with the VC check.

Summary Our analysis shows that contextual factors captured by the CI framework affect the degree to which participants judged a VC practice acceptable. Some combinations of sender, subject, recipient of the VC information and the condition/constraint of the transfer (transmission principle) have a statistically significant effect on the perceived acceptance of the information flow that these parameters define. This aligns with prior work that leverages CI to evaluate privacy violations in other contexts [16, 213]. Notably, the subject parameter of the information flow is particularly important, as it distinguishes re-sharing practices. Our participants found VC-related information re-sharing practices less acceptable than their providing VC directly to recipients.

4.3.4 Different Views on VCs: Qualitative Analysis

The open-ended comments accompanying the vignettes provide insight into the motivating factors behind the stated perceptions of appropriateness. Our thematic analysis of the free texts reveals three main attitudes.

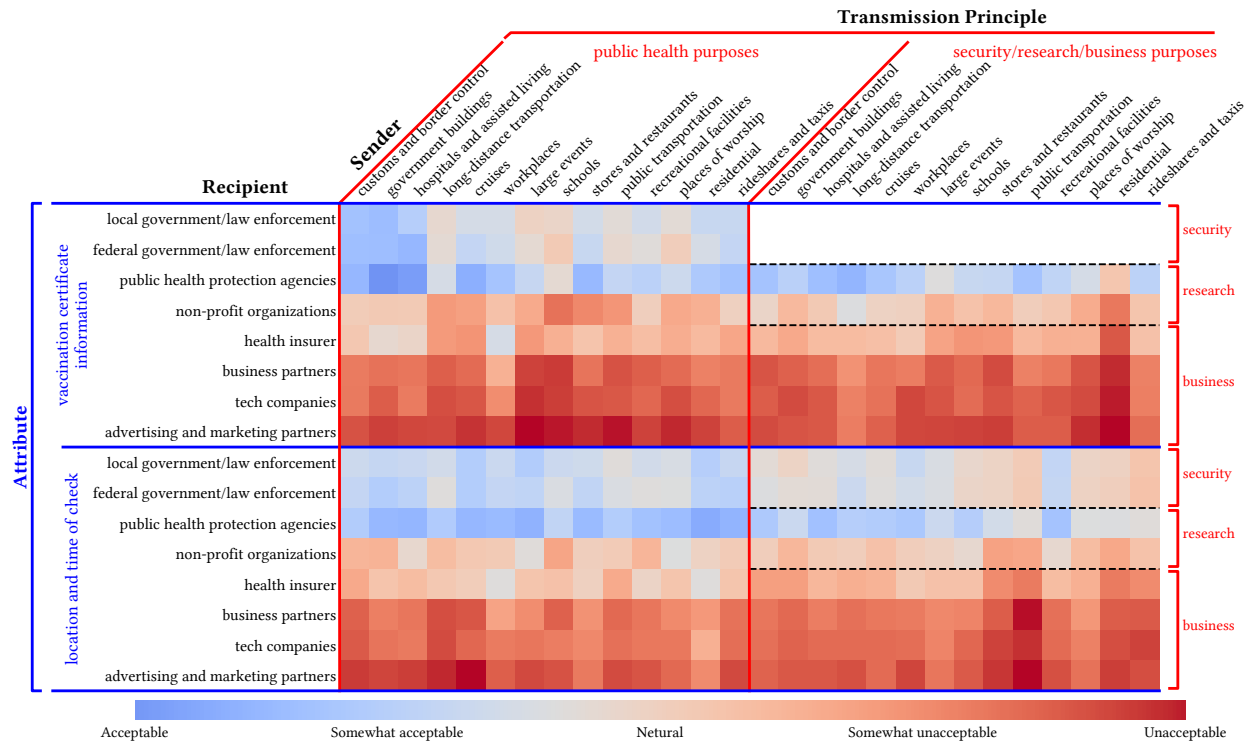


Figure 4.5: A heat map of the average of all participants’ responses under a combination of four CI parameters (sender, recipient, attribute, and transmission principle). For instance, the color of the top left cell represents the acceptance level of the information flow—customs and border control agencies share their customers’ vaccination certificate information with the local government for public health purposes.

In favor of VCs

Over half of participants (57%) noted that VCs would make them feel safer or curb the spread of the virus. Some (12.5%) also mentioned VCs can show proof and prevent counterfeit CDC cards, or carrying digital VCs are easy and safe from losing them. 7.6% of participants referred to communitarian ethics in helping protect others and their community, while others (3.8%) saw no difference from existing practices like showing IDs to buy alcohol. 2.5% of participants commented that VCs would serve as an incentive to get more people vaccinated.

Opposing VCs

Some (11.7%) participants indicated they regard VCs as an invasion of their privacy or, more generally, a restriction on their personal freedom. Others (7.3%) believed that deciding on whether to receive the vaccine should be left to their individual discretion instead of being imposed by the government or some other organization. 6.6% of participants considered vaccination status as private medical information and thus information that should not be shared with anyone other than their doctors. 4.3% claimed that VCs are illegal and/or unconstitutional and/or violated their HIPAA rights. Concerned about potential harms, some (6.2%) perceived VCs as a form of government overreach, compared the practice to identity control measures such as the ones under the Nazi regime (“Papers, please”).

A few participants (4.7%) referred to information privacy, indicating that they were not comfortable with some information included in VCs or would not want such data retained or shared. Some (4.0%) noted that employing VCs would result in discrimination against the unvaccinated.

Context-sensitive views

37% of participants expressed mixed reactions and considerations dependent on the contexts of VC information sharing. For example, 11.2% thought that private businesses are free to require VCs at their own discretion, and 7.6% were against requiring VCs at places to which people need access, such as public transportation and stores. 5.7% of participants believed that other methods such as mask wearing, negative COVID tests, and occupancy limits should also be accepted if some would not want to present their VCs. 4.6% of participants mentioned the need to accommodate people who may not be able to receive vaccines when deploying VCs. 3.9% of participants thought that VCs are particularly controversial and could elicit strong objections and potentially violent behaviors.

4.4 Clustering Analysis

Our study also presents a snapshot of a spectrum of privacy expectations with regard to vaccination certificate usages. To identify potential patterns in emerging privacy expectations towards VCs, we cluster responses by using their acceptance levels as feature vector representing each participant. We applied agglomerative clustering with ward linkage on the feature vectors to identify groups of like-minded participants. We selected the number of clusters based on analyzing the dendrogram [231]. Figure 4.6 illustrates the resulting seven clusters (i.e., privacy profiles) of each cluster’s acceptance

	Cluster#1: size=47	Cluster#2: size=47	Cluster#3: size=78	Cluster#4: size=65	Cluster#5: size=173	Cluster#6: size=110	Cluster#7: size=367
teachers return to in-person learning	-2.00 0.00	-1.87 0.34	-0.69 1.02	-0.08 1.24	1.83 0.39	1.51 0.55	1.93 0.27
students return to in-person learning	-2.00 0.00	-1.74 0.77	-1.04 1.01	-0.88 0.98	1.58 0.60	0.52 1.30	1.84 0.44
apply for a job	-2.00 0.00	-2.00 0.00	-1.04 1.07	-1.09 0.95	0.40 1.26	0.91 0.87	1.63 0.66
apply for a job in hospitals	-1.89 0.31	-1.81 0.45	0.14 0.99	1.05 1.22	1.88 0.41	1.40 0.73	1.95 0.23
apply for a job in assisted living	-1.89 0.31	-1.77 0.43	0.13 0.97	1.55 0.64	1.87 0.37	1.23 0.93	1.96 0.19
US nations enter foreign country	-2.00 0.00	0.38 1.34	-0.01 1.12	1.60 0.58	1.86 0.43	1.50 0.52	1.98 0.15
foreign travelers enter US	-2.00 0.00	0.64 1.22	0.13 1.22	1.51 0.79	1.87 0.37	1.23 0.80	1.99 0.10
take an international flight	-2.00 0.00	-0.81 1.53	-0.62 1.06	1.52 0.59	1.77 0.62	1.41 0.65	1.99 0.14
rent an apartment	-2.00 0.00	-2.00 0.00	-1.24 0.98	-1.58 0.56	-1.12 0.76	0.23 1.11	1.48 0.70

Figure 4.6: Profiles associated with a 7-cluster model. Each cluster profile contains 2 columns: the left one displays the mean acceptance level (acceptable=2, unacceptable=-2). The right column shows the standard deviations, ranging from 0 to 2.

levels. Figure 4.6 illustrates the resulting seven clusters: Cluster#1 ($N = 47$) and Cluster#7 ($N = 367$) represent polar extremes of acceptance (Cluster#7) and non-acceptance (Cluster#1) to all nine vignettes. The low standard deviations in these two clusters point to the homogeneity among cluster members. Cluster#5 ($N = 173$) represents acceptance to the majority of vignettes at least somewhat acceptable, except the scenario of using VCs to rent an apartment which is somewhat unacceptable and neutral support for showing VCs to apply for jobs. Cluster#2 ($N = 47$) resembles Cluster#1's perceptions in all but three scenarios—presenting VCs at border and customs to enter the US or a foreign country and taking an international flight. Finally, Cluster#3 and #4 are a mixed bag. Cluster#3 ($N = 78$) represents nearly neutral responses ($mean = -0.01$ — 0.14) towards scenarios related to presenting VC at border, or to apply for a position in hospitals and assisted living facilities. Cluster#4 ($N = 65$) on the one side is against using VC for in-person learning, applying for a general position, and renting an apartment; on the other side, it shows support for vignettes related to assisted living and hospitals, as well as international air travel and at the border.

4.5 Discussion

As we write this paper, VMs and VCs remain a highly contentious and politically polarizing subject. Faced with the new and highly infectious omicron variant, many governments around the world have introduced vaccination mandates or the use of vaccination certificates across a number of different contexts [157, 180, 207]. The intensity and polarization of the debate is vividly reflected in the views expressed by the participants in our study. At the one extreme, a handful of participants left profanities in the free-text responses, aimed at the authors whom they mistakenly thought were conducting research to shore up support for VC mandates and deployments. At the other extreme, a few participants left equally strong responses about people's collective responsibility to protect one another, asserting that those who refuse vaccinations are selfishly neglecting their responsibility.

Aside from the extremes, at an aggregate level, the percentage of people who find appropriate many of the VC sharing scenarios presented to them, could be taken as potential support for a fairly broad VC mandate. A closer look, however, reveals a more nuanced picture in which contextual factors had significantly affected participants' attitudes. It mattered whether the VC information is shared with the school to facilitate in-person classroom, with a grocery store owner or with a gym operator as a condition of admittance, or with a customs agent to enter a country. The recipient

with whom VC data is shared, the purpose(s) for sharing, as well as guarantees (or lack thereof) about the processing of VC information all have a significant effect on people’s acceptance of VC deployments. It is worth noting, too, that our study found the **subject** parameter of the information flow to be important, lending credence to our initial question about first-hand use versus re-sharing practices. When the values for all the parameters are clearly stated, our results indicate a negative sentiment towards requiring VCs for access to essential services and activities, places of worship, and apartment buildings. Further, perhaps not surprisingly, the practice of re-sharing VC information is perceived as largely inappropriate. These empirical results illustrate the importance of organizing a survey like this one by systematically sampling different contextual values, especially when it comes to understanding people’s acceptance of information flows associated with different possible VC deployments and their implications.

Finally, as posited by the CI theory [164], newly-formed information flows that challenge established norms can affect the ultimate realization of a range of societal values such as equality, equity, and civil liberties. The assessment of the appropriateness of new flows includes: 1) a cost and benefit analysis of the information flow related to all the affected parties: Who benefits? What risks are involved? 2) a review of moral and ethical values such as fairness, autonomy, and informational harm; 3) considerations around how the new information flow contributes to fulfilling the “context-specific values, ends and purposes” [164].

The qualitative analysis of the open-ended responses in Section 4.3.4 reveals that the ethical and societal values indeed are part of the normative assessment of the perceived appropriateness of VCs. The open-ended comments included different aspects related to the appropriateness assessment. We observed the weighing of public health interests against the expectations of freedom and privacy in various contexts. Many participants reported viewing enhanced public health as a societal benefit, while some were concerned about potential harm brought by heightened government surveillance. Some participants also expressed concerns about their bodily autonomy, the violation of personal freedom, and the intrusion of privacy on their health information, while others also warned of potential discrimination against the unvaccinated and restrictions on their rights to access essential facilities such as stores and hospitals.

4.6 Results from an Additional Survey with a Large Sample

We were able to run a slightly revised version of our COVID vaccination certificate study again with a large sample under the support of the COVID States Project [77]. The larger sample size enabled us to examine the impact of various demographic features on people’s perceived acceptance levels of vaccination certificates in different usage contexts. Conducting the study once more after 16 months provides a rare opportunity to analyze any changes in perceptions that may have occurred due to the evolving circumstances of the COVID situation.

4.6.1 The COVID States Project

This sample was collected under the COVID States Project [77] data wave #25 between October 6th, 2022 to November 9th, 2022 by PureSpectrum, an online panel management platform. Respon-

dents were recruited from all 50 states and Washington DC, with flexible state-level demographic quotas applied for gender, age, race, and ethnicity. Additional sample augmentations were used to oversample African American, Asian, and Hispanic respondents. Moreover, additional demographic information, including income level, relationship status, education level, employment status, parenthood with children under the age of 18, zip code, state, urbanicity, political affiliation, and political ideology, was provided. The survey also gathered information related to COVID, such as whether respondents had tested positive for COVID, the severity of their illness (if applicable), their vaccination status, and their future vaccination plans. The full survey text of our vaccination certificate module, along with the module questions for the demographic and COVID-related data we received, can be found in Appendix B.2.

4.6.2 Participants and Demographics

We received data from a total of 13,774 participants, but we excluded those who failed to answer our attention check question correctly. Therefore, the analysis presented in this section was conducted with the remaining 10,631 participants. Demographic information (i.e., age, gender, race, and education level) can be found in Table 4.4. Comparing with our Prolific sample collected in July 2021 (Table 4.2), our sample is skewed towards more females and older individuals.

One of the objectives of the COVID States project is to identify and assess potential variations among states. Figure 4.7 displays the number of participants recruited from each state, while Figure 4.8 illustrates the ratio of $\frac{\text{Percentage of sample population from state}}{\text{Percentage of US Census population from state}}$. A ratio closer to 1 indicates that the state population in our sample is representative of the state's actual population percentage in the US population. A ratio greater than 1 indicates that the state population in our sample is oversampled relative to its actual population percentage, while a ratio less than 1 indicates that the state population in our sample is under-sampled. For example, even though California has the most significant number of participants (634 in Figure 4.7), it is still underrepresented relative to its population percentage in the United States. The same holds for other populous states such as Texas, Florida, and New York. In contrast, some states, such as New Hampshire, Nebraska, Wyoming, and New Mexico, are oversampled. This suggests that while the sample contains respondents across the United States, it is not a precise representation of the US population in terms of states.

4.6.3 Acceptance Levels toward Various Usage Scenarios

Figure 4.9 shows the acceptance levels that our participants expressed towards each of the 22 usage scenario on vaccination certificates. When compared to the results obtained in July 2021 (Figure 4.4 and Figure 4.3), there has been a significant drop in the acceptance of vaccination certificates, with the percentage of respondents who found it somewhat acceptable or acceptable decreasing from 70% to 49%. On the other hand, the percentage of those who found it somewhat unacceptable or unacceptable increased from 24% to 36%. Although there has been a general shift towards unacceptability, the relative order of acceptance levels across different usage scenarios has remained almost the same. For example, scenarios involving showing VCs to enter a country or access assisted living facilities, board cruise ships or airplanes are still among the most acceptable scenarios. Conversely, scenarios involving showing VCs to rent an apartment or visit apartment

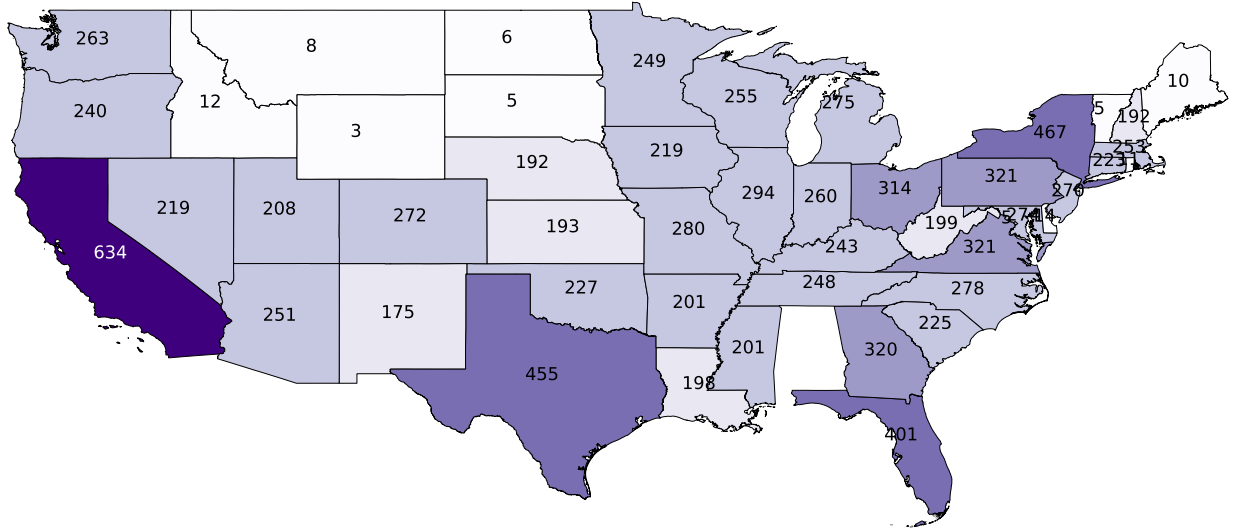


Figure 4.7: U.S. map labeled with the number of participants recruited in each state

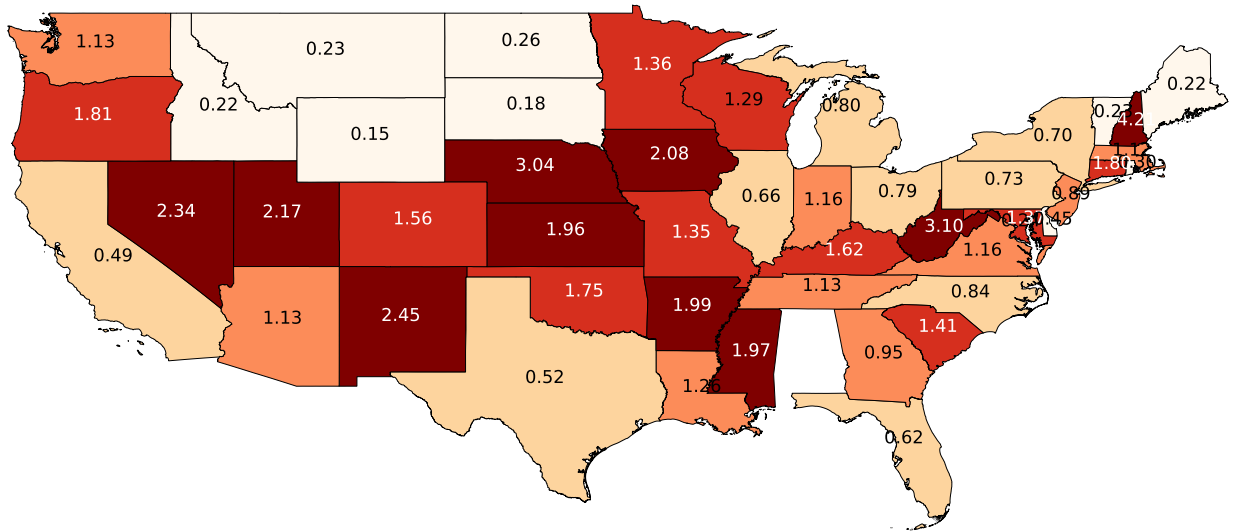


Figure 4.8: Darker red shading indicates oversampling and lighter yellow and white shading indicates the number of participants from this state is under-sampled.

Gender		Age		Race		Education	
Female	61.0%	18–27	10.6%	Asian	5.6%	Some High School or Less	2.6%
Male	37.5%	28–37	13.4%	African American	11.5%	High School Graduate	22.1%
Other	1.4%	38–47	12.7%	Caucasian	72.0%	Some College	26.0%
Declined	0.1%	48–57	13.7%	Hispanic	8.1%	College Degree	36.7%
		58+	49.5%	Other	2.8%	Graduate Degree	12.7%

Table 4.4: Demographics of our study participants $N = 10,631$

buildings, places of worship, stores, or apply for a job are still among the least acceptable scenarios. This suggests that our results are largely consistent with our prior study, but overall acceptance has decreased, possibly due to the decrease of disease severity and the winding down of the COVID pandemic. This highlights the importance of conducting periodic assessments of public attitudes towards technologies such as VCs, which can be influenced by external factors such as changes in public health circumstances.

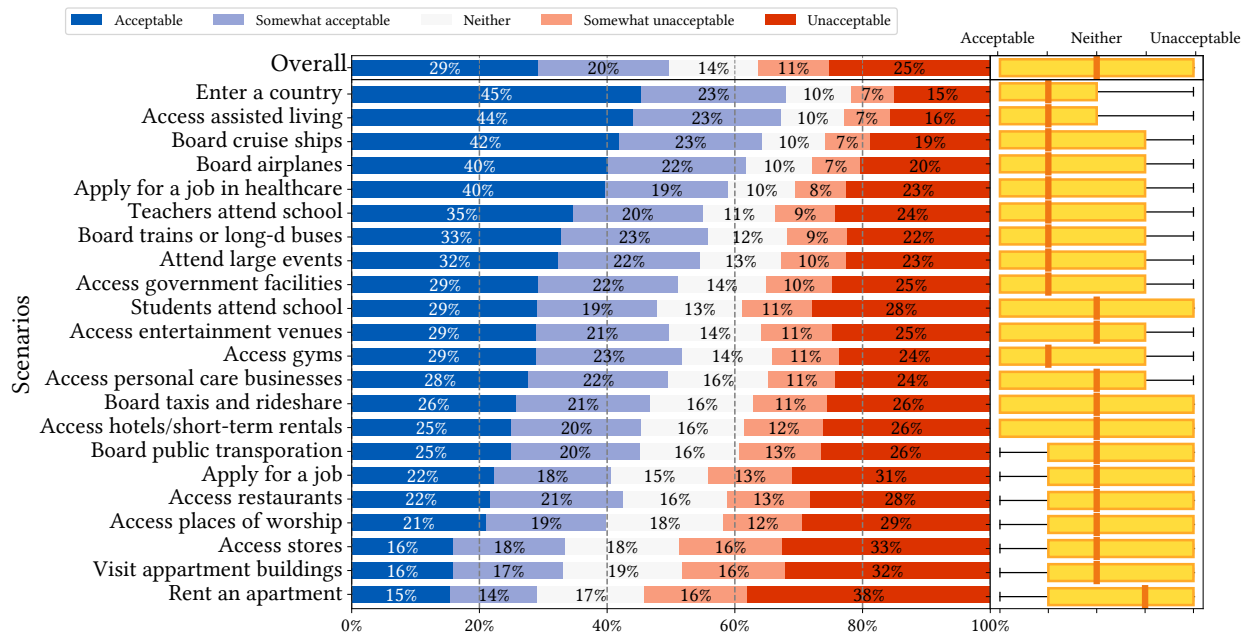


Figure 4.9: Reported acceptance levels for VC vignettes sorted by the most number of acceptable response percentage. The box plots on the right indicate the variances of the acceptability scores. The top row shows an overall acceptance level across all vignettes.




Figure 4.10: Dendrogram of agglomerative clustering with a cut-off at 9 clusters

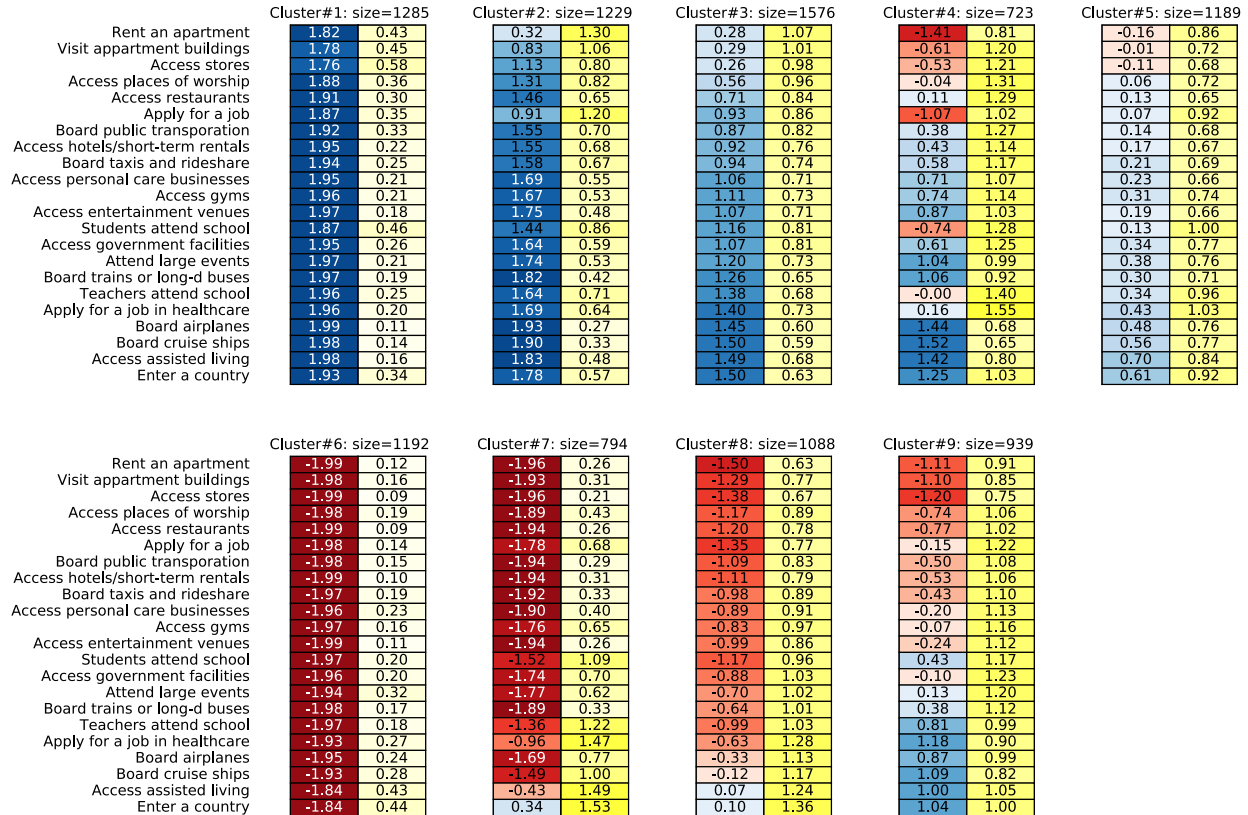


Figure 4.11: Profiles associated with a 9-cluster model. Each cluster profile contains 2 columns: the left one displays the mean acceptance level (acceptable=2, unacceptable=-2). The right column shows the standard deviations, ranging from 0 to 2. The rows are ordered based on Figure 4.9.

4.6.4 Clustering

Cluster selection

Our study also presents a snapshot of a spectrum of privacy expectations with regard to vaccination certificate usages. To identify potential patterns in emerging privacy expectations towards VCs, we cluster responses by using their acceptance levels as feature vector representing each participant. We applied agglomerative clustering with ward linkage on the feature vectors to identify groups of like-minded participants. We selected the number of clusters based on analyzing the dendrogram (Figure 4.10).

Cluster Profiles

Figure 4.11 illustrates the resulting nine clusters (i.e., privacy profiles) of each cluster's acceptance levels. Cluster#1 ($N = 1,285$, 12.8%) and Cluster#6 ($N = 1,192$, 11.9%) represent polar extremes of acceptance (Cluster#1) and non-acceptance (Cluster#7) to all 22 vignettes. The low standard deviations in these two clusters point to the homogeneity among cluster members. Notably, the

current study's unacceptable cluster comprises about 11.2% of the sample size, which is a notable increase from prior study where the unacceptable cluster accounted for only 5.3%. Cluster#2 ($N = 1,229$, 12.3%) represents acceptance to the vast majority of vignettes being at least somewhat acceptable, except for the scenario of showing VCs to apply for jobs and to visit apartment buildings, which is considered somewhat unacceptable and neutral support for renting an apartment. Cluster#7 ($N = 794$, 7.9%) resembles Cluster#6's perceptions in all but three scenarios—presenting VCs at border and customs to enter a country, presenting VCs to apply for jobs in healthcare, and showing VCs to access assisted living facilities. Overall, the clusters follow a similar pattern compared to our clustering results from July 2021, as shown in Figure 4.6.

4.7 Summary of Main Contributions

- We used an established CI-based vignette survey methodology to analyze data from a US-based demographically-stratified sample ($N = 890$) about how they perceive sharing VC information with various recipients in different contexts such as education, health, or public transportation, under different conditions and for various purposes.
- Our analysis revealed that perceived appropriateness is contextual and varies depending on CI's five parameters for information flows (i.e., sender, attribute, subject, recipient, transmission principle). There is also a significant difference in acceptance of first-hand information sharing compared to later re-sharing and re-purposing of originally collected information. Overall, we found that information re-sharing with entities other than public health agencies is widely viewed as unacceptable.
- Our study illustrated how Contextual Integrity (CI) provides an effective framework for approaching controversial societal practices, such as VC deployment. It suggested that the multifactorial insights that CI yields can inform richer and more nuanced responses to challenges confronting society in today's fight against COVID-19, and potentially other similar challenges going forward.
- We obtained a snapshot of the different types (or privacy profiles) of our survey participants representing a spectrum of various attitudes towards vaccination certificates.
- The comparison between our large-scale re-run of the study and the previous one highlights the importance of conducting periodic assessments of public attitudes towards technologies such as VCs, which can be influenced by external factors such as changes in public health circumstances.

This work was published at FAccT 2022 [257].

Chapter 5

Evaluating the Effectiveness of Mobile App Privacy Labels: How Usable are Today's Mobile App Privacy Labels?


5.1 Overview

Another domain that has given rise to complex privacy management challenges is smartphones. Mobile apps, which have played a key role in the broad adoption of smartphones, are known to collect a broad range of sensitive data about people and frequently use this data for purposes that are not solely limited to delivering the app's core functionality. This includes in particular the use of sensitive data such as location for advertising and marketing purposes. As everyone knows, users seldom, if ever, read the text of privacy policies, which in principle should help them learn about the data collection and use practices of the apps they are considering downloading on their phones. For a number of years, researchers in usable privacy have advocated the adoption of privacy nutrition labels—or privacy labels for short, that would provide users with clear and succinct summaries of data practices that are likely to be most relevant to them [113, 115, 116]. About two years ago, the iOS app store and a little later the Google Play Store decided to adopt such labels, requiring app developers to now provide privacy labels for their apps when they publish them in their app stores. In this and the next chapter, we investigate to what extent these recently introduced labels actually deliver on their promise of offering users more effective and practical ways of learning about the data practices of mobile apps. This includes looking at whether people are aware of the availability of these labels, whether they understand what they mean and whether they can effectively use information provided in these labels to support their privacy decisions (Chapter 5). This also includes looking at whether and to what extent labels introduced by the iOS App Store and the Android Google Play Store actually address people's most typical privacy concerns (Chapter 6).






In this chapter, we present research that suggests that iOS privacy nutrition labels currently play a limited role in helping users better manage their mobile app privacy. We identify areas of misunderstanding of and dissatisfaction with iOS privacy nutrition labels that hinder their usability and effectiveness as privacy notices. We discuss areas where app privacy labels could be improved


App Privacy [See Details](#)

The developer, **DoorDash, Inc.**, indicated that the app's privacy practices may include handling of data as described below. For more information, see the [developer's privacy policy](#).













 **Data Used to Track You**

The following data may be used to track you across apps and websites owned by other companies:

-  Purchases
-  Location
-  Contact Info
-  Identifiers
-  Usage Data

 **Data Linked to You**


The following data may be collected and linked to your identity:

-  Purchases
-  Financial Info
-  Location
-  Contact Info
-  Contacts
-  User Content
-  Search History
-  Browsing History
-  Identifiers
-  Usage Data
-  Diagnostics
-  Other Data








Privacy practices may vary, for example, based on the features you use or your age. [Learn More](#)

App Privacy [See Details](#)

The developer, **Chipotle Mexican Grill**, indicated that the app's privacy practices may include handling of data as described below. For more information, see the [developer's privacy policy](#).

 **Data Not Linked to You**

The following data may be collected but it is not linked to your identity:

-  Purchases
-  Financial Info
-  Location
-  Contact Info
-  User Content
-  Usage Data
-  Diagnostics

Privacy practices may vary, for example, based on the features you use or your age. [Learn More](#)

Figure 5.1: Screenshots of compact privacy labels from DoorDash (left) and Chipotle (right) in the iOS App Store

App Privacy

The developer, DoorDash, Inc., indicated that the app's privacy practices may include handling of data as described below. This information has not been verified by Apple. For more information, see the [developer's privacy policy](#).

To help you better understand the developer's responses, see [Privacy Definitions and Examples](#).

Privacy practices may vary, for example, based on the features you use or your age. [Learn More](#)



Data Used to Track You

The following data may be used to track you across apps and websites owned by other companies:

- Purchases**
Purchase History
- Location**
Precise Location
Coarse Location
- Contact Info**
Physical Address
Email Address
Name
- Identifiers**
User ID
Device ID
- Usage Data**
Product Interaction



Data Linked to You

The following data, which may be collected and linked to your identity, may be used for the following purposes:

Third-Party Advertising

- Purchases**
Purchase History
- Location**
Precise Location
Coarse Location
- Contact Info**
Physical Address
Email Address
- Identifiers**
User ID
Device ID
- Usage Data**
Product Interaction

Developer's Advertising or Marketing

- Purchases**
Purchase History
- Location**
Coarse Location
- Contact Info**
Physical Address
Email Address
Name
- Contacts**
Contacts
- Search History**
Search History
- Browsing History**

Browsing History

- Identifiers**
User ID
Device ID
- Usage Data**
Product Interaction
Advertising Data
- Other Data**
Other Data Types

Analytics

- Purchases**
Purchase History
- Location**
Precise Location
Coarse Location
- Contact Info**
Physical Address
Email Address
Name
Phone Number
- User Content**
Customer Support
- Search History**
Search History
- Browsing History**
Browsing History
- Identifiers**
User ID
Device ID
- Usage Data**
Product Interaction
Advertising Data
- Diagnostics**
Crash Data
Performance Data
Other Diagnostic Data

Product Personalization

- Purchases**
Purchase History
- Financial Info**
Payment Info
- Location**
Precise Location
Coarse Location
- Contact Info**
Physical Address
Email Address
Name
Phone Number
- User Content**
Emails or Text Messages
Photos or Videos
Customer Support
- Search History**
Search History
- Identifiers**
User ID
Device ID
- Usage Data**
Product Interaction
- Diagnostics**
Crash Data
Performance Data
Other Diagnostic Data
- Other Data**
Other Data Types

Other Purposes

- Contacts**
Contacts

App Privacy

The developer, Chipotle Mexican Grill, indicated that the app's privacy practices may include handling of data as described below. This information has not been verified by Apple. For more information, see the [developer's privacy policy](#).

To help you better understand the developer's responses, see [Privacy Definitions and Examples](#).

Privacy practices may vary, for example, based on the features you use or your age. [Learn More](#)



Data Not Linked to You

The following data, which may be collected but is not linked to your identity, may be used for the following purposes:

Analytics

- Usage Data**
Product Interaction

Other Usage Data

- Diagnostics**
Crash Data
Performance Data
Other Diagnostic Data

App Functionality

- Purchases**
Purchase History
- Financial Info**
Payment Info
- Location**
Precise Location
Coarse Location
- Contact Info**
Physical Address
Email Address
Name
Phone Number
- User Content**
Customer Support
- Usage Data**
Product Interaction

(b) Chipotle's privacy label in the iOS App Store

(a) DoorDash's privacy label in the iOS App Store

and provide specific suggestions.

This study was published at PoPETS 2022 [256].

5.2 Method

Specifically, we conducted semi-structured interviews with 24 iOS users over Zoom between mid January and early March of 2022 to explore their experiences, perceptions, and understanding of Apple’s privacy labels inside the iOS App Store. This study was approved by our institutional review board.

5.2.1 Recruitment and Screening

We recruited participants online through postings on about two dozen local Craigslist and sub-Reddit forums for geographic locations throughout the U.S. The recruitment posts did not mention or refer to privacy. Potential participants responded to our recruitment posts by filling out a short screening survey to confirm their eligibility (age 18 or older, able to speak English, located in the United States). Participants who had downloaded one or more apps from Apple’s App Store in the past three months were qualified, and some were invited to participate in the follow-up interview. We employed a purposive sampling method [236] to ensure a diverse sample of participants based on age, gender, and occupation. Among the 148 people who completed the screening survey, only one person was disqualified due to not having downloaded an app in the past three months.

5.2.2 Interview Protocol

Participants completed the online consent form prior to the scheduled Zoom meeting. At the start of the Zoom session, participants were informed that they could stop the interview at any time or decline to answer a question, and then they were given the opportunity to ask the researcher any questions. They were instructed not to disclose any personally identifiable information.

The lead author conducted all 24 interviews, which on average lasted 64 minutes. At the end of the interview, each participant filled out a brief post-survey with additional demographic questions. All interviews were recorded via Zoom and transcribed by a commercial transcription service with participants’ consent. We sent each participant a \$25 Amazon.com gift card via email. The screening survey, interview script, and post-survey can be found in Appendix C.

We first asked participants about their experience using an iPhone and downloading apps, and asked them to walk us through a recent experience downloading an app. Then, we asked participants whether they had ever wondered about or investigated what information apps collected about them and whether privacy was an important factor when downloading apps.

Later, we asked participants to share their iPhone screen through Zoom for an interactive session. In this activity, we asked them to visit the App Store and read the compact (Figure 5.1) and detailed (Figure 5.2b and Figure 5.2a) privacy labels of two apps in a randomized order. For each app, we asked some specific questions related to the privacy labels, such as their understanding of terms in the labels (e.g., “Data Used to Track You,” “Identifiers,” “Product Interaction,” “Product

Personalization”), and their interpretations of the data practices disclosed (e.g., whether the app might share their data with third-party companies for advertising purposes). We also asked them to compare several similar terms (e.g., “data linked to you” and “data used to track you”) and explain the differences, if there were any.

After participants completed all the questions for both apps, we asked about their general perceptions of the privacy labels, including whether they found these labels to be useful or not, what they liked or disliked about these labels, and whether they would pay attention to these labels in the future. In addition, we also asked whom they considered to be the source of information presented in the labels (e.g., the app developer or Apple). We finished the interview by asking participants about their general privacy concerns and behaviors (e.g., whether they had read a privacy policy or not, whether they had experienced any of their data being misused).

5.2.3 Interview Design and Piloting

We carefully designed our interviews and set the question order so as to minimize any priming of participants. We iteratively piloted and refined the interview protocol with 5 volunteer participants and 1 recruited participant. Our interview protocol is designed to learn about participant awareness of privacy labels before we mention them, and then to learn about participant understanding of the labels.

In choosing which privacy labels to show, we originally designed the protocol to let participants view the privacy labels of an app they recently downloaded and an app they use frequently. After piloting, we decided to fix the apps that participants reviewed due to the unpredictability of apps and large variances of these apps’ privacy labels. We considered one pair of most downloaded apps for each of four app categories (Shopping, Social Networking, Finance, and Food&Drink). Two pilot participants expressed a strong preference for a particular social networking app for privacy reasons. Finance apps also elicited greater privacy concerns from participants, potentially making the results more app-specific. We decided to use Chipotle (Figure 5.2b) and DoorDash (Figure 5.2a) because the two apps are similar in nature, but their privacy labels are very different and together cover a variety of terms and topics introduced in Apple’s privacy labels.

5.2.4 Data Analysis

All transcribed interviews were cleaned up by the lead author and analyzed using inductive coding [29]. The lead author met with the research team several times to discuss the first few interview transcripts and generate an initial codebook. The lead author then coded the rest of the data individually using the codebook. During the process, the research team met as needed to improve the codebook and to discuss any perceived ambiguities. Given the qualitative and exploratory nature of the study, these methods were deemed sufficient [150]. The final codebook includes 204 codes across 64 categories. We released the codebook and the redacted interview transcripts with codes via the Open Science Framework¹.

¹<https://osf.io/47kzt/>

#	Age	Gender	Education	Tech Exp	Occupation	Employment Status	iPhone Usage	# of Apps	Recent App Download	Chipotle	DoorDash	Order
N1	60-69	F	Bachelor's	No	Event planner	Retired	4	82	<1 week	No	No	D
N2	40-49	F	Bachelor's	No	Administrator	Full-time	10	79	<1 week	No	Installed	C
N3	30-39	F	Bachelor's	No	Administrator	Full-time	10	103	<1 day	No	Installed	D
N4	18-29	M	>Master's	No	Neuroscientist	Full-time	14	115	<1 day	No	Installed	D
N5	50-59	F	>Master's	No	Administrator	Full-time	10	71	<1 month	No	Installed	C
N6	18-29	M	Bachelor's	No	Hair Stylist	Full-time	6	36	<1 week	No	No	D
N7	30-39	F	Bachelor's	No	Contracting	Full-time	7	135	<1 week	No	Installed	C
N8	40-49	F	Some college	No	Homemaker	Homemaker	4	102	<1 day	No	Installed	D
N9	40-49	F	Master's	No	Project mngr	Full-time	7	75	<1 day	Installed	No	C
N10	30-39	F	Master's	No	Operation mngr	Full-time	14	126	<1 day	No	Installed	D
N11	30-39	F	Bachelor's	No	Hair Stylist	Full-time	5	180	<1 day	No	No	C
N12	40-49	F	Master's	Yes	Tech	Full-time	10	67	<1 day	Installed	No	D
N13	18-29	F	Associate's	No	Head of HR	Full-time	12	107	<1 week	No	Installed	C
N14	40-49	M	Bachelor's	Yes	Software trainer	Full-time	6	56	<1 week	No	Installed	D
N15	40-49	M	Bachelor's	No	Office mngr	Full-time	12	249	<1 week	Installed	No	C
N16	50-59	F	Bachelor's	No	Art advisor	Full-time	10	73	<1 week	No	No	D
N17	50-59	M	Some college	No	Banquet server	Full-time	6	161	<1 week	Installed	Installed	C
N18	18-29	M	Bachelor's	Yes	Urban planner	Part-time	6.5	57	<1 month	Installed	Installed	D
N19	40-49	M	Master's	Yes	Options trader	Full-time	12	97	<1 day	Installed	Installed	C
N20	40-49	M	Master's	No	HR Director	Full-time	15	319	<1 month	No	No	C
N21	30-39	F	Some college	No	Dental asstnt	Part-time	8	231	<1 day	Installed	Installed	D
N22	18-29	M	Bachelor's	No	CWO	Part-time	8	19	<1 month	No	No	C
N23	30-39	M	Bachelor's	Yes	Reseller	Full-time	12	324	<1 day	Installed	Installed	C
N24	18-29	M	Bachelor's	No	Waiter	Full-time	10	121	<1 week	Installed	Installed	D

Table 5.1: Participant demographics. “Tech Exp” refers to whether they have held a job or degree in computer science or a related field; “iPhone Usage” refers to the number of years using an iPhone; “# of Apps” refers to the number of apps installed on their phone as shown in iPhone Settings; “Recent App Download” refers to when participants reported having last downloaded an app on to their phone; “Order” refers to the app that participants visited first in the App Store with “D” representing DoorDash and “C” for Chipotle.

Due to the qualitative nature of this work, we try to avoid using exact numbers but adopt a consistent terminology to convey the relative sense of the frequency of major themes, similar to prior work [66, 95]. We use the terminology shown in Figure 5.3 to characterize the frequency of participant responses.




Figure 5.3: Terminology used to present relative frequency of themes.

5.2.5 Demographics

Among all participants, 11 were male, and 13 were female. They represent a diverse background in terms of their age, education, technology experience (whether they said “yes” to the question “Have you ever held a job or received a degree in computer science or any related technology field?”), employment status, and their general understanding of what companies are doing with their data.

We present participants' demographic information and some descriptive statistics about their iPhone and app usage in Table 5.1.

Our study participants were experienced iPhone users with a median of 10 years of usage. Ten of them downloaded one or more new apps within a day of the interview, 10 within the previous week, and 4 within the previous month. Participants estimated having on average 59 apps on their phone, far below the actual average number of 124 apps, which was obtained from their iOS Settings during the interview.

5.2.6 Limitations

Our study focuses on iOS privacy labels viewed on iPhones and is qualitative in nature based on a purposive sample recruited from location-specific online fora. Our sample skews more educated and younger than the general U.S. population. We aim to describe some of the challenges lay users might encounter as they interact with Apple's privacy labels in the App Store without making quantitative or generalizable claims. We recruited iPhone users with iOS 14 or above so that participants would have all potentially been exposed to the iOS privacy labels, but there could be a sampling bias resulting from targeting such a population.

We did not explore other devices in Apple's ecosystem (e.g., iPad, Mac). However, since Apple uses the same privacy label system with identical terminology and structures but slightly different layouts, many of our results could be reasonably extended to labels on other types of devices. In addition, we chose only two apps (DoorDash and Chipotle) for our study. Future work could investigate whether different apps might have induced different perceptions.

Finally, our study focused on the U.S. population and did not consider other cultural backgrounds. Future work could explore the potential impact of different languages used in the labels or expand the population coverage to account for other cultural factors.

5.3 Results

We report participants' perceptions about app privacy and Apple's privacy labels, their misunderstandings about the labels, and their suggestions for improvements.

5.3.1 Perceptions about App Privacy

We present three insights from the first part of our interview prior to introducing participants to the app labels.

Privacy Is Rarely Considered Prior to Installation

We asked participants to describe their recent or typical process of downloading apps from the App Store. Most of our participants said they already know what app to download when they visit the app store, either through recommendations from friends, articles, or ads. Some participants reported trying several apps and deleting unwanted ones. Some participants also reported searching

for keywords in the app store and selecting an app. Privacy was rarely the reason to (not) choose an app during the downloading phase except for a few participants; in comparison, utility, reviews, and cost are the top factors that participants considered before their download. As described by N19 when asked about whether privacy was a reason to (not) choose an app: “Not at that stage. No. I might download the app and then decide, oh, this is too much. Maybe I might delete it. But when I’m looking for an app, no.”

Even though privacy was rarely considered when participants were downloading apps, many participants did report having privacy concerns regarding specific apps. Some described removing newly downloaded apps because specific personal information was requested during sign-up. Some indicated that they had deleted apps, such as WhatsApp, Facebook, and T-mobile, after a data breach or learning about privacy concerns regarding these apps.

About half of the participants were not concerned about app privacy, as N10 acknowledged: “I try to ignore that and push that outta my head.”

Most Have Questions about App Privacy but Lack Usable Sources for Answers

Most participants reported having questions related to app privacy. N23 provided examples of the types of questions he had in mind: “I do think about like, what information are they taking from me? How does it affect me in my life?”

About half of the participants said they would use Google to find answers. For instance, N21 explained trying to use Google: “Google searching, I have tried, but... it wasn’t like a professional opinion. There were all other consumers who had wondered that same thing that I was wondering and what they think it does.”

N23 mentioned using Google to find answers:

If there’s an app that might be asking for permission for something and I’m like, wait, does this seem right? I will go to a third party, Google or Twitter, so like that, and just do a little bit of research on my own... make sure that what they’re asking for is actually legitimate.

Many participants considered privacy policies or terms of services as places to look for answers but reported frustration with them. For instance, N21 voiced her resignation:

I mean, they do have what they use it for in the terms and conditions. But that’s something for a lawyer to look at. I need layman’s terms. I need people’s terms. So I really don’t understand the terms and the conditions. I just hope that they use my information for the best.

Some participants reported looking into apps’ data collection practices in iPhone’s privacy settings or being informed by iPhone privacy prompts, as detailed by N20:

The only time I get concerned... when the app has a little pop-up, you know, when I’m using it, and it says this app will collect personal information about you and when would you like it to do so. And it’s like, all the time, when you’re using it, or never.

Some participants considered the app store as a place to find privacy-related information, as N14 described: “I would hope that in the app store that the part of the description of the app itself would have that type of stuff.”

Most Unaware of Privacy Labels

Even though the app privacy labels had been in the App Store for more than a year at the time of our interviews, most of our participants had not seen or read them. Among our 24 participants, a few participants said they had previously read an app privacy label in the iOS App Store. One of them was likely mistaken based on his description of what he had read. Some participants said they became aware of the existence of app privacy labels while scrolling past them in the App Store but did not stop and read them. Most were unaware of the labels. Many participants reported not scrolling down on app pages to see the labels, as N21 commented, “Don’t think I go all the way down there.” Others simply did not see them. For instance, N20 acknowledged, “No. If I did, then I glazed over it. This is the first time I’ve ever consciously seen it.”

5.3.2 Perceptions of Privacy Labels

This section reports how participants perceived the app privacy labels after they examined the App Privacy section for both apps.

Most Found Labels Useful

Most participants reported finding the labels useful. For example, N4 said the labels compared favorably to other types of privacy notices:

I think it is useful because as society at its whole and people individually are caring more about their privacy. So it makes sense that companies and app providers are forced to actually display this stuff in a way to the customer that is not completely incredibly difficult to understand like in a 50-page ToS [terms of service], for example.

N7 considered the label useful but also noted how inconspicuous it is:

I mean it’s useful, I think, if you specifically know what you’re looking for. I would think for most people they don’t know that this exists, you know, like they might just scroll through and again, see the comment, like this app’s amazing, or how many stars, or when it was last updated, but I don’t know that people really know that this much information is provided in that section.

On the other hand, some participants did not consider the labels useful. For example, N14 was not satisfied with the labels: “It just definitely feels like it’s like a company fulfilling a requirement and not necessarily like trying to tell the consumer what’s happening.” Similarly, N13 noted, “It’s so vague that... it’s like a ‘there there, don’t worry, we’re telling you exactly what we’re doing’, but in reality, you’re not telling me nothing.”

N2 acknowledged that the labels would be unlikely to impact her use of apps, “I guess it’s kind of like a nice to know, but at the end of the day... I probably wouldn’t run and delete the app.”

Most Stated Intention to Use Labels

Most participants reported that they would refer to the labels in the future. N15 described:

It makes me alarmed the amount of information that's being collected about me through apps. And I'm gonna take a serious look at it after this. I'm gonna look at it and I'll take action. I'll make sure that if there's something easy I can do to set what I'm allowing on my all apps on my phone, if I can do it in one easy step, that's what I'll do. But if not, if I have to go through each app, then I'll have to do that.

Some participants reported that they would look at the privacy labels for apps that they are not familiar with or to compare apps, as described by N18, "If I'm looking for a certain app... if I had choices among other apps that were of similar caliber... this [privacy label] might influence my decision somewhat."

A few participants even deleted apps during the interview or planned to delete apps upon completing the interview.

Those participants who said they did not find the labels useful were also less likely to say they would use them in the future.

Some Mistakenly Assumed Apple's Role in Producing the Labels

A few participants thought Apple and app developers together produced the labels, while some participants thought the labels were provided by Apple, as N8 commented, "It says it hasn't been verified by Apple. So for me, that's confusing because I would assume that... it was Apple all this time." Most assumed the labels were provided by the app developers.

Many participants correctly deduced that Apple did not review the information submitted by app developers; some learned this from the disclaimer at the top of the labels (Figure 5.2b and Figure 5.2a). For instance, N23 said, if he hadn't read the disclaimer:

Given that it's on the app store description page, I would assume it would be information either provided by or directly vetted by Apple because they seem to exert a pretty tight control over things that go on the app store itself at least. And I know that they also do some amount of reviewing of each version of each app that gets up.

The majority believed (wrongly) that Apple had reviewed or verified the information in the labels. N7 explained:

[Apple is] allowing that app, that product on their system.... So I think Apple, if they're approving that app and they're behind it, then I would think they should be checking [the privacy label].

Many Participants Would Not Trust the Labels

The majority of participants reported trusting the labels, because some trust Apple. N24 explained, "I trust that their App Store wouldn't be showing me misinformation." Some believed that they had no reason not to trust the labels or had to trust it, as N20 explained, "Why not? I'm going to trust that. It is because I have no evidence that it's not [trustworthy]."

On the other hand, many participants did not trust the labels. N10 explained her reasoning:

Because of all these things listed that I don't even understand but know that they're wrong, I understand enough to know that it's wrong and it's not okay. I just know that

they put it on here because it's mandatory that they do, but they're just doing it because they have to, not because they want me to know. There's nothing with good intentions.

Some participants were concerned about the vagueness of the labels, as N13 explained, "I trust it as far as I could throw it. I mean, it's reliable but very vague. So it's just like... I don't know to what extent you're really doing this."

Both labels included an "other" category, which concerned some participants and contributed to their concerns about vagueness. N4 noted, "I would not trust it a hundred percent. So just because of wild cards, like other data." We discuss this further in Section 5.3.3.

5.3.3 Misunderstandings of Privacy Labels

In this section, we report the misunderstandings that our participants expressed as they systematically looked at two privacy labels following the interview protocol in the Appendix. We focus on overall understanding, confusing terminology, and vague language.

Overall Understanding

At a high level, about half of our participants had misconceptions about the requirements for what is disclosed in an app privacy label, and about half were confused about the label structure itself.

What needs to be disclosed About half our participants mistakenly assumed the label included all app data collection and usage. However, Apple does not require disclosure of all app data collection: disclosure is optional for data that are not used for tracking or advertising purposes, collected infrequently, and in the app's user interface [11].

On the other hand, about half of our participants said they did not believe or were not sure that the privacy labels show all of the app data collection and handling. They were concerned that companies were only disclosing what they had to disclose. N14 explained:

It almost seems like the developer, the companies only show you what they're forced to show you. And it's possible that the development of data capturing features and functionality is faster than the regulation to regulate it. And they only do so most of the time if they're in fear of getting in trouble for it.

A few participants were puzzled by why only some of the three boxes were displayed. N18 noted:

The DoorDash app, it had "data linked to you", you know, "data used to track you"... whereas for Chipotle, they're just saying what's not linked to you.... I understand what they're not using to link to me, but then what are they using to link to me? Right. You can't just assume that because it's not on here, that means they aren't doing it.

Confusing label structure The compact privacy label for each app, as shown in Figure 5.1, shows up to three boxes for three categories of data: data used to track you, data linked to you, and data not linked to you. Within each box is a list of specific data types, each accompanied by an icon. If there is no data for a particular category, that box is omitted. Users who click on one of the boxes

or the “See Details” link at the top of the label are taken to a more detailed view that shows the same three categories. The “data used to track you” category shows specific data types, with an even more detailed list under each type. However, the other two categories are presented differently, with lists of purposes of data use under each category, and then specific types of data under each purpose, as shown in Figure 5.2b and Figure 5.2a. This structure was not readily apparent to about half of our participants, who were confused about how to find the purposes and did not understand why specific data types were shown multiple times. For example, after viewing the DoorDash label, which discloses data being used for multiple purposes, some participants said they thought the DoorDash label contained redundant information. For example, N21 noted, “A lot of the information was repeated.... I don’t see the need to keep repeating the same that it’s gonna collect my purchase, my location, stuff like that.”

When asked about what the purpose heading “app functionality” means on the DoorDash label, N10 said:

It’s a word that I figure out the meaning to, but I don’t know how it applies to all these things listed under it. So I really can’t even guess.

Similarly, N3 expressed confusion when examining the Chipotle label, and noted that she had similar concerns about the DoorDash label, “There’s not really a purpose honestly... because they all kind of say the same thing. So like with the last one too, like they were all like basically a lot of repetitive information.”

Terminology Caused Confusion

All participants expressed confusion about one or more terms used in the labels or gave interpretations of terms that did not match Apple’s definitions. Only a few participants noticed the link on the label to review the actual privacy term definitions (the second paragraph of Figure 5.2b and Figure 5.2a), and even fewer of those clicked to view those definitions.

“Tracking” is overloaded Apple’s definition of tracking indicates that data can be used to link with third-party data for targeted advertising or to share data with data brokers. However, this was not always clear to participants. For instance, when looking at the tracking section of the DoorDash label (Figure 5.2a) N19 could not tell whether data would be used for advertising purposes: “It doesn’t say anything about ads? It doesn’t say what exactly they will use this data for.” Some participants did not associate tracking with data brokers aggregating their information across other websites and companies.

Some participants who glanced at the definition of “data used to track you” were surprised to find that tracking involves sharing information with other parties rather than just collecting data on either the user’s location or website usage. For example, N1 said tracking was “tracking me like where I go,” while N5 described tracking as, “I think my usage, the frequency of ordering or using the site and identifying where I am. I think that’s all tracking. I think tracking is what are my habits.” After being told that the definition included sharing, she said, “I don’t like that at all. Like now that you say it, I’m thinking I’m deleting this app [DoorDash]. I think it’s too intrusive.”

N7 was also surprised and concerned by the definition of “data used to track you”:

Term	Definition
Data linked to you	Data that is linked to your identity (via your account, device, or other details). “Personal information” and “personal data” as defined under relevant privacy laws are considered linked to you. In order for data to not be linked to you, the developer must avoid certain activities after they collect it: 1) They must make no attempt to link the data back to your identity. 2) They must not tie the data to other data sets that enable it to be linked to your identity.
Data not linked to you	Data that is not linked to your identity (via your account, device, or other details).

Table 5.2: Apple’s definitions of data linked to you and data not linked to you [12]

It’s concerning. I will say that I didn’t realize. I kind of believed that they would track like, oh, she looks up dogs. So then just random dog things would pop up like ads or like, you know, so pick up kind of like what I was searching. But seeing this, it goes way deeper than that, you know, they’re actually collecting the information, which to me is kind of scary.

Confusion around “data (not) linked to you” The three boxes in the compact label seemed to confuse participants, despite the explanations included in each box.

Partially due to not understanding tracking, most participants could not explain the difference between “data used to track you” and “data linked to you,” or their explanation was inconsistent with Apple’s definitions, shown in Table 5.2. N2 admitted, “Well I don’t know what the difference is.... Maybe they’re just required to have both sections.”

N4 was confused about the same data categories shown under these two headings in the DoorDash label, and he wrongly assumed “data linked to you” implied data not being shared:

I was under the perception that data that is linked to me has more identifiers that make me non-anonymous, for example, my physical address and my name. But then I read that this is also in “data used to track me.” So out from the descriptions “the following data might be used to track you across apps” and “the following data, which may be collected and linked to your identity may be used for purposes.” Difficult for me to understand. I can imagine that this data is perhaps just put into a database to create a profile of me, but I don’t a hundred percent understand what it is.”

“Data not linked to you” on the Chipotle label also confused about half of the participants; most participants were confused after they saw contact information shown under this heading as in Figure 5.2b. The inclusion of contact information in this category is possibly an error, as N13 observed:

It states it’s not linked to you, but obviously it is linked to you because it’s your personal information, like your address, your email address, your phone number, and your name. These are all very personal things that are specifically you, as opposed to them stating this is like completely anonymous. So I don’t know if they’re like lying, if that’s like

Term	Definition
Analytics	Using data to evaluate your behavior, including to understand the effectiveness of existing product features, plan new features, or measure audience size or characteristics.
App Functionality	Such as to authenticate you in the app, enable features, prevent fraud, implement security measures, ensure server uptime, minimize app crashes, improve scalability and performance, or perform customer support.
Product Personalization	Customizing what you see, such as a list of recommended products, posts, or suggestions.

Table 5.3: Apple’s definitions of analytics, app functionality, and product personalization purposes [12]

Term	Definition
Third-party Advertising	Such as displaying third-party ads in the app or sharing data with entities who display third-party ads.
Developer’s Advertising or Marketing	Such as displaying the developer’s own ads in the app, sending marketing communications directly to you, or sharing data with entities who will display ads to you.

Table 5.4: Apple’s definitions of advertising-related purposes [12]

a fine line as in, “we collect this information for our app purposes only, and then the information that we share with other companies that’s not linked to you is you know what we are sharing.”

When users observe potentially erroneous information in the labels it may undermine trust in the labels, as noted by N21, “Now I feel like you’re lying to me because I don’t see how you’re collecting my name but then telling me it’s not linked to me.”

Entangled and overlapping definitions Most participants were particularly confused about the terms, app functionality, analytics, and product personalization, shown in Table 5.3. They either admitted that they did not know the definitions of these terms or gave their own definitions in which they mixed up the terms.

For example, N19 gave a definition for app functionality that confuses it with analytics: “By collecting this data about me and gazillion other people, they analyze it and they might change the functionality based on the patterns that they see.”

N1 mixed up app functionality with product personalization, describing app functionality as “customizing my user content.” N3 tried to define analytics as “for them to see how like the app is working, I think,” admitting, “I don’t really know what analytics is. And I don’t really even know what the difference between a lot of these.”

A majority of participants could not tell the difference between third-party advertising and

developer’s advertising, shown in Table 5.4. N8 described her confusion:

The advertising and the third party? Yeah, that’s a good question because it kind of feels like it means the same thing. I feel like the third-party advertising is somebody other than DoorDash. And then the developer’s one is, I don’t know, to tell you the truth. I don’t know the difference.

Some participants mistakenly thought that developer’s advertising did not involve third parties, as N4 suggested:

Theoretically, the difference should be that third-party advertisement has the purpose of monetizing me and making money out of me by showing me ads, whereas developer’s advertising or marketing is for developing a better product in the end. So it’s not shared with a third party, but it’s shared with DoorDash itself and when DoorDash wants to create a better product out of it.

Unfamiliar terms for frequently used data types Three frequently used terms describing data types often confused participants.

The data category “user content,” for which Apple does not provide a definition, was called out as confusing by about half of the participants. Some were also confused by the term “customer support” below the bold heading for user content in both labels. For example, not knowing what the heading meant, N10 took a guess: “User content and customers support. Maybe they’re like blocking me from getting in touch with customer support.... I don’t know. That’s ridiculous. I really dunno what it is.”

Some participants did not understand the term “identifiers,” for which Apple does not provide a definition, and a few looked to the icon next to the term for answers, noting it looked like some sort of a photo ID. N4 explained, “I don’t know what identifier is, but based on the icon... I’m guessing like identity-related information could be phone, name, email, etc.” On the other hand, N10 was confused by the icon, “What are identifiers? That looks like a license.”

“Product interaction” was also confusing to many participants, and a few of them erroneously believed that product refers to items they viewed or bought, as N18 explained:

I think the product refers to like an item, like anything tangible, right? Cause you know, for DoorDash, you’re purchasing items through them. Typically it’s like food or drink. So I guess to that end, what specific restaurants am I ordering from frequently? What foods am I buying frequently if I am ordering online, stuff like that.

Other data type terms that confused participants included “usage data,” “diagnostics,” and “coarse location.”

Term	Definition
Product Interaction	Such as app launches, taps, clicks, scrolling information, music-listening data, video views, saved place in a game, video, or song, or other information about how you interact with the app.

Table 5.5: Apple’s definition of product interaction [12]

Term	Definition
Browsing history	Information about the content you have viewed that is not part of the app, such as websites.
Search history	Information about searches performed in the app.

Table 5.6: Apple’s definitions of browsing and search history [12]

Vague Language on Labels

The “other” category is alarming Almost all participants responded with confusion when they saw terms in the label related to other data types or other purposes since they did not know what these other categories could entail. About half of the participants said these terms made them anxious or decreased the utility of the label. N3 asked, “Why even write all the details, if you’re just gonna say other?” Similarly, N22 commented:

I don’t like it. This section seems to be like an extra section for people who want to be clear on what’s being used, and then they give vague answers. Anyway, it doesn’t make it pointless, but it makes it less useful.

N4 considered this a way for companies to make sure they covered everything that was required:

I don’t think that a company would blatantly lie to me here. This, in my opinion, would be bad intent of the company. And I don’t think that the company would do this nowadays because it would be very, very stupid if this comes out like Dieselgate, you know, but as long as they can put something in there, which is other data, I can imagine that they are covered anyways, you know, because what is other data it can put in anything.

The scope of browsing and search history is unclear Most participants made assumptions about browsing and search history that did not align with Apple’s definitions, shown in Table 5.6, or became confused or concerned when they learned about the definitions.

N19 looked at the DoorDash label and tried to interpret the meaning of browsing history: “Is it the DoorDash website’s browsing or any browsing? Is it the DoorDash app or any app? I don’t know.” The majority of participants were concerned when they learned that browsing history includes “content you have viewed that is not part of the app.” N4 commented, “So then I think this should be a bigger topic that is communicated overall, that just by having food delivered to you, you are showing the world what you browse on your smartphone.”

Even though search history is defined to be only within the app, some participants assumed it referred to other searches as well. For instance, N17 commented:

I would hope that it’s only keeping track of searches on my apps. In other words, DoorDash, I go on there and type in “tacos.” It will come up with tacos. But it doesn’t actually say that. So it says “search history.” It could also be keeping track of all the searches I do on Google.

User content: emails, texts, photos, or videos Many participants were concerned or not sure about to what extent the data listed under user content could be used in the DoorDash label. For instance, N22 commented, “I’d hope that this user content section is just stuff in the app for customer support purposes. But if it’s looking at your emails or texts or photos or videos outside of the app, that’s very disconcerting.”

Many participants were disturbed by this kind of data collection, as N8 explained:

User content, yeah. They’re gonna check how I use the app, how many purchases I’m making, where I am at the time that I’m using it. I don’t know what they would need emails or text messages or photos or videos. I don’t know why they would need that. I didn’t know they would have access to my photos. That’s kind weird. That’s kind of creepy. I don’t know what food has to do with my photos.

Most participants were concerned to see data collection that they perceived as unrelated to the purpose of the app, such as their contacts being used by DoorDash.

Term	Definition
Emails or text messages	Including subject line, sender, recipients, and contents of the email or message.
Photos or videos	Your photos or videos.
Audio data	Your voice or sound recordings.
Customer support	Data you generate during a customer support request.
Other user content	Any other content you generate.

Table 5.7: Apple’s definitions of terms under “User Content” [12]

5.3.4 Suggested Improvements

Although most participants liked the label concept, they had a number of concerns, as discussed in the previous sections, as well as ideas for improvements. In this section, we report on improvements to the labels suggested by study participants.

Better Structure for Data and Purposes

As discussed in Section 5.3.3, participants found the structure of the label confusing, and many had the misconception that some sections were redundant. Underlying this confusion is that the label designers mapped a multi-dimensional space of data types and purposes onto a list-based label representation. A matrix or tabular approach might offer a more compact and intuitive representation, as N5 suggested: “I do think those last two sections [app functionality and product personalization] were very redundant, and so I think they could do it by like a table with a bunch of checkboxes....”

Easy Access to Definitions and Contextualized Examples of Data Collected

The majority of participants suggested making the definitions of label terms more accessible. Currently, the definitions of many of the terms are available through a link from the detailed view to a web page with the definitions all in one place. However, many participants wanted to see each term linked to its definition, perhaps appearing through a hover. For example, N4 explained:

It would be cool if this information would be hot linked there, you know, like there is this symbol with an “i” in it, which means information, you know, that would be cool. So like browsing history, for example, click, you know, I’m like, “Ooh, what does this mean?”

Others, including N8, suggested using “terminology that’s just easier to understand.” Many participants suggested providing concrete examples of the data being collected. N20 suggested:

If I click user ID... I know what a user ID is, but tell me which user IDs are you tracking. So Is it me, my wife, and my kids? ... What exactly are you tracking? You know, is it my phone? Is it my watch? Is it my iPad? You know what’s all really linked to user ID?

N17 also suggested that specific examples should be listed for “other” categories to ease concerns about what might be included.

Embedded Actions and Controls

Many participants voiced their frustration with the labels due to their lack of controls and suggested controls to turn off some of the data collection. N19 explained:

I don’t like it that it’s too long and you can’t really take any action on it. It’s really just informational and you can’t really turn it on off, etc. at this level.... I’d rather just be able to turn off whatever is not required.... If any of this tracking is optional, I wanna be able to turn it off.

N2 described in detail what she envisioned:

Maybe by letting you check off things that you don’t want included or make it easy to like opt out of all of this. Well, every app asks you when you install it if they can like track and share your stuff, so make it easier. So I don’t have to learn or go hunting.... It could be like a checkbox or a radio slider or something....

Access to More Information

Participants suggested other topics related to app privacy that they would like to see added to the label.

Some participants asked for information about data retention, as N18 articulated, “How long is this information stored for, right. If there was a clear understanding that your information’s gonna be stored for 30 days, that would probably give me a lot more solace than not knowing right now.”

Some participants were interested in knowing to whom the apps are sending their information. N7 said she would like to know “where exactly is it being [sent]? Is it being sold? Is it being just shared back and forth so that there’s this hub that everybody uses?”

A few participants wanted to know where they could have questions answered or read more detailed information. N3 wondered:

So where would I go to like ask somebody or chat with somebody or like, there should be like another link that takes you to dive in deeper if you wanted to know, because I don't know what I would even like, there's no contact information. And who would I even ask about this? It's kinda useless if I do have a question and I don't know who to ask them, it kind of seems a little useless.

Finally, a few users wanted to better understand the privacy-related implications of using the app and any measures the app was taking to protect their privacy. N14 asked, "How would you see this occur or affect you? Like because of having this app, these are the things that are happening to you, like you're seeing targeted ads, you know. It relates to the user more."

5.4 Discussion

With hundreds of thousands of mobile apps now featuring privacy labels in iOS 14, these labels are for the first time available at scale to mobile app users. The introduction of privacy labels is an important step towards empowering users to better understand mobile app data practices that matter to them. At the same time, in their current deployment and configuration, these labels are not as usable or effective as they could be.

5.4.1 Helping Users Comprehend Complex App Privacy Practices

After examining the labels in our study, almost all participants learned new things that they did not know before and appreciated the existence of the labels. Also, about half of the participants regarded the privacy labels as useful and most reported being likely to use them in the future.

However, the labels suffer from confusing terms and definitions (see Section 5.3.3), which led to a range of misunderstandings. In addition, vague language (e.g., "other" category, user content) impede participants' understanding of the actual data practices (see Section 5.3.3). These findings clearly demonstrate that Apple's privacy labels still fail to fully support user comprehension of the disclosed app privacy practices.

The linear structure of the labels, which is presented differently in the compact and detailed views, seems to do a poor job of communicating the multi-dimensional space of data practices where multiple categories of data (each of which is represented in a multi-level hierarchy) are used for multiple purposes. A tabular representation may be more compact and intuitive [114, 196], although the small form-factor of mobile devices may present design challenges. Additional work is needed to better understand which label elements are most important to users so that the compact version might focus on those elements.

Our findings on end-user misunderstandings extend recent studies that showed how app developers often struggle with privacy label definitions (e.g., interpretation of terms such as "tracking") and how this hampers their ability to create accurate labels [83, 134]. Our study focused on lay users who lack technical expertise and experienced a high level of confusion.

The addition of links or hover text to provide more ready access to definitions of terms and examples might aid comprehension.

With the recent rollout of Android privacy labels, we have observed that definitions of terms such as “tracking” are not completely consistent on the Android and iOS platforms. It would be helpful if the industry were to adopt standard terms and definitions for privacy labels, empirically tested with both developers and lay users. An earlier multi-stakeholder effort led by the U.S. Department of Commerce resulted in a standard set of terms for app transparency, but did not include terms describing purposes of use and the terms were not updated after user testing found them to be confusing to both experts and lay users [19].

5.4.2 Improving Privacy Labels’ Saliency

Even though the privacy labels were introduced in Apple’s App Store over a year before our study, the majority of our participants were still unaware of them. Our finding shows that the discoverability of privacy labels on each individual app’s page in the App Store is low, even for participants who said they were concerned about mobile app privacy. As currently deployed, users have to scroll past several sections, including images, Ratings & Reviews, and What’s New, before finding the App Privacy section. Our findings on discoverability are corroborated by prior research that has shown that the location and timing of privacy labels and indicators can have a large impact on whether users pay attention to them [20, 63, 96].

Our results suggest the need for more prominent placement of privacy labels, consistent with recommendations to display concise privacy notices in salient ways [62]. Alternatively, it would be beneficial to add standardized indicators (e.g., links, icons) to signal the existence of these labels during users’ app installation decision-making process.

There is also a need for additional mechanisms to bring users’ attention to privacy labels for apps that users already have on their phones. For example, iOS privacy nudges [13] about background app data collection, just-in-time app permission requests, and iPhone permissions setting interfaces are potential places to include links to the privacy labels that would increase both awareness and convenience.

5.4.3 Promoting Privacy Labels’ Role in App Privacy Management

Another key complaint from participants is that the privacy labels do not offer control options (see Section Section 5.3.4). Some participants reported being disappointed that even after learning the information presented by the labels, they were not provided with any actionable steps they could take. Information on the labels is not readily accessible in the permission settings (i.e., the permission manager) where users decide which permissions to grant to each app. It would be helpful if the permissions manager included the relevant information for each app that appears in the app store label. For instance, the iOS app tracking permission could be incorporated into each label with a toggle control.

In addition, the controls offered to users in the current permission manager are not aligned with the information conveyed in the privacy labels. For example, while a privacy label might inform users that their location information might be used by the app for multiple purposes such as for

the app’s core functionality as well as for advertising purposes, users do not have the option to grant an app access to their location for one purpose and not for another (e.g., granting access for the core functionality but not for advertising purposes). A few participants were puzzled by DoorDash never requesting the Contacts permission despite listing it in the privacy label. Even worse, when users deny a particular permission, for example, location access, some apps might still be able to extract location-related information from IP addresses, metadata associated with uploaded user photos, WiFi connections, etc [59, 86]. Such misalignments between the disclosures made in privacy nutrition labels and privacy controls made available to users create another potential source of confusion.

Furthermore, if users have already selected the global setting to turn off app requests to track, it is unclear whether any of the tracking indicated in the “Data Used to Track You” section could happen or not. It might be helpful to include a toggle to allow users to turn off tracking directly in the label and indicate appropriately whether the user has previously configured that setting.

Privacy labels are shown within the descriptions of individual apps in the App Store, but no functionality is provided to enable users to compare apps or look for equivalent apps with less invasive or more desirable data practices. The App Store should enable users to search for apps that meet certain privacy criteria, for example, filtering similar game apps that do not collect any location information or picture editing apps that do not involve sharing user information with data brokers.

5.4.4 Reducing User Burden in App Privacy Management

Ultimately, privacy labels are designed to empower users to quickly find answers to some of their most common questions and save them the time and effort that would be required if they had to read the text of privacy policies. Even though privacy labels offer the promise of providing users with more succinct and more effective notifications, given the large number of apps on each user’s phone, it is unrealistic to expect users to go through the privacy labels for each app one at a time. Prior work using machine learning and natural language processing techniques to automatically extract and analyze disclosure statements from the text of privacy policies [190, 235, 260], including privacy question answering functionality [189], has been technically challenging. With the help of these standardized notices, it will be more feasible to automatically extract relevant privacy disclosures, which in turn can support chatbot functionality to quickly address users’ questions or refer them to parts of the labels pertaining to their questions.

Another way to decrease user burden is to leverage the operating system or a personal privacy assistant to act on behalf of users instead of relying on users to manually configure every app setting. Users could be selectively notified about the types of data collection disclosures that they would like to be reminded about and only show users relevant disclosure information that they personally care about [139, 218, 248].

5.5 Conclusion

While iOS app privacy nutrition labels offer the first wide-scale deployment of standardized short-form privacy notices, our qualitative interview study highlights the barriers that prevent these labels from achieving their desired impact when it comes to actually helping users. Findings from this work provide the basis for concrete recommendations to refine existing labels, potentially delivering benefits to millions of smartphone users, as well as informing the design and effective deployment of similar privacy labels on other platforms (e.g., Android) and in other domains (e.g., websites, Internet of Things).

Chapter 6

Evaluating the Effectiveness of Mobile App Privacy Labels: To What Extent do Apple and Google Privacy Labels Address People’s Privacy Questions?

6.1 Overview

In Chapter 5, we aimed to evaluate the effectiveness of mobile app privacy nutrition labels by investigating whether users are aware of the labels, whether they use them, whether they correctly interpret the information provided by the labels, and more generally whether they can effectively use the labels to support their privacy decisions. However, independently of how usable privacy labels are, another key question is to what extent these labels actually address the privacy questions mobile app users have. This is the question we examine in the present chapter.

Our study relies on the analysis of a corpus of 1,750 privacy questions [189] collected from mobile app users using Amazon Mechanical Turk¹. We compare the content of this corpus with the information provided by mobile app privacy labels in both the iOS App Store and Google Play Store and examine whether these labels actually address people’s mobile app privacy questions. Our analysis indicates that mobile app users have rather diverse privacy questions, and although there are differences between iOS labels and Google Play labels, an important percentage of people’s privacy questions are not answered or only partially addressed by these labels. Our findings suggest that existing mobile app privacy labels can benefit from improvement to better address people’s most typical privacy questions. However, given the already complex nature of these labels and the diversity of people’s privacy questions about their mobile apps, it is also clear that privacy labels can only go so far. The information in privacy labels may benefit from being organized differently (e.g., using expandable tabular formats), and may need to be supplemented with additional functionality such as privacy question answering functionality (e.g., [189, 190]).

Our study relies on the analysis of a corpus of 1,750 privacy questions [189] collected from

¹<https://www.mturk.com>

mobile app users, using Amazon Mechanical Turk². We compare the content of this corpus with the information actually provided by mobile app privacy labels in both the iOS app store and the Google Play Store and look at whether these labels actually address people’s mobile app privacy questions. Our analysis indicates that mobile app users have rather diverse privacy questions, and while there are differences between iOS labels and Google Play labels, an important percentage of people’s privacy questions are not answered or only partially addressed by these labels, whether iOS and Google Play Store labels. Our analysis provides a roadmap for the possible refinement of existing labels if one were to try and improve the way in which they cover people’s most typical privacy questions. Because existing mobile app privacy labels are already rather complex and because people’s privacy questions about their mobile apps are also quite diverse, this research also suggests that privacy labels can only go so far in addressing the diverse privacy questions people have. Information in privacy labels may benefit from being organized differently (e.g., expandable tabular formats) and may need to be supplemented with additional functionality such as privacy question answering functionality (e.g., [189, 190]).

The research presented in this chapter was published at USEC 2023 [228].

6.2 A Dataset of Privacy Questions Users Have About Mobile Apps

As part of a study of mobile app privacy question answering functionality, Ravichander et al. [189] collected a dataset (the “PRIVACYQA” dataset) of privacy questions that people had for a diverse sample of 35 mobile apps. The set of apps were selected to include well-known apps and apps with smaller install bases, also covering a broad range of application categories across the Google Play Store. The study, which involved recruiting Amazon Mechanical Turkers, asked each participant to provide five free text questions per application related to a subset of 35 mobile apps. The study was designed to elicit questions that mattered to participants as they were presented with the name, description, and navigable screenshots of the app as shown in the Google Play Store. The resulting dataset comprises 1,750 questions. Though the authors cannot make any hard claim about how representative this dataset is, it provides a sufficiently diverse collection of privacy questions to warrant comparison with the content of the mobile app labels. Our study leverages this publicly available dataset and explores to what extent these questions can be addressed by iOS and Google Play mobile app labels.

6.3 Methodology

We selected this specific dataset [189] because the questions in this corpus were elicited in a context intended to mimic that of a user examining an app in an app store. Participants were presented with information about an app, including its name, description, and navigable screenshots, similar to

²<https://www.mturk.com>

Question Theme	Types of Questions Under This Theme	Count	%
Data collected	Does the app collect PII, location, search history, payment, texts, health, calls, IP, calendar, other? Is any information recorded?	364	22.1%
App security	How secure is the app? Is my payment information secure with the app? How will my password be stored?	199	12.1%
Sharing	Is my data shared, with whom, and what data is shared?	151	9.2%
Selling	Is my data sold, to whom, and what data is sold?	141	8.6%
Permissions	Any permission required to run this app? Does it have access to my camera or access to my microphone?	140	8.5%
App-specific privacy	Is my status in the app visible to other users?	128	7.8%
Purpose	Will the app use my data for marketing purposes? Why do you need those permissions?	95	5.8%
Who has access	Do app company employees have access to my data? Can the government request my data?	73	4.4%
Privacy risks	Will the microphone secretly be turned on to listen to my surroundings?	64	3.9%
Retention	Will my data be saved permanently? For how long is my data kept?	56	3.4%
Privacy controls	Can I make my profile private? Is there a way to opt out of data sharing?	49	3.0%
Retained method	Will it store any information on my phone? How do you store my data and information?	40	2.4%
Account required	Do I need an account to use this app? Do I have to sign in using a social media account?	37	2.2%
External access	Does [APP] look at other stuff on my phone besides in app? Does the app have access to financial apps I use?	32	1.9%
Deletion	Do I have any rights as far as whether I want my account info deleted?	31	1.9%
Privacy protections	What safeguards does the app use to protect the privacy of my data?	29	1.8%
Privacy policy	Is there a privacy policy? Where can I read your privacy policy?	9	0.1%
Cookies policy	Do you use or collect cookies?	9	0.1%
		Total:	1647

Table 6.1: Themes identified, types of questions under each theme, and the number of questions under this theme

what one would find in an app store, and were instructed to ask privacy-related questions about the app.

Our first goal of this study was to understand the nature and topics of questions asked by users in the corpus [189]. We applied thematic analysis as an organizational tool to classify and describe the questions, as well as a process to interpret, connect, and transform the questions into themes [119]. The lead author first familiarized herself with the data by reading all 1,750 questions. Subsequently, the lead author coded all questions, generated an initial codebook, and met with the second author several times to refine the codebook. The lead author then re-coded all questions individually using the finalized codebook. Given the qualitative and exploratory nature of the study, these methods were deemed sufficient [150]. The final codebook includes 18 themes with 67 codes. The themes and example questions are shown in Table 6.1. Most questions were labeled with one code, while 60 questions were annotated with more than one code, totaling 1,647 codes. The thematic analysis and the generated themes help us to better understand the corpus and also facilitate our next step.

The second objective of this study was to evaluate whether users' privacy questions could

respectively be answered by the iOS and Google Play privacy labels. To minimize the impact of app developers' inaccuracies in specifying privacy labels or of apps that may not have been published on both platforms or lack privacy labels, we made the assumption that developers utilized the privacy labels optimally to disclose their apps' privacy practices. This means we did not examine the actual privacy labels within the app stores, but instead evaluated if the iOS and Google privacy labels have the capability to address user privacy questions.

Both authors analyzed and discussed each of the 67 codes (sub-themes) to determine if questions under each sub-theme could be answered using the labels provided by Google or Apple. We randomly sampled example questions for each sub-theme, compared them to the definitions of the Google and Apple labels [11, 89], and reached a consensus on whether those questions could be answered or addressed. We deemed a question fully addressed by the app labels if any part of the label, including definitions, contained implicit or explicit answers to that question. We provide further explanations of implicit answers in Section 6.4.3. We considered a question partially addressed by the app labels if the presence or absence of a label section provided relevant information but not a complete answer. Table 6.2 shows examples of answerable and partially addressed question themes and the corresponding label or definition snippets that can be used to answer these questions. Following the analysis process, the lead author conducted an evaluation of all questions in the data set to determine whether each question could be answered or partially addressed using either the Google or Apple labels.

6.4 Results

6.4.1 Question Themes

We present all 18 question themes resulting from the thematic analysis in Table 6.1. The theme with the highest frequency, accounting for 22% of all questions, pertains to the data being collected by apps. These questions, typically phrased as “Does this app collect X?” were interpreted as “Can this app collect X?” considering that privacy labels do not necessarily indicate actual data practices but rather the potential for data collection. For instance, an app might be able to collect the user's GPS location as long as the user does not deny the app access to their GPS location. Approximately one-sixth of these *data collected* questions, totaling 60, pertain to what types of data are collected. Over 20 questions address the issue of whether the app collects data at all, including questions such as “Do you keep my data and upload to your server?” A handful of questions pertain to whether the collected data is anonymous or not. The remaining *data collected* questions are related to whether specific data types of data are being collected, such as search history, contacts, usage data, etc.

Approximately 12% of questions in the corpus are related to app security, encompassing a variety of topics such as the overall security of the app, inquiries about recent security breaches, technical questions such as whether data is encrypted or whether security protocols are being used, how the app handles passwords, or whether payment data is secure. It is worth noting that even though participants received prompts to ask privacy-related questions, they asked security-related questions as well, indicating that they view their security questions as legitimate privacy questions.

The third and fourth most frequently asked question themes were about data sharing and selling,

respectively accounting for 9.2% and 8.6% of all questions. Participants wondered whether their information is shared/sold, what type of information is shared/sold, and to whom. The fifth theme revolves around the types of permissions that apps might need, such as whether a specific permission (e.g., camera, microphone) is accessed or the necessary permissions for the app to function properly. The sixth theme pertains to specific privacy questions related to the functionalities and features of the app. For instance, one question regarding the *TripAdvisor* app says, “Can I review stuff without having my name attached?” Another question about the app *Recipe, Menu & Cooking Planner* reads, “Will anyone see the recipes that I upload?”

Together, these top six themes (one-third of all themes) add up to approximately two-thirds of all questions asked. The remaining 12 themes are listed in Table 6.1 and account for just one-third of all the questions.

6.4.2 Question Themes Mostly Answered by Labels

Table 6.3 shows a summary of the question themes that can or cannot be answered by the labels. As seen in the table, both the iOS and Google labels include information on the collected data categories and can answer most of the questions. Note already that not all questions can be answered by the labels and that this varies between Apple and Google. For instance, Google labels specifically mention the collection of IP addresses and calendar information, while iOS labels do not include these data categories. Google labels also allow developers to indicate what data types are optional “where a user has control over its collection and can use the app without providing it” [89], therefore answering a few of questions related to what data are required to use the app. A handful of questions with regards to whether the collected data is anonymous can be answered using only Apple’s privacy labels, as these labels include a section on “data not linked to you.”

While Google labels contain information on security, iOS labels do not mention security at all. App developers can declare optionally in Google labels that their app “has been independently validated against a global security standard... MASA (Mobile Application Security Assessment)” [89]. This review³ covers a wide range of security-related topics [9], addressing many questions in the app security theme, such as password handling and encryption. However, the review does not cover all the security questions the participants had. For instance, it does not indicate how payment information is stored or cannot help answer questions about whether an app has had a breach in the past. We considered user questions such as “Is the app secure?” or “What protection do you offer against hackers?” answered if the Google label for an app indicates that an optional review has been conducted.

Both labels provide information on whether an app shares data and the types of data being shared, but neither label directly states with whom the data is shared, only referring to third parties in general. Apple labels only require disclosure of data sharing when it is used for advertising or “tracking,” while Google labels require developers to disclose any non-first-party sharing, unless it is for legal purposes or if the data is anonymous. No data sharing needs to be reported if the action is clearly a “user-initiated action” with clear disclosure and user consent [89] or “it is clear to the

³<https://github.com/appdefensealliance/ASA/blob/main/MobileAppSecurityAssessment/MobileSecurityGuide.md>

user what data is collected” [11].

Both Google and Apple labels address most questions related to the purpose of data collection. While they adopt slightly different definitions of purposes, both include categories such as app functionality, analytics, personalization, and advertising or marketing. Questions such as “Why is my data needed?,” “Will you use my data for advertising?,” and “What does the app do with my personal information?” can be answered by both labels. However, some questions about the specifics of how data is used for personalization or advertising, such as “How are features personalized?,” are not addressed by either label.

Only Google labels contain information related to data deletion [89]. However, some questions pertain to the deletion of specific data types rather than the complete removal of a user’s information. Such questions cannot be answered.

6.4.3 Implicit Answers

Two themes contain questions that cannot be directly answered directly from the labels, but the label definitions contain implicit answers. Questions under that “data collected” theme, such as “Do you keep the data of mine and upload to your company?” can be inferred from the fact that both labels ask developers to declare user data that is transferred out of users’ devices, implying that the data listed on the labels is uploaded to servers. As per both labels, data solely residing on users’ devices are not considered to be “collected.” Similarly, while the labels do not use the term “selling,” questions about whether data is sold are addressed by information provided under Apple’s “data used to track you” section as Apple defines tracking to include user data sharing with a data broker.

6.4.4 Question Themes Not Addressed by Labels

The lower half of Table 6.3 lists the themes identified in Table 6.1 that are not addressed by either Google’s or Apple’s labels, as evidenced by the zeroes under the Google and Apple columns. The most frequently asked theme (8.5% of all questions) pertains to permissions, with participants asking whether an app accesses specific permission(s) of the phone. It is important to note that “accessing” information in an app does not equate to collecting that information. Data collection only occurs when the information leaves the device. In other words, label entries about data collected by an app do not allow us to answer questions about permissions used by an app. This is the case for both iOS and Google.

Other questions related to permissions, including questions on the necessary permissions needed in order for the app to function (e.g., “What type of permissions does the app need to operate?”), can also not be answered by either label. About 4.4% of questions ask who has access to their information in general or specifically inquire about whether specific entities such as the government or employees of the app company may have access. 3.4% of questions pertain to data retention. 2.8% of questions are related to whether the app requests external access to other apps, accounts, or data outside the app. 2.2% of questions are about whether users are required to create an account or use a social media account to use the app. Nine questions (0.6%) are related to how the app handles cookies. These questions are relatively easy to answer not only because they request factual answers

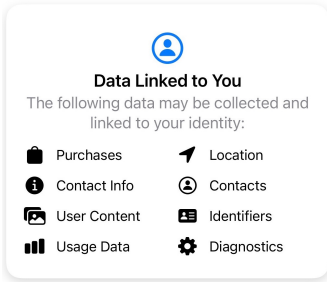
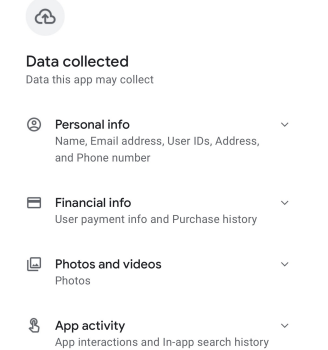
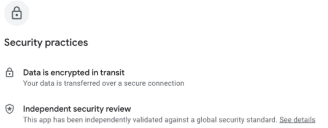
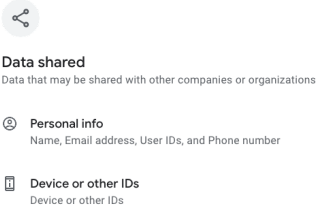
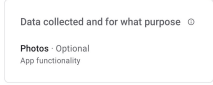
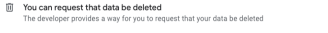
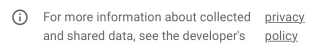
Question Theme	User Question Example	Apple Label	Google Label				
Data collected	<p>What kind of data does [APP] collect?</p> <p>Do you keep the data of mine and upload to your company?</p>	 <p>"Collect" refers to transmitting data off the device in a way that allows you and/or your third-party partners to access it for a period longer than what is necessary to service the transmitted request in real time.</p>	 <p>Data collected Data this app may collect</p> <ul style="list-style-type: none"> Personal info Name, Email address, User IDs, Address, and Phone number Financial info User payment info and Purchase history Photos and videos Photos App activity App interactions and In-app search history <p>Not in scope for data collection</p> <ul style="list-style-type: none"> The following use cases do not need to be disclosed as collected: <ul style="list-style-type: none"> On-device access/processing: User data accessed by your app that is only processed locally on the user's device and not sent off device does not need to be disclosed. 				
App security	Are you certified to be secure?	N/A	 <p>Security practices</p> <ul style="list-style-type: none"> Data is encrypted in transit Your data is transferred over a secure connection Independent security review This app has been independently validated against a global security standard. See details 				
Sharing	Is information shared with any third parties?	<table border="1"> <thead> <tr> <th>Purpose</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>Third-Party Advertising</td> <td>Such as displaying third-party ads in your app, or sharing data with entities who display third-party ads</td> </tr> </tbody> </table>	Purpose	Definition	Third-Party Advertising	Such as displaying third-party ads in your app, or sharing data with entities who display third-party ads	 <p>Data shared Data that may be shared with other companies or organizations</p> <ul style="list-style-type: none"> Personal info Name, Email address, User IDs, and Phone number Device or other IDs Device or other IDs
Purpose	Definition						
Third-Party Advertising	Such as displaying third-party ads in your app, or sharing data with entities who display third-party ads						
Selling	Which information, if any, does the app sell to third parties?	<table border="1"> <thead> <tr> <th>Purpose</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>Third-Party Advertising</td> <td>Such as displaying third-party ads in your app, or sharing data with entities who display third-party ads</td> </tr> </tbody> </table>	Purpose	Definition	Third-Party Advertising	Such as displaying third-party ads in your app, or sharing data with entities who display third-party ads	N/A
Purpose	Definition						
Third-Party Advertising	Such as displaying third-party ads in your app, or sharing data with entities who display third-party ads						
Purpose	How does this app utilize my data?	<p>App Functionality</p> <ul style="list-style-type: none"> Purchases Purchase History Other Purposes Contacts Contacts 	 <p>Data collected and for what purpose</p> <ul style="list-style-type: none"> Photos - Optional App functionality 				
Deletion	Can I remove all my data if I choose not to use this app again?	N/A	 <p>You can request that data be deleted The developer provides a way for you to request that your data be deleted</p>				
Privacy Policy	Where can I read your privacy policy?	The developer, Google LLC, indicated that the app's privacy practices may include handling of data as described below. For more information, see the developer's privacy policy .	 <p>For more information about collected and shared data, see the developer's privacy policy.</p>				

Table 6.2: Sample user questions and corresponding privacy label entry in the iOS and Google Play Stores. N/A means that the question does not have a label addressing it.

but also because they are generally app-agnostic. The remaining question themes are harder to answer in general.

Around 7.8% of questions pertain to specific app functionality. For instance, a question about

the DNA genetic testing app 23andMe reads, “If my genetic data turns out to be unexpected, can my family see it?” About 3.9% of questions address concerns about privacy violations or potential privacy risks, such as “Does having this on my device create a privacy concern?” Approximately 3% of questions are about privacy controls offered by the app, such as “Can I selectively block scripts on pages that I feel are invading my privacy?” when referring to the *Cake Web Browser* app. Another 2.8% of questions relate to how the app is protecting users’ privacy, such as “Can you guarantee my privacy while playing your game?” The questions under these themes are often app-specific and request more sophisticated answers.

6.4.5 A Comparative Summary of iOS and Google Labels

Our evaluation found that only around 40% of the question themes could be answered by the iOS or Google Play privacy labels, as shown in the top half of Table 6.3. Specifically, 43.2% of questions could be answered by Google Play labels, while 38.6% could be answered by iOS labels. The questions that could be answered by Google Play labels but not by iOS labels pertained to 1) additional data types (such as IP, calendar, and calls), 2) security-related questions, 3) whether the app data can be deleted, 4) optional tags for data that is not necessary for users to provide in order to use the app. In contrast, Apple’s labels provided more information related to data selling, which was not addressed by Google’s labels. Overall, Google’s labels addressed more questions than iOS labels.

Question Theme	# of Questions	Answered by Google	Answered by Apple	Answerable or not
Data collected	364	325	310	Could be answered or partially addressed by labels
App security	199	161	0	
Sharing	151	119	122	
Selling	141	0	125	
Purpose	95	87	87	
Deletion	31	18	0	
Privacy policy	9	7	7	
Permissions	140	0	0	Not answerable by labels
App-specific privacy	128	0	0	
Who has access	73	0	0	
Privacy risks	64	0	0	
Retention	56	0	0	
Privacy controls	49	0	0	
Retained method	40	0	0	
Account required	37	0	0	
External access	32	0	0	
Privacy protections	29	0	0	
Cookies policy	9	0	0	
Sum	1647	717 (43.6%)	651 (39.5%)	

Table 6.3: Questions can be answered by Google Play or iOS privacy labels

6.5 Discussions

6.5.1 Limitations

Our study investigates the crowd-sourced privacy questions in a public dataset. Even though the questions elicited are specific to the apps present in the dataset, the broad selection of apps and the questions, when analyzed as a whole, can to some extent reflect users' questions and concerns about apps. We cannot and are not making generalizable claims about our findings since our analysis mainly serves as an exploratory starting point. This chapter focuses on the scope of the label contents, and other issues, such as the usability problems or whether labels are reliable or factual, are outside of the scope of this chapter. Instead, we try to shed light on the potentially missing elements in label design and the unmatched mental models of users. As privacy researchers, we can only provide an upper limit when assessing whether labels address users' questions, assuming a complete and perfect understanding of the labels. Our results do not indicate whether users looking at the labels would find that their questions had been answered. As discussed in Chapter 5, the usability issues of the current labels could impede users' ability to really answer these questions, despite the answers being part of the privacy labels or explanations. Prior findings such as that users are confused by the technical definitions and the flattened label structure as mentioned in Section 5.3.3 will further reduce users' understanding of the labels. Further research is also needed to evaluate the effectiveness and efficacy of privacy labels in addressing the questions of users with varying levels of technical expertise, taking into account both the usability and scope of the labels.

6.5.2 Missing Key Information

Our analysis revealed that many question themes were not addressed by the iOS or Google privacy labels. This highlights the need for additional information to effectively address mobile users' privacy concerns or questions.

Recipient of Information

Participants wanted to learn who has access to their information and also whom their information is shared with or sold to. They also asked about access to their information by the government or the employees at the app company. This underscores the importance of disclosing the recipients of information, which aligns with the principles of Contextual Integrity [162] stating that it is imperative to disclose the recipient of the information flow. Therefore, privacy labels should include information about entities with whom user data is shared or sold. They should also address common questions, such as explaining to users whether the government or app company employees can access users' data. For example, messaging app Signal⁴ clearly states on its website that its end-to-end encryption keeps users' conversations secure and that no one, including the government or Signal employees, can read their messages or listen to their calls.

⁴<https://signal.org/en/>

App Permissions

Before the introduction of run-time Android permissions, the Google Play Store used to display a list of permissions that users needed to agree to before downloading an app. This information was no longer in the Play Store since Android 6. Google has changed its stance on including the permission list in the privacy label and currently does not include it [31]. Our analysis suggests that a good number of user questions pertain to what permissions the app needs or has access to, particularly about location, camera, microphone, and contacts permissions. Currently, users can only view the requested permissions for an app after installing it.

Users would also like to know the retention of their information, the availability of privacy controls, and the privacy risks of installing or using these apps.

6.5.3 Implicit Answers and Mismatching Mental Models

Our analysis reveals that a few question themes only have implicit answers, which might not be apparent to regular users.

Definition of Data Collection

Participants used terms like “store,” “save,” and “keep” when asking about data retention and whether their information is being stored. These questions often took the form of “Are you storing any of my information?” or “Do they collect my data and upload it?” This suggests that some users may not equate data collection with storing user data on servers, and it might be beneficial to emphasize that data collection is taken outside of the users’ devices or stored on servers.

Data Selling

One area of concern among participants was whether their data could be sold or not, with 8.6% of the questions related to this issue. This specific concern is addressed in the recent consumer privacy regulations in California (CCPA/CPRA), which require data controllers to disclose whether and with whom they may be sharing users’ data and to also provide users with privacy options to opt out of such selling. Even though users seem to want to know specifically about selling, neither the iOS nor the Google privacy labels readily use the term “selling,” making it difficult for users to find answers to these questions. For instance, Google requires disclosing what data is shared with third parties but does not require disclosing the purpose of the sharing or with whom the data is being shared. Apple’s privacy labels come closer to disclosing whether data is sold under the definition of CCPA/CPRA by introducing the concept of “tracking,” which focuses on sharing data with third parties in return for some type of consideration. However, Apple does not explicitly use the word “selling,” making it difficult for users to understand what is being disclosed and in particular whether their data is being sold [11].

What about Security

Google’s privacy labels already contain security information of an app, currently including whether “data is encrypted in transit” and “optional security review.” App developers can claim in the Google labels that for an app, “data is encrypted in transit: your data is transferred over a secure connection.” This, however, only seems to pertain to a very small number of questions—only 3 out of 199 security questions are about how secure data is during transit. Other aspects of security, such as whether user password is encrypted, are of more importance to users. Although the optional security review covers a wide range of topics, it might be unclear to users what such a security review entails. Furthermore, it is worth noting that this review is optional and not adopted by many apps.

6.5.4 Privacy Question Answering Functionality

Privacy labels are an important step towards the standardization of data practice disclosures. Prior work found that most users in an interview study reported that they like the concept of privacy labels in the Apple app store [256]. These labels also open the door for compliance analysis [120, 122].

Decreasing User Burden

Even though privacy labels are designed to help users quickly grasp the important data collection and usage practices without them having to read the text of privacy policies, current labels can already be overwhelming for some apps. For instance, the DoorDash iOS privacy label contains 106 entries of data types organized around 5 purposes and 2 sections. Concurrently, our results show that users have a rather diverse set of privacy questions, with more than half of these questions unlikely to be addressed in current labels. These two findings reveal a challenging tension, with labels appearing already overwhelming yet failing to address a substantial percentage of privacy questions typical users can be expected to have. The paper specifically identifies additional information that one might consider including in labels if one would like to have a better chance of answering people’s typical privacy questions in Section 6.5.2.

Given the amount of label information for each app and the large number of apps on each user’s phone, it is unrealistic to expect users to go through the privacy labels for each app on their phone. There is a need to reduce user burden and to help users quickly locate privacy information that they care about. Future research might want to explore the use of machine learning and natural language processing techniques to automatically extract and analyze standardized notices as a way of providing users with chatbot functionality to quickly answer their questions or refer them to parts of the labels pertaining to their questions.

App Specific Questions

Our analysis also reveals that many user questions pertain to querying about available privacy controls and app-specific privacy information, which fall beyond the scope of privacy labels. However, these concerns are still relevant to users. Recent development of advanced question-answering chatbots, trained on large language models, presents new research opportunities to

provide users with personalized answers to their privacy-related questions regarding specific apps. By doing so, users can make informed decisions without feeling overwhelmed by excessive privacy details. Utilizing these advanced chatbots to answer privacy questions can ease the burden on users to navigate complex privacy information. Further research is necessary to assess the feasibility, accuracy, and comprehensiveness of the answers provided by these chatbots.

6.6 Conclusion

We conducted a thematic analysis on a dataset of privacy questions that represent the questions mobile app users typically have about mobile apps. We evaluated whether these questions can be answered, whether partially or completely, by both iOS privacy labels and Google’s data safety labels. Our results indicate that an important percentage of people’s privacy questions are not answered or only partially addressed by today’s labels. Our analysis provides a roadmap for the possible refinement of existing labels if one were to try and improve the way in which they cover people’s most typical privacy questions. Because existing mobile app privacy labels are already rather complex and because people’s privacy questions about their mobile apps are also quite diverse, this research also suggests that privacy labels can only go so far in addressing the diverse privacy questions people have. Information in privacy labels may benefit from being organized differently (e.g., expandable tabular formats) and may need to be supplemented with additional functionality such as privacy question answering functionality (e.g., [189, 190]).

6.7 Summary of Main Contributions for Chapter 5 and 6

- We reported on the first in-depth interview study with 24 lay iPhone users to investigate their experiences, understanding, and perceptions of Apple’s privacy labels.
- We uncovered misunderstandings of and dissatisfaction with the iOS privacy labels that hinder their effectiveness, including confusing structure, unfamiliar terms, and disconnects between the labels and the permission settings and controls available to users.
- We identified areas where mobile app privacy labels might be improved and proposed suggestions to address shortcomings to make them more understandable, usable, and useful.
- We analyzed a corpus of privacy questions collected from mobile app users on Amazon Mechanical Turk to determine to what extent these mobile app labels could answer users’ privacy concerns and questions. Our analysis revealed that mobile app users have diverse privacy questions.
- Our analysis further indicates that an important percentage of people’s privacy questions are not answered or only partially addressed in today’s labels, though there are differences between iOS labels and Google Play labels.
- Our study highlights the need for further improvement in the usability and scope of privacy labels to help users effectively utilize them. Because of the diversity of people’s privacy questions, it also suggests that current privacy labels have their own limit in helping answer

people's privacy questions and that novel ways of presenting this information might be needed. The development of privacy question answering functionality offers another avenue for addressing people's privacy questions and could be used to supplement, or possibly one day just replace, privacy labels.

This study was published at PoPETS 2022 and USEC 2023 [253, 256].

Chapter 7

Conclusion

7.1 Summary of Contributions

The focus of this dissertation was to study the nuanced and complex nature of individuals' diverse privacy attitudes and concerns towards the deployment of several recent technologies (video analytics technologies, COVID vaccine certificates, and mobile app privacy labels) and analyze the implications of our findings on what it takes to effectively inform people about and give them adequate control over the collection and use of their information. Our work involved a succession of human subject studies, each intended to capture key elements of people's privacy attitudes across a number of different contexts. Our research provides further evidence of the important role played by key contextual attributes on people's privacy attitudes <insert reference to contextual integrity here and also to earlier work we have done such as: [122,123, 126] and also including the work with Mike Benisch and Patrick Kelley: <https://link.springer.com/article/10.1007/s00779-010-0346-0>>. It also documents the tension between the need to provide people with sufficient details and control over the collection and use of their information and the quickly unrealistic burden on users this might generate. We discuss ways in which this user burden could be mitigated with the introduction of standardized APIs and privacy assistants that could help users manage what otherwise be highly repetitive privacy decisions and the manual communication of these decisions to systems with which they interact. We further explored how machine learning techniques, including simple clustering techniques, can help organize people into groups of like minded-users and help build profiles that capture many of their privacy preferences. We discussed how such profiles could be edited by users to ensure they most closely align with their individual preferences and how the resulting individual profiles could then be used to communicate people's preferences such as opt-in/opt-out decisions to different systems, as users keep on encountering similar contexts such as similar video analytics deployments.

The main contributions of this thesis include:

- The development of detailed models of people's privacy expectations and preferences across a broad cross-section of realistic data collection and use practices associated with video analytics deployments, COVID-19 vaccination certificate deployments, and mobile app privacy notices. This includes the identification of key contextual parameters influencing

people's privacy expectations and preferences across these scenarios.

- Beyond the identification and modeling of privacy expectations and preferences that reflect attitudes of broad cross-sections of the population, this dissertation also explores differences in attitudes within the population and shows how a relatively small number of clusters of like-minded users can often be identified. These clusters in turn can be used to predict many aspects of people's privacy attitudes. In the context of attitudes towards the deployment of COVID certificates, they can help us gain a deeper understanding of how different segments of the population feel about different deployment scenarios and help inform public policy decisions about what is likely to be perceived as acceptable in some contexts and what may not be.
- These same types of clusters can be used to generate recommendations to help people more effectively and efficiently configure otherwise unmanageable collections of privacy decisions such as opt-in/opt-out decisions required by some regulations in the context of videoanalytics deployments. As such these models can contribute to reconciling the tension between user burden and the ability of people to effectively control the collection and use of their data..
- Our findings also shed new light on the unrealistic burden currently placed on users when it comes to managing their privacy across common mobile app and video analytics deployment scenarios. We argue that these findings provide strong support for the introduction of additional regulation that would require the availability of mechanisms, APIs, and protocols designed to reduce user burden. We proceed to detail some of such mechanisms in the context of IoT, mobile app, and web browsing scenarios.

7.2 Ongoing Work for the Label Study

In light of the usability issues identified in Chapter 5 (e.g., confusion about the label structure, lack of accessible definitions), we have undertaken to re-design iOS app privacy labels. Our goal is to see whether it is possible to make the labels clearer and easier to use. The modifications we are exploring include: 1) adopting a grid design with expandable lists, 2) using colors to differentiate between data linked to the user and data not linked to the user, and 3) making definitions easily accessible to users by adding clickable information icons with an "I" (for "Information"). We have created prototype labels that we are experimenting with, as depicted in Figure 7.1, for two apps, Candy Crush and Venmo. To evaluate the effectiveness of our prototype labels and gain insights into users' perceptions, we have developed a between-subjects interview protocol. Participants were randomly assigned to either view the original iOS privacy labels or the prototype labels and were then asked questions related to app data collection practices and their understanding of the labels, as well as for any suggestions for improvement. Currently, we are conducting pilot interviews to refine the protocol, which should enable us to make any necessary adjustments before conducting a larger study to gather more data and insights.

App Privacy

The developer, King, indicated that the app's privacy practices may include handling of data as described below. For more information, see the [developer's privacy policy](#).

To help you better understand the developer's responses, see [Privacy Definitions and Examples](#)

Data	Purpose	Functionality ⁽ⁱ⁾	Personalization ⁽ⁱ⁾	Analytics ⁽ⁱ⁾	Developer Ads ⁽ⁱ⁾	3rd-party Ads ⁽ⁱ⁾	Other ⁽ⁱ⁾	Tracking ⁽ⁱ⁾
▸ Identifiers		Dark Red	Dark Red	Dark Red	Dark Red	Dark Red	Grey	Dark Red
▸ Usage Data		Dark Red	Dark Red	Dark Red	Dark Red	Dark Red	Grey	Dark Red
▸ Diagnostics		Dark Red	Grey	Dark Red	Grey	Grey	Grey	Dark Red
▸ Contact Info		Dark Red	Grey	Grey	Dark Red	Dark Red	Grey	Dark Red
▸ Location		Dark Red	Dark Red	Dark Red	Dark Red	Dark Red	Grey	Dark Red
▸ User Content		Dark Red	Dark Red	Dark Red	Dark Red	Grey	Grey	Dark Red
▸ See More		Dark Red	Dark Red	Dark Red	Dark Red	Grey	Grey	Dark Red

What do the colors mean?

Dark Red	Any data could be collected and linked to your identity
Light Red	Some data could be collected and linked to your identity
Dark Blue	Any data could be collected but not linked to your identity
Light Blue	Some data could be collected but not linked to your identity
Grey	Not collected

(a) Candy Crush

App Privacy

The developer, Venmo, indicated that the app's privacy practices may include handling of data as described below. For more information, see the [developer's privacy policy](#).

To help you better understand the developer's responses, see [Privacy Definitions and Examples](#)

Data	Purpose	Functionality ⁽ⁱ⁾	Personalization ⁽ⁱ⁾	Analytics ⁽ⁱ⁾	Developer Ads ⁽ⁱ⁾	3rd-party Ads ⁽ⁱ⁾	Other ⁽ⁱ⁾	Tracking ⁽ⁱ⁾
▸ Identifiers		Dark Red	Dark Red	Dark Red	Dark Red	Dark Red	Grey	Dark Red
▸ Usage Data		Dark Red	Dark Red	Dark Red	Dark Red	Dark Red	Grey	Dark Red
▸ Diagnostics		Dark Blue	Grey	Dark Blue	Grey	Grey	Dark Blue	Dark Red
▸ Contact Info		Dark Red	Dark Red	Dark Red	Dark Red	Dark Red	Grey	Dark Red
▸ Location		Dark Red	Dark Red	Dark Red	Dark Red	Dark Red	Grey	Dark Red
▸ User Content		Dark Red	Dark Red	Dark Red	Dark Red	Dark Red	Dark Red	Dark Red
▸ See More		Dark Red	Dark Red	Dark Red	Dark Red	Dark Red	Dark Red	Dark Red

What do the colors mean?

Dark Red	Any data could be collected and linked to your identity
Light Red	Some data could be collected and linked to your identity
Dark Blue	Any data could be collected but not linked to your identity
Light Blue	Some data could be collected but not linked to your identity
Grey	Not collected

(b) Venmo

Figure 7.1: Examples of our privacy label prototypes

7.3 Challenges and Future Work

Although new technologies, such as the ones examined in this dissertation, have undoubtedly improved our lives by contributing to greater convenience, public health, and productivity, their increasing ubiquity and capabilities are also creating significant privacy risks. Unfortunately, the regulations in place to protect users' privacy often lag behind emerging technologies and lack adequate enforcement capabilities. It is often only after privacy violations have occurred that regulators take action, leaving individuals vulnerable in the meantime. It is increasingly important for companies and governments to take proactive steps to safeguard consumer privacy and protect users against the potential privacy risks introduced by new technologies. An important element in informing these proactive steps involves developing a deep and systematic understanding of people's privacy expectations and preferences across new deployment contexts. Through our research, we

have shown that individuals' privacy attitudes and concerns tend to be rather diverse and context-dependent, making it particularly challenging to design and implement usable privacy mechanisms such as effective notification and control mechanisms.

- In the case of Internet of Things (IoT) and video analytics developments, our results showing that individual's preferences for notification of video analytics deployments are quite diverse suggest that different people would select different setting configurations, with some individuals preferring to be systematically informed about each deployment and prompted to manually decide whether to opt in or out, while others preferring more selective notification settings and greater delegation of opt-in/opt-out decisions. Our findings can also inform the design of privacy assistants that help users manage privacy decisions related to the deployment of video analytics and other IoT technologies. Using privacy assistants, users could configure privacy profiles, which would capture their preferences across common taxonomies of contextual attributes. These profiles could be based on recommendations generated by machine learning algorithms, which the user could review and edit. As users repeatedly run into similar video analytics or IoT scenarios, their privacy assistants would rely on standardized APIs to recognize these scenarios, retrieve their user's privacy preferences for the corresponding scenario (e.g., opt-in/opt-out preferences), and communicate these preferences to the IoT systems on behalf of the user. This would relieve users from the tedious task of manually communicating the same preferences over and over again. Similarly, privacy assistants could be configured to notify their users only about those deployments users care to be notified about, in addition to possibly configuring any available opt-in/opt-out settings in accordance with their individual preferences. Based on our findings, it is evident that different users would likely select different configurations of their privacy assistants, enabling them to modulate the types of notifications they want to receive and the types of opt-in/opt-out decisions they might want to make manually.
- Our findings indicate that a one-size-fits-all approach to the deployment and implementation of COVID vaccination certificates may not be suitable for accommodating the varied and context-dependent privacy attitudes of individuals. Although our study indicates a general acceptance of vaccination certificates, the results of our vignette CI survey suggest that there is still a negative sentiment towards mandating VCs for accessing essential services and activities, places of worship, and apartment buildings. This highlights the importance of conducting surveys like ours by systematically sampling various contextual values to understand people's acceptance of different possible VC deployments and their implications. Our study highlights the importance of policymakers taking into account the diverse and context-dependent nature of privacy attitudes when designing policies related to the use of vaccination certificates.
- When examining whether notices are effective and practical in providing users with answers to their questions, a recent development that we investigated was privacy nutrition labels. While our study showed that privacy labels were generally a positive step towards increased transparency, the current form of these labels suffers from several usability issues. In addition, a significant number of user questions remain unanswered by these labels, which means that they are not providing users with the information they need. Moving forward, it is essential to

continue to improve and refine privacy labels to ensure they are more useful and usable for users. Future research could explore ways to address the limitations of these labels, such as increasing their scope and addressing a broader range of user questions. However, it is also important to note that privacy labels have limitations regarding the types of questions they can answer. For example, they may not be able to address app-specific privacy questions that users may have. As such, it is crucial to consider other approaches to privacy communication that can complement the use of privacy nutrition labels, such as privacy question answering functionality [189, 190].

The growth of technology has reached a point where users are overwhelmed by the number of privacy options available, making it challenging for them to manage privacy decisions manually. Therefore, it is essential to design and leverage technologies that can facilitate users' decision-making process. By carefully designing and leveraging such technologies, we can alleviate the user burden associated with configuring notice and choice functionality, thereby empowering them to regain control of the collection and use of their personal information in an increasingly digital world.

Appendix A

Understanding Privacy Expectations and Preferences of Video Analytics Technology

A.1 Scenarios

Table A.1: Scenario text shown to participants. *Controller* being a variable that would be instantiated with the name of the venue participants were visiting. Texts inside curly brackets display all retention options associated with this scenario. Texts inside square brackets can be inserted to specify sharing practice or the detection of whom people are with.

Purpose	Scenario Text
Generic Surveillance	Some places like <i>Controller</i> have started to deploy video surveillance cameras to deter crime . [This footage can be shared with law enforcement .] Assume that you are captured by such a camera, and {1)the raw footage is kept for 30 days , 2) it is unclear for how long the raw footage is kept }.
Petty Crime	Some places like <i>Controller</i> have started to deploy video surveillance cameras to deter crime . These cameras are equipped with software that can automatically detect and record petty crime (e.g. pickpocketing, car break-ins, breaking store windows). When a suspicious scene is believed to have been detected, it is recorded for further analysis (possibly including facial recognition) and kept for 30 days. Otherwise the data is immediately discarded . [This footage can be shared with law enforcement .] Assume that you are captured by such a camera.
Known Criminal	Some places like <i>Controller</i> have started to deploy video surveillance cameras with facial recognition software. This software can identify and track known shoplifters, criminals, and bad actors . Assume that <i>Controller</i> engages in this practice, and {1) the raw footage is discarded immediately with the analysis results being kept for 30 days 2) all the data (raw footage and analysis results) is kept for 30 days 3) the raw footage is kept for 30 days . Assume also that it is unclear for how long the analysis results are kept }.

Count people	Some places like <i>Controller</i> have started to deploy video surveillance cameras with anonymous face detection software. This software can estimate the number of customers in the facility in order to optimize operation , such as personnel allocation. Assume that <i>Controller</i> engages in this practice and {1) the raw footage is discarded immediately , and it is unclear for how long the analysis results are kept 2) the raw footage is kept for 30 days , and it is unclear for how long the analysis results are kept 3) it is unclear for how long all the data (raw footage and analysis results) is kept }.
Jump Line	Some places like <i>Controller</i> have started to deploy video surveillance cameras with facial recognition software. This software can identify patrons in line and push individualized offers to skip the wait-line for a fee . [This software can also record your presence and who you are with .] Assume that <i>Controller</i> engages in this practice and {1) all the data (raw footage and analysis results) is discarded immediately 2) the raw footage is discarded immediately with the analysis results being kept for 30 days 3) the raw footage is discarded immediately . Assume also that it is unclear for how long the analysis results are kept 4) the raw footage is kept for 30 days . Assume also that it is unclear for how long the analysis results are kept }.
Targeted (Anon) Ads	Some places like <i>Controller</i> have started to deploy video surveillance cameras with anonymous face detection software. This software can estimate customers' race and ethnicity in order to offer tailored deals and coupons . Assume that <i>Controller</i> engages in this practice and {1) all the data (raw footage and analysis results) is discarded immediately 2) the raw footage is discarded immediately with analysis results being kept for 30 days 3) all the data (raw footage and analysis results) is kept for 30 days 4) the raw footage is kept for 30 days . Assume also that it is unclear for how long the analysis results are kept }.
Targeted (IDed) Ads	Some places like <i>Controller</i> have started to deploy video surveillance cameras with facial recognition software. This software can match detected faces against individual customer profiles in order to offer tailored deals and coupons . [This software can record your presence and who you are with .] Assume that <i>Controller</i> engages in this practice and {1) the raw footage is discarded immediately with analysis results being kept for 30 days 2) all the data (raw footage and analysis results) is kept for 30 days 3) the raw footage is kept for 30 days . Assume also that it is unclear for how long the analysis results are kept }.

Sentiment (Anon)	Ads	Some places like <i>Controller</i> have started to deploy video surveillance cameras with anonymous face detection and emotion analysis software. This software can estimate customers' age, gender and ethnicity, and analyze their reactions to items displayed. This software is used to generate tailored deals and coupons for different demographic groups. Assume that <i>Controller</i> engages in this practice and {1) all the data (raw footage and analysis results) is discarded immediately 2) the raw footage is discarded immediately with analysis results being kept for 30 days 3) all the data (raw footage and analysis results) is kept for 30 days 4) the raw footage is kept for 30 days . Assume also that it is unclear for how long the analysis results are kept }.
Sentiment (IDed)	Ads	Some places like <i>Controller</i> have started to deploy video surveillance cameras with facial recognition and emotion analysis software. This software recognizes people, and analyzes their reactions to items displayed. Then the software matches detected faces against individual customer profiles to send tailored deals and coupons to their phones. [This software can record your presence and who you are with .] Assume that <i>Controller</i> engages in this practice and {1) the raw footage is discarded immediately with analysis results being kept for 30 days 2) all the data (raw footage and analysis results) is kept for 30 days 3) the raw footage is kept for 30 days . Assume also that it is unclear for how long the analysis results are kept }.
Rate Service		Some places like <i>Controller</i> have started to deploy video surveillance cameras with anonymous emotion analysis software. This software can gauge customer satisfaction with the service provided by its employees. They can use the results for employee evaluation and training purposes . Assume that <i>Controller</i> engages in this practice and {1) all the data (raw footage and analysis results) is kept for 30 days 2) the raw footage is kept for 30 days . Assume also that it is unclear for how long the analysis results are kept 3) it is unclear for how long all the data (raw footage and analysis results) is kept }.
Rate Engagement		Some places like <i>Controller</i> have started to deploy video surveillance cameras with facial recognition and emotion analysis software. This software can identify each patron, and measure their engagement at the facility. [This software can be used to record your presence and also identify who you are with .] Assume that <i>Controller</i> engages in this practice and {1) the raw footage is discarded immediately with the analysis results being kept for 30 days 2) all the data (raw footage and analysis results) is kept for 30 days 3) the raw footage is kept for 30 days . Assume also that it is unclear for how long the analysis results are kept 4) it is unclear for how long all the data (raw footage and analysis results) is kept }.

Face as ID	Some places like <i>Controller</i> have started to deploy video surveillance cameras with facial recognition software. This software can identify faces to replace ID cards . [This software can record your presence and who you are with .] Assume that Controller engages in this practice, and {1) the raw footage is discarded immediately . Assume also that it is unclear for how long the analysis results are kept 2) all the data (raw footage and analysis results) is kept for 30 days 3) the raw footage is kept for 30 days . Assume also that it is unclear for how long the analysis results are kept 4) it is unclear for how long all the data (raw footage and analysis results) is kept }.
Track Attendance	Some companies have started to deploy video surveillance cameras with facial recognition software. This software can track the work time attendance of its employees . [This software can record your presence and who you are with .] Assume Controller engages in this practice, and {1) the raw footage is discarded immediately . Assume also that it is unclear for how long the analysis results are kept 2) all the data (raw footage and analysis results) is kept for 30 days 3) the raw footage is kept for 30 days . Assume also that it is unclear for how long the analysis results are kept 4) it is unclear for how long all the data (raw footage and analysis results) is kept }.
Word Productivity	Some companies have started to deploy video surveillance cameras with emotion analysis and facial recognition software. This software can detect the mood of its employees and predict their productivity . [This software can record your presence and who you are with .] Assume that your workplace engages in this practice, and {1) all the data (raw footage and analysis results) is kept for 30 days 2) the raw footage is kept for 30 days . Assume also that it is unclear for how long the analysis results are kept 3) it is unclear for how long all the data (raw footage and analysis results) is kept }.
Health Predictions	Some eatery chains like <i>Controller</i> have started to deploy video surveillance cameras with emotion analysis and facial recognition software. This software can detect your mood and record data about your orders [and who you are with]. [This information can be shared with health insurance providers . The health insurance providers could use such data to estimate your likelihood of developing depression, diabetes, and obesity , which can impact your health insurance premium .] Assume that <i>Controller</i> engages in this practice, and {1) all the data (raw footage and analysis results) is kept for 30 days 2) the raw footage is kept for 30 days . Assume also that it is unclear for how long the analysis results are kept 3) it is unclear for how long all the data (raw footage and analysis results) is kept }.

Medical Predictions	Some medical facilities have started to deploy video surveillance cameras with emotion analysis and facial recognition software. This software can automatically detect some physical and mental health problems . [This information can be shared with health insurance providers , and impact your health insurance premium .] Assume that <i>Controller</i> engages in this practice, and { 1) all the data (raw footage and analysis results) is kept for 30 days 2) the raw footage is kept for 30 days . Assume also that it is unclear for how long the analysis results are kept 3) it is unclear for how long all the data (raw footage and analysis results) is kept }.
---------------------	--

Table A.1: Full Scenarios

A.2 Evening Review

[Show a map, timestamp and scenario for each notification]

- We asked: How surprised would you be about [PLACE] engaging in this data practice? At the time, you indicated that you would find this _____. Why?
- We asked: How comfortable would you feel about [PLACE] engaging in this data practice? At the time, you indicated that you would find this _____. Why?
- We asked: How would you want to be notified as you enter [PLACE]? At the time, you indicated that you _____. Why?
- If you had the choice, would you allow or deny this data practice?
- Based on the data practice description above, do you believe the footage in which you appear could be made available to third parties for analysis with facial recognition?
- Please indicate how much you agree or disagree with each of the following statements.
 - I feel that I benefit from this data practice
 - I feel that [PLACE] benefits from this data practice
 - I feel that the data practice enhances public safety
- How would you feel about the raw footage being shared with the following entities?

A.3 Post Study Survey

- What is the first thing that comes to your mind when you think about facial recognition technology?
- In what context(s) do you find the use of facial recognition technology to be particularly beneficial? (Enter up to 5 types of contexts)
- In what context(s) do you find the use of facial recognition technology to be particularly concerning? (Enter up to 5 types of contexts)

- Do you feel that you have a general understanding of where this type of technology is likely to be used and why?
- Please rate your comfort level when visiting stores and other locations that use facial recognition technology.
- How likely would you be to intentionally avoid stores that use facial recognition technology?
- Has your level of concern about facial recognition technology changed over the course of the study?
- 10 IUIPC Questions
- Show scenarios: Petty Crime/Sentiment Ads(IDed)/Health Predictions
 - Within what timeframe, do you believe this data practice will be commonplace?
 - Would you like to be notified about this data practice?
 - What sensitive information do you think could be inferred from this data collection practice?
 - How concerned would you be about this sensitive information being inferred? Why?
 - How likely would you be to avoid visiting those places following the introduction of this data practice?
 - What do you think is a reasonable timeframe for those places to retain the footage they capture of you?
 - In what manner would you like to receive notification about those places' use of this data practice?

A.4 Interview Scripts

- Introduction: Thank you for agreeing to this interview. My name is _____. I will be audio-recording our session. How are you doing today? Just to fresh your memory. You started the study around [DATE], and finished the study around [DATE]. For this interview, we will be asking you some additional questions and clarifications about your experience during this study.
- Where did you find about our study?
- When did you download the app?
- Did you find participating in this study to be demanding?
- On average, how much time would you say you spent answering our questions each day?
- Were there days when you didn't receive any prompts?
- On the whole, do you feel that we covered most of the interesting places you went to during the course of the study, or would you say we missed some interesting places? If so, which interesting places did we miss? Would you expect cameras to be present at these places and what do you think these cameras could be doing?
- While going through the evening reviews, did you ever feel that you wished you could modify some of the answers you provided during the day? If so, can you specifically remember some of the scenarios and in which way you would have liked to modify your answers (e.g., less

surprised or more surprised, less comfortable or more comfortable?)

- [CHECK DATA] For scenarios where we only collected your answers in the evening, because you didn't have time to answer them when the scenario occurred. Do you believe that you might have given different answers if you had responded at the time we first prompted you? If so, how different would your answer have been and why?
- [SHOW INSTANCES] Do you remember when you did not answer those scenarios on site / in-situ, why you could not answer them, and what you were you doing at the time?
- If you remember, each scenario came with two questions designed to check whether you had carefully read the description of the scenario. Do you remember those?
- Did you find that answering these questions could easily be defeated, or did you actually have to carefully read the scenarios to answer the questions? Feel free to tell us that the questions were easy to guess without reading the scenarios. We are trying to understand to what extent these questions help, or to what extent they are just not terribly useful.
- How often did you think that the scenarios we described matched actual video collection practices at the places you were visiting?
- Did you actually look for cameras, or start paying more attention to cameras?
- Have you discussed the study or scenarios with others?
- On the whole, do you feel that you have grown more concerned or less concerned about the types of video analytics scenarios used in our study? Or would you have you remain equally concerned or unconcerned?
- If you remembered, there are a lot of scenarios you encountered as part of this study, were there scenarios that you found particularly surprising? Were there some scenarios that you found particular concerning? Or would you say that all these scenarios are to be expected and do not feel particularly concerning?
- Do you feel that, if you were to retake the study and be presented with the same scenarios, most of your answers would be the same? If some of your answers are likely to be different, could you identify some of the scenarios for which you would likely have different answers?
- Questions/Clarifications related to the interviewee's post-study and evening answers (different case by case)

Appendix B

A Contextual Integrity Analysis of Vaccination Certificates

B.1 Full Survey Text

Consent Form

- I am age 18 or older.
- I have read and understand the information above.
- I want to participate in this research and continue with the survey.

Introduction

With ongoing COVID-19 vaccination efforts, governments and other organizations around the world have proposed the use of “vaccination certificates” as a way to verify that a person has been vaccinated against the coronavirus, received a negative test or has recovered from the virus. Some vaccination certificates are already in use today. Researchers at Carnegie Mellon University are conducting a study to understand people’s opinions and perceptions of these vaccination certificate proposals. Please answer the survey honestly. There are no right or wrong answers to any of the questions.

Fist-hand Information Sharing: Vaccination Passport Vignettes

- Pre-COVID, how often did you visit [place]?
- Assume that you have a vaccination certificate similar to the one below.
- Template: [Recipient] ask [Sender] to show [Subject+Attribute] to [Transmission Principle]. Would such a practice be acceptable?
- Example: [Gyms] ask [members] to show [their vaccination certificates] to [gain indoor access]. Would such a practice be acceptable? Please explain.



The entity scanning this QR code can view your full name, date of birth, and your Vaccination Certificate Status (valid, invalid, expired)

Figure B.1: An example vaccination certificate shown to survey participants.

- If such a certificate were to be required to [gain indoor access], how much more likely would you be to go to [gyms] over the next 6 months? Please explain.

Fist-hand Information Sharing: Vaccination Mandate Vignettes

- Passengers are asked to show their vaccination certificates to airline companies to take an international flight. Is this acceptable?
- Foreign travelers are asked to show their vaccination certificates to customs and border controls to enter the United States. Is this acceptable?
- Us nationals are asked to show their vaccination certificates to customs and border controls to enter a foreign country. Is this acceptable?
- Teachers are asked to show their vaccination certificates to schools (K-12 and higher education)

to return to in-person learning. Is this acceptable?

- Students are asked to show their vaccination certificates to schools (K-12 and higher education) to return to in-person learning. Is this acceptable?
- Job applicants are asked to their show vaccination certificates to employers to be considered for a job. Is this acceptable?
- Job applicants are asked to show their vaccination certificates to employers to apply for jobs in hospitals. Is this acceptable?
- Job applicants are asked to show their vaccination certificates to employers to apply for or retain jobs in assisted living facilities. Is this acceptable?
- Potential renters asked to show their vaccination certificates to building management to rent an apartment. Is this acceptable?
- Word count: Please select the answer choice with the largest number of words in the list below.

VC Information Re-sharing Vignettes

- Template: Would it be acceptable for [Sender] to share [Subject Attribute] with the following entities for [Transmission Principle]?
Example: Would it be acceptable for [recreational services or facilities (e.g., bars, gyms, salons)] to share [information on a person's vaccination certificate] with the following entities [for public health purposes such as contact tracking]?

Vaccination Certificate Questions

- Do you agree or disagree with the following statements?
The government (federal or state)
 - should promote vaccination against COVID-19.
 - has no right to impose vaccination certificates.
 - should issue vaccination certificates and require them to be used in different contexts.
- Which entity do you consider trustworthy to develop a vaccination certificate? Please select all that apply. Please explain.
- Would you prefer to have a single certificate issued by the federal government and recognized by everyone or different certificates issued by different organizations from which you can choose?

COVID Related Questions

- Have you been vaccinated against COVID-19?

- Have you contracted COVID-19?
- Do you personally know anyone who got seriously ill due to COVID-19?
- If vaccination certificates were to be used, would you be more or less likely to get vaccinated?

Demographics

- What is your age?
- What is your gender?
- What is the highest level of education you have completed?
- What was your total household income before taxes during the past 12 months?
- What is your marital status?
- Which of the following best describes your primary occupation?
- Please specify your ethnicity.
- In general, would you describe your political views as ___?
- Have you ever held a job in assisted living facilities or hospitals?
- Do you use a smartphone?
- In which state do you currently reside [drop-down]?
- Which category best describes where you live?
- Have you used the following tools in the past year? Please select all that apply.
- Not including this survey, approximately how many surveys related to privacy or security have you completed in the past year?
- Anything else you'd like to say about the situation and/or your concerns?

B.2 Full Survey Text

Demographic Questions

- Respondent age (in years) — Numeric value
- Gender as provided by vendor (M/F only) [Text items]
 - Female
 - Male
- How would you describe your gender identity? [Text items]
 - Female
 - Male
 - Genderqueer

- Other
- Race as provided by vendor (select one) [Text items]
 - African American
 - Asian American
 - Hispanic
 - Native American
 - Pacific Islander
 - White
 - Other
- What racial or ethnic group best describes you? (Please select all that apply)
 - 1 = Asian American
 - 0 = Not Asian American
- What racial or ethnic group best describes you? (Please select all that apply)
 - 1 = African American
 - 0 = Not African American
- What racial or ethnic group best describes you? (Please select all that apply)
 - 1 = Hispanic
 - 0 = Not Hispanic
- Education level (numerical)
 - 1 = Some High School or Less
 - 2 = High School Graduate
 - 3 = Some College
 - 4 = College Degree
 - 5 = Graduate Degree
- Education level (5 categories) [Text items]
 - Some High School or Less
 - High School Graduate
 - Some College
 - College Degree
 - Graduate Degree
- Raw household income as provided by vendor — Numeric value
- What was the total combined income of your household for the past year?
 - 1 = Under 10K

- 2 = 10k to under 15k
- 3 = 15K to under 25K
- 4 = 25K to under 35K
- 5 = 35K to under 50K
- 6 = 50K to under 75K
- 7 = 75K to under 100K
- 8 = 100K to under 150K
- 9 = 150K to under 200K
- 10 = 200K and over
- Which of the following best describes your current employment status? [Text items]
 - Full-time
 - Part-time
 - Self-employed
 - Unemployed
 - Home-maker
 - Student
 - Retired
 - Gig/Contract
- Relationship status from vendor [Text items]
 - Single
 - Engaged
 - Living with partner
 - Married
 - Divorced
 - Widowed
- Parent of children under 18
 - 0 = No kids
 - 1 = Has kids
- Political party
 - Republican
 - Democrat
 - Independent
 - Other

- Political affiliation
 - 1 = Strong Republican
 - 2 = Republican
 - 3 = Leaning Republican
 - 4 = Independent
 - 5 = Leaning Democrat
 - 6 = Democrat
 - 7 = Strong Democrat
- In general, do you think of yourself as...
 - 1 = Extremely liberal
 - 2 = Liberal
 - 3 = Slightly liberal
 - 4 = Moderate, middle of the road
 - 5 = Slightly conservative
 - 6 = Conservative
 - 7 = Extremely conservative
- Do you think of yourself as closer to the...
 - 1 = Republican Party
 - 2 = Democratic Party
 - 3 = Neither
- Do you consider yourself to be a...
 - 1 = Strong Democrat
 - 2 = Not very strong Democrat
- Do you consider yourself to be a...
 - 1 = Strong Republican
 - 2 = Not very strong Republican
- Respondent ZIP code — Text entry
- Respondent county (from ZIP) — Text entry
- State name—Text entry
- US State numerical census code 1 = AL 2 = AK 4 = AZ 5 = AR 6 = CA 8 = CO 9 = CT 10 = DE 11 = DC 12 = FL 13 = GA 15 = HI 16 = ID 17 = IL 18 = IN 19 = IA 20 = KS 21 = KY 22 = LA 23 = ME 24 = MD 25 = MA 26 = MI 27 = MN 28 = MS 29 = MO 30 = MT 31 = NE 32 = NV 33 = NH 34 = NJ 35 = NM 36 = NY 37 = NC 38 = ND 39 = OH 40 = OK 41 = OR 42 = PA 44 = RI 45 = SC 46 = SD 47 = TN 48 = TX 49 = UT 50 = VT 51 = VA 53 = WA 54

= WV 55 = WI 56 = WY 60 = AS 66 = GU 69 = MP 72 = PR 78 = VI

- Two-letter state code—Text entry
- USDA Rural-urban Continuum Codes, 2013
 - 1 = Metro - Counties in metro areas of 1 million population or more
 - 2 = Metro - Counties in metro areas of 250,000 to 1 million population
 - 3 = Metro - Counties in metro areas of fewer than 250,000 population
 - 4 = Nonmetro - Urban population of 20,000 or more, adjacent to a metro area
 - 5 = Nonmetro - Urban population of 20,000 or more, not adjacent to a metro area
 - 6 = Nonmetro - Urban population of 2,500 to 19,999, adjacent to a metro area
 - 7 = Nonmetro - Urban population of 2,500 to 19,999, not adjacent to a metro area
 - 8 = Nonmetro - Completely rural or less than 2,500 urban population, adjacent to a metro area
 - 9 = Nonmetro - Completely rural or less than 2,500 urban population, not adjacent to a metro area
- NCHS Urban-Rural Classification 2013
 - 1 = Large central metro
 - 2 = Large fringe metro
 - 3 = Medium metro
 - 4 = Small metro
 - 5 = Micropolitan
 - 6 = Non-Core
- Have you ever been diagnosed with coronavirus (COVID-19)?
 - 1 = Yes, I was diagnosed by a medical professional
 - 2 = No, I was not diagnosed but I think I may have it now
 - 3 = No, I was not diagnosed but I think I had it previously and recovered
 - 4 = No, I was not diagnosed and I do not think I ever had it
 - 5 = I am not sure
- Have you been tested for coronavirus (COVID-19)?
 - 1 = Yes, and I tested positive for COVID-19 at least once
 - 2 = Yes, and I tested negative for COVID-19 every time
 - 3 = No, I wanted to but was not able to get a test
 - 4 = No, I never tried to get tested
- How severe is/was the illness?
 - 4 = Very severe

- 3 = Somewhat severe
- 2 = Not too severe
- 1 = Not at all severe
- Have you fully recovered from COVID-19?
 - 1 = Yes, I have fully recovered
 - 2 = No, I still have some symptoms
 - 3 = No, I do not have any symptoms but my test is still positive
- Have you received a COVID-19 vaccine?
 - 1 = Yes, one dose
 - 2 = Yes, two doses
 - 3 = No
 - 4 = Yes, three doses
 - 5 = Yes, four or more doses
- How likely would you be to get vaccinated against COVID-19 in the future?
 - 5 = Extremely likely
 - 4 = Somewhat likely
 - 3 = Neither likely nor unlikely
 - 2 = Somewhat unlikely
 - 1 = Extremely unlikely
- Have you received a COVID-19 booster shot (that is, an additional vaccine dose to increase immunity that may have waned over time)?
 - 1 = Yes, one booster shot
 - 2 = Yes, two booster shots
 - 3 = Yes, three or more booster shots
 - 0 = No
- How likely would you be to get a COVID-19 vaccine booster shot in the future?
 - 5 = Extremely likely
 - 4 = Somewhat likely
 - 3 = Neither likely nor unlikely
 - 2 = Somewhat unlikely
 - 1 = Extremely unlikely
- Do you plan to get a COVID-19 booster...
 - 1 = As soon as possible
 - 2 = After at least some people I know have already received it

- 3 = After most people I know have already received it
- 4 = I would not get the COVID-19 booster

Vaccination Certificate Questions

With ongoing global COVID-19 vaccination efforts, governments and organizations have proposed and adopted the use of “vaccination certificates” or other equivalent certification, as a proof of COVID vaccination or recent recovery from the virus. Please answer the following survey questions regarding the use of vaccination certificates in various scenarios. Please assume that “vaccination certificate” refers to a proof (whether digital or not) that someone has been vaccinated or an optional proof that a person has recently recovered from a COVID virus infection.

Would the following practices be acceptable or unacceptable? 5 = Acceptable 4 = Somewhat acceptable 3 = Neither acceptable nor unacceptable 2 = Somewhat unacceptable 1 = Unacceptable

- Government buildings (e.g., DMVs, courthouses) ask visitors to show their vaccination certificates to gain indoor access.
- Assisted living facilities ask visitors to show their vaccination certificates to gain indoor access.
- Places of worship ask visitors to show their vaccination certificates to gain indoor access.
- Apartment building management asks visitors to show their vaccination certificates to gain indoor access.
- Large event organizers ask attendees to show their vaccination certificates to gain indoor access.
- Gyms ask members to show their vaccination certificates to gain indoor access.
- Restaurants and cafes ask customers to show their vaccination certificates to gain indoor access.
- Stores and supermarkets ask customers to show their vaccination certificates to gain indoor access.
- Entertainment venues (e.g., movie theaters, museums) ask customers to show their vaccination certificates to gain indoor access.
- Personal care businesses (e.g., nail salons, barber shops) ask customers to show their vaccination certificates to gain indoor access.
- Hotels and short-term rentals (e.g., Airbnb) ask customers to show their vaccination certificates to gain indoor access.
- Cruise companies ask passengers to show their vaccination certificates to board.
- Public transportation operators ask passengers to show their vaccination certificates to board.
- Taxi drivers or rideshare drivers (e.g., Uber drivers) ask passengers to show their vaccination certificates to board.
- Long-distance bus or train companies (e.g., MegaBus, Amtrak) ask passengers to show their

vaccination certificates to board.

- Airline companies ask passengers to show their vaccination certificates to board an international flight.
- Customs and border controls ask travelers to show their vaccination certificates to enter a country.
- Schools (K–12 and higher education) ask teachers to show their vaccination certificates to attend school.
- Employers ask job applicants to show their vaccination certificates to apply for a job.
- Employers ask job applicants to show their vaccination certificates to apply for a job in healthcare (e.g., assisted living facilities, hospitals).
- Apartment building management asks renters to show their vaccination certificates to be considered for an apartment.
- Attention question: In this question, please select the second to last option, namely “somewhat unacceptable.”

Appendix C

Usability of iOS App Privacy Labels

C.1 Screening Questionnaire

- You can check your iOS version by going to Settings>General>About, and looking at “Software Version”. Please type in the version number exactly as it appears in the “Software Version” section. [free-text]
- I am at least 18 years old, reside in the US, and am a regular user of an iPhone with iOS 14 or above.
- Please select the most applicable answer. I downloaded one or more apps from the app store —.
 - In the past week
 - In the past month
 - In the past 3 months
 - More than 3 months ago [disqualified]
- Roughly how many new apps have you downloaded from the app store yourself over the past 3 months?
 - None [disqualified]
 - Somewhere between 1 and 10
 - Likely more than 10
- Have you ever done any of the following in the past? Please select all that apply.
 - Uninstalled or stopped using an app or service because of the types of data the app collects about you or how that data is used
 - Reviewed an app’s privacy settings, namely what data it requests access to
 - Read (partially or fully) an app’s privacy policy or end-user license agreement

- Used a VPN or Tor for non-work-related reasons on your phone or other device
- Decided not to download an app after looking at its privacy information in the app store
- None of the above
- What is your age? [free-text]
- What is your gender?
 - Male
 - Female
 - Non-binary
 - Prefer to self disclose
- What is your occupation? [free-text]
- Please enter your email. Your email will only be used to contact you to set up a time for the study and to pay you, if you are selected to participate in this study. If you are not selected, your responses to the survey (including your email address) will be deleted within 3 days of the completion of recruitment for the study. Your email address will not be shared with anyone and will be stored separately from your other study data.

C.2 Interview Scripts

- Introduction: Thank you for meeting with me today. This interview is being conducted for research at Carnegie Mellon University to better understand how people interact with mobile apps in the Apple App Store. We will ask you to answer some questions and view some information in the App Store. This session should take no more than 1 hour to complete, and will be recorded via Zoom. Upon completion of the study, you will receive \$25 in the form of an Amazon Gift Card that will be sent to you via email. You will be asked to share your iPhone's screen via Zoom at some point during the session to enable us to follow what you are doing on your phone as you visit the App Store.
- Please answer our questions truthfully and as thoroughly as possible. If in doubt, feel free to ask me for clarification at any point during the interview. I want to emphasize that there are no right or wrong answers. Our goal is simply to understand your opinions and thought processes. You may stop the interview at any point, or choose to not answer a question, or take a break if you wish. Please do not reveal any private or personally-identifiable information about yourself or others during the interview. If you accidentally reveal any personal information, please let me know so that I can remove it from the recording. Do you have any questions at this time?
- Part 1: General Questions about App Usage
 - For how long have you been using an iPhone? (Prompt: Any particular reason why you chose an iPhone?)
 - To the best of your knowledge, approximately how many apps do you have on your

iPhone? (estimates are expected) [After getting the estimate, give them instructions to look up the actual number Settings>General>About>Applications]

- When was the last time that you downloaded a new app on your phone?
- Could you describe a recent experience when you decided to download an app on your phone, starting from how you discovered the new app all the way to what happened when you used it for the first time? (to the extent they went all the way - some people can stop halfway and decide not to set up an account or may even change their mind and remove the app)
- What are some of the typical factors that influence which apps you download on your phone? (Prompt: app reviews, brand, ratings, security, ranking, data privacy, your friends) Have you ever compared different apps before deciding which one to download (Prompt: what types of things have you compared?)? Was data privacy ever a reason that you chose or did not choose an app?
- Part 2 : Information Seeking
 - Have you ever wondered what information apps collect about you?
 - [If they say they have] How would you go about finding out what information an app collects about you and what the app does with the information?
 - [follow-up] Have you ever actually done that?
 - [If they say they have not] If you were to look for this, how could you possibly find out what information an app collects about you and what the app does with that information? Prompt: media? friends/family? experts? privacy policies/EULA? permission settings? Other? Do you think those sources are reliable? App store? Have you ever looked
- Part 3: Label Comprehension — 2 Scenarios
 - [Instructions] This section requires you to share your screen with us via zoom. Which device are you on, your iPhone or your desktop/laptop? Please open the App Store app on your phone before you start sharing your screen. Please silence your notifications and remove anything confidential from your screen. [If participant is not on iOS 15, show instructions] Please enable Do Not Disturb on your device to prevent unexpected notifications by going to “Settings”>“Focus” and then “Do Not Disturb” and turn on the top toggle. [Remind if on their phone using the following sentence] Please note that we will be able see snapshots of your apps when you switch from Zoom to the App Store app. So please make sure that there is nothing sensitive displayed on your screen before you start sharing your screen.
 - [Show instructions on screen if needed] Great! Now that you are sharing your screen.
 - Could you search for the [Doordash or Chipotle] app? Are you familiar with this app? Could you describe what this app does? Have you used [APP] before? [If yes] how often do you use it?
 - Please scroll down to the “App Privacy” section. Do you remember ever seeing or reading an “App Privacy” section like this one before? (follow-up: if yes, ask about

their experience with privacy labels; when did you last see one? For which app? What was the context? Did you find the information useful? Did you end up downloading the app? If not, why not? Do you typically look for this information before downloading a new app?)

- Please take some time to read this “App Privacy” section. What do you think of the section you are seeing? What do you think this section is for? (impression testing)
- Please click the “See Details” at the top right corner and take some time to read this as well. Let’s go through the app privacy information section you just looked at systematically. Please answer the following questions based on what you see in this “App Privacy” section:
 - Starting at the top where it says "Data used to Track you", what do you think “data used to track you” means?
 - Do you know what (online) tracking is?
 - What do you think of it? Do you find it to be a useful practice or are you possibly concerned about it? Please explain.
 - In this section there is information about “identifiers.” Do you know what an identifier is? If yes, what are (other) examples of identifiers?
 - What do you think of the fact that this app may use your identifiers to track you across apps and websites owned by other companies? Do you find it to be a useful practice or are you possibly concerned about it? Please explain.
 - In this section there is information about “product interaction” under “usage data.” Do you know what product interaction entails? How do you think Doordash tracks you via product interaction(s)?
 - [If the user is looking at Doordash]
 - What do you think of it? Do you find it to be a useful practice or are you possibly concerned about it? Please explain.
 - [Tracking] Do you think Doordash is allowed to share your location with a third-party company that would combine your location obtained from Doordash and location data from other apps and websites to build a history of your whereabouts?
 - What do you think of it? Do you find it to be a useful practice or are you possibly concerned about it? Please explain.
 - [Tracking] Do you think Doordash uses data collected from other companies (including websites, apps, and offline services) to decide what ads to show you? If so, what data do you think the app uses? How can you tell? What do you think of it? Do you find it to be a useful practice or are you possibly concerned about it? Please explain.
 - If you scroll down a bit you should see a heading for “data linked to you,” previously there was another heading for “data used to track you”. What is the difference between “data used to track you” and “data linked to you”?

- The next heading is Third-party advertising. If you scroll down a bit you should see another heading Developer’s Advertising or Marketing. What do you think those headings denote? Do you think that these headings refer to different practices and, if so, what are the differences?
- What do you think “third-party advertising” means?
- Do you know what targeted advertising is?
- Do you find it to be a useful practice or are you possibly concerned about it? Please explain. (There is no right or wrong answer. We’re just curious to understand how you feel about these practices)
- Do you know what Developer’s Advertising or Marketing means? What do you think is the difference, if there is, between “third-party advertising” and “developer’s advertising or marketing”?
- Below that you will see Browsing History. What do you think “browsing history” covers?
- Do you think browsing history includes content the user has viewed that is not part of the app, such as websites?
- Are you concerned or not concerned about this data being collected?
- If you keep scrolling you will see “other data”, what do you think “other data” include?
- If you scroll down you will see the Analytics heading. What do you think your data being used for “Analytics” purposes means?
- What do you think of it? Do you find it to be a useful practice or are you possibly concerned about it? Please explain.
- [Analytics] Do you think Doordash uses data to understand or analyze your behavior (e.g., to develop new features, to measure audience characteristics)? If so, what data do you think the app uses? What do you think of it?/How do you feel about it? Do you find it to be a useful practice or are you possibly concerned about it? Please explain.
- If you scroll down you will see the Product Personalization heading. Do you know what that is? Do you find it to be a useful practice or are you possibly concerned about it? Please explain.
- If you keep scrolling down you will see the heading “app functionality,” What do you think your data being used for “App Functionality” purposes means? What do you think of all the other purposes beyond app functionality?
- What do you think of it? Do you find it to be a useful practice or are you possibly concerned about it? Please explain.
- If you keep scrolling you will see “other purposes,” what do you think about your contacts being used for “other purposes”? Do you find it to be a useful practice or are you possibly concerned about it? Please explain.
- Is there any information here that you do not understand?

- Is there any information that is not present but you would like to know about?
- [For participants assigned to the Chipotle]
- Let's go through the app privacy information section you just looked at systematically. Please answer the following questions based on what you see in this "App Privacy" section: Starting at the top where it says "Data not linked to you", what do you think "data not linked to you" means? The next heading is Analytics. If you scroll down a bit you should see another heading App Functionality. What do you think those headings denote?
- [Analytics] Do you think Chipotle uses data to understand or analyze your behavior (e.g., to develop new features, to measure audience characteristics)? If so, what data do you think the app uses?
- How do you think Chipotle collects "contact info" such as "Name" and "Email address" without the data being linked to you?
- Do you think Chipotle collects any data that is linked to you? For example, when you place an order?
- Do you think this App Privacy section includes all the data and all the usage of your data that Chipotle collects about you? Please explain.
- Is there any information here that you do not understand?
- Is there any information that is not present but you would like to know about?
- Perceptions of Privacy Labels
 - In general, how do you feel about the information provided here (i.e., the information under the "App Privacy" section)?
 - Do you find this type of information useful? Why/Why not? Do you find it easy to understand? Do you find it well organized?
 - Do you feel you understand most of the information provided in that section?
 - Did you learn about things you didn't know or you pretty much knew everything in that section? Or somewhere in-between?
 - What do you like or dislike about that section? How can this section be improved?
 - Do you think you might pay attention to this information later or you will probably not be looking at this?
 - [If yes] Would this information influence your decision to download the app one way or another?
 - [If no] Why do you think you will not be looking at this?
 - [if participant saw labels before] After downloading an app, did you ever find yourself going back to this section in the app store?
 - Under what circumstances? For what reason? (Prompt: to answer privacy questions you might have had)

- If you had to guess, who do you think provided the information under “App Privacy”?
- [Follow-up] Do you believe that this information is provided directly by Apple or vetted by Apple? Or by the app developers? Why/why not?
- If it’s the latter, do you believe that it has been reviewed by Apple?
- [Follow-up] Do you believe it is done manually or the result of some automated processing?
- [if participant saw labels before] If you ever looked at this type of information in the past, did you ever have that question in mind? Who did you assume provided the information in the labels? Did you think the information could be trusted?
- Do you think the information provided in this section is reliable? How likely are you to trust this information? Why/why not?

[Remind participants to stop sharing their screen]

- Part 5: General Privacy Concerns / Behaviors
 - Have you ever read the privacy policy (partially or fully) of a mobile app? What do you think of them? Prompt: When you read a privacy policy, what do you typically do? How much do you typically understand the privacy policies you read?
 - Have you ever regretted downloading or using an app because of data privacy issues? [Follow-up]: Did you take any further actions because of these regrets such as changing your privacy settings, uninstalling the app or limiting your use of it?
 - In the past, have you ever had your personal data misused or compromised in general? If so, what happened? What about data related to apps or web services?
- Wrap-up Alright, I have asked all the interview questions. Is there anything else you would like to say? Any questions at this point? Or any comments? Now I will be asking you to fill out a short survey. The survey link is pasted inside the chat. You will receive the e-gift card via email soon after you complete the survey. Feel free to disconnect now. Thank you so much for your participation!

C.3 Post-Interview Questions

- How much control do you think you have over the data that companies collect about you?
- How concerned are you, if at all, about how companies are using the data they collect about you?
- How much do you feel you understand what companies are doing with the data they collected about you? Prompt: a great deal, some, very little, nothing
- What is the highest level of education you have completed?
- Have you ever held a job or received a degree in computer science or any related technology field?

- Which of the following best describes your employment status?
- Privacy behavior: Have you used the following tools in the past year? Please select all that apply.
 - Antivirus software
 - Ad blockers
 - VPNs
 - Private browsing
 - Cookie/tracker blockers (e.g., Ghostery)
 - Tor Browser
 - Private search engines (e.g., DuckDuckGo)
 - PrivacyDog
 - Other, please specify

C.4 Codebook

Code categories are shown in bold type, with the list of codes in that category following. For a more detailed version with code descriptions, see <https://osf.io/47kzt/>.

App Privacy

- **iPhone usage length:** less than 5 years, at least 5 and less than 10 years, at least 10 years
- **Why use iPhone**
- **iPhone # of apps estimate:** ≤ 40 , 50–100, > 300
- **iPhone actual # of apps:** < 50 , 50–100, 101–200, > 200
- **Recent app download time:** within 1 day, within a week, within a month
- **App download process:** search keywords in App Store, search in Google, download multiple apps and then delete unwanted, learn specific app from ads, learn specific app from recommendations, generally know the app the download when in the App Store
- **Factors considered when downloading apps:** cost, reviews, ratings, utility, descriptions, brand/trust, rewards program, space/battery, number of downloads or reviews, bugs
- **Concerns about app privacy:** yes, no, consider privacy before downloading, remove app out of privacy concern, privacy concern for newly downloaded apps, privacy concern about Facebook related apps, privacy concern only for important apps
- **Whether have questions about app data collection** yes, no, specific question
- **Where to learn about app data collection** Google, terms of services/privacy policies, iPhone privacy settings, iPhone privacy prompts, in-app privacy settings, the App Store, media, downloaded data
- **Looking for app privacy in the App Store:** no

- **Seen app privacy section in the App Store before:** no, yes, yes but only aware of its existence
- **DoorDash app used before:** yes, no
- **Chipotle app used before:** yes, no

Label Understanding & Perception

- **Tracking:** understanding, confusion, concerned, not concerned, useful, not useful, mixed feelings
- **Tracking by identifiers:** concerned, not concerned, mixed feelings
- **Tracking implies aggregating location data:** yes, clear from the label, not clear from the label, concerned
- **Tracking used for advertising:** yes, clear from the label, not clear from the label
- **Data linked to you:** understanding, confusion, useful
- **Difference between data used to track you and data linked to you:** understanding, confusion
- **Data not linked to you:** understanding, confusion
- **Contact under data not linked to you:** confusion
- **Targeted advertising:** understanding, concerned, not concerned, mixed feelings, useful
- **Third party or developer advertising:** understanding, confusion
- **App functionality:** understanding, confusion, useful
- **Analytics:** understanding, confusion, useful, not concerned
- **Product Personalization:** understanding, confusion, useful, concerned
- **Identifiers:** understanding, confusion, concerned
- **Device id:** understanding, confusion
- **Product interaction:** understanding, confusion, mixed feelings
- **Browsing or search history:** understanding, confusion, concerned, not concerned
- **Other category:** understanding, confusion, concerned, not concerned
- **User content:** concern, confusion
- **Contacts used by DoorDash for other purposes:** concerned, mixed
- **Jargon confusion:** usage data, crash data, diagnostics, coarse location, purchase
- **App has no need for listed data on label**
- **Label first impression**
- **Confusion about label structure**
- **Confusion about label sections**
- **Label useful:** yes, no
- **Understood most of the labels?:** yes, no, in-between

- **Learned new things from labels?:** yes, no, in-between
- **Future use of labels:** yes, no, depends
- **Labels impacting later decision to download apps:** yes, no, depends
- **Labels include all app data collection practices:** yes, no, not sure
- **What participants like about the labels:** existence, increased transparency, other
- **What participants dislike about the labels:** vagueness, long and/or repetitive, use of jargon
- **How to improve the labels or what participants would want the labels to include:** add accessible definitions, add specific and contextualized examples of data collected, add privacy controls, add whether the data is being shared or sold and/or with whom, add data retention, explain in details what they do with the data and/or justification, add how data privacy is protected, arrange purposes in a table format, add contact info for further questions, other
- **Do participants trust the labels?** yes, no, depends, reason for yes, reason for no
- **Labels provided by:** app developers, Apple, both Apple and app developers, not sure
- **Labels reviewed or verified by Apple?** neither, only reviewed, verified, not sure
- **If labels are reviewed, how?** automated processing, manual review, both
- **Like Compact label**
- **Participants expect labels to be interactive**
- **Participants think labels are required**
- **Labels lack oversight or guarantee**

Privacy Attitudes and Experiences

- **Resignation**
- **Trade-off**
- **Not concerned about privacy:** generally unconcerned, nothing to hide, low perceived risk, privacy as a secondary task
- **Privacy concern:** generic, desire to remain personal autonomy, feeling watched, risks
- **Usable privacy:** user burden high, frustration with privacy policies
- **Privacy protection behavior**
- **More concerned after reading labels**
- **Past experience with personal data being misused or compromised?** no, data breaches, fraudulent activity on bank accounts, identity theft, accounts hacked
- **Aware of turning off tracking on iPhone:** yes, no, confusion

Bibliography

- [1] Augmented mental health: Revolutionary mental health care using emotion recognition. <https://www.augmentedmentalhealth.com/blog/augmented-mental-health-revolutionary-mental-health-care-using-emotion-recognition>, May 2018. Accessed: 2020-12-15. 3.2.2
- [2] Chinese man caught by facial recognition at pop concert. <https://www.bbc.com/news/world-asia-china-43751276>, April 2018. Accessed: 2020-12-15. 3.2.2
- [3] Mobile fact sheet. Technical report, Pew Research Center, April 2021. URL <https://www.pewresearch.org/internet/fact-sheet/mobile/>. 1.1
- [4] Alessandro Acquisti and Ralph Gross. Imagined communities: Awareness, information sharing, and privacy on the facebook. In *International Workshop on Privacy Enhancing Technologies*, pages 36–58. Springer, 2006. 2.2
- [5] Alessandro Acquisti and Jens Grossklags. Privacy and rationality in individual decision making. *IEEE Security & Privacy*, 3(1):26–33, 2005. 2.1.3
- [6] Alessandro Acquisti, Idris Adjerid, Rebecca Balebako, Laura Brandimarte, Lorrie Faith Cranor, Saranga Komanduri, Pedro Giovanni Leon, Norman Sadeh, Florian Schaub, Manya Sleeper, et al. Nudges for privacy and security: Understanding and assisting users’ choices online. *ACM Computing Surveys (CSUR)*, 50(3):1–41, 2017. 2.1.2, 3.7.5
- [7] Idris Adjerid, Alessandro Acquisti, Laura Brandimarte, and George Loewenstein. Sleights of privacy: Framing, disclosures, and the limits of transparency. In *Proceedings of the Ninth Symposium on Usable Privacy and Security, SOUPS ’13*, New York, NY, USA, 2013. Association for Computing Machinery. ISBN 9781450323192. doi: 10.1145/2501604.2501613. URL <https://doi.org/10.1145/2501604.2501613>. 2.1.1, 2.1.2
- [8] Marshall Allen. Health insurers are vacuuming up details about you — and it could raise your rates. <https://www.propublica.org/article/health-insurers-are-vacuuming-up-details-about-you-and-it-could-raise-your-rates>, July 2018. Accessed: 2020-12-15. 3.2.2
- [9] App Defense Alliance. Mobile application security assessment. URL <https://appdefensealliance.dev/masa>. 6.4.2

- [10] Hazim Almuhammedi, 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 *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*, pages 787–796, 2015. 2.1.2, 2.1.3, 3.7.5
- [11] Apple. App privacy details on the app store. <https://developer.apple.com/app-store/app-privacy-details/>, 2021. 5.3.3, 6.3, 6.4.2, 6.5.3
- [12] Apple. Privacy definitions and examples. <https://apps.apple.com/story/id1539235847>, 2021. (document), 5.2, 5.3, 5.4, 5.5, 5.6, 5.7
- [13] Apple. About privacy and location services in iOS and iPadOS. <https://support.apple.com/en-gb/HT203033>, February 3, 2022. 5.4.2
- [14] Noah Apthorpe, Yan Shvartzshnaider, Arunesh Mathur, Dillon Reisman, and Nick Feamster. Discovering smart home internet of things privacy norms using contextual integrity. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.*, 2(2), July 2018. doi: 10.1145/3214262. URL <https://doi.org/10.1145/3214262>. 2.3
- [15] Noah Apthorpe, Yan Shvartzshnaider, Arunesh Mathur, Dillon Reisman, and Nick Feamster. Discovering smart home internet of things privacy norms using contextual integrity. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.*, 2(2), July 2018. doi: 10.1145/3214262. URL <https://doi.org/10.1145/3214262>. 2.2
- [16] Noah Apthorpe, Sarah Varghese, and Nick Feamster. Evaluating the contextual integrity of privacy regulation: Parents’ IoT toy privacy norms versus COPPA. In *28th USENIX Security Symposium (USENIX Security 19)*, pages 123–140, Santa Clara, CA, August 2019. USENIX Association. ISBN 978-1-939133-06-9. URL <https://www.usenix.org/conference/usenixsecurity19/presentation/apthorpe>. 2.3, 4.2.1, 4.2.4, 4.3.3
- [17] Rachel Bachman. Your gym’s tech wants to know you better. <https://www.wsj.com/articles/your-gyms-tech-wants-to-know-you-better-1497281915>, June 2017. Accessed: 2020-12-15. 3.2.2
- [18] Sarah Pulliam Bailey. Skipping church? Facial recognition software could be tracking you. <http://www.washingtonpost.com/news/acts-of-faith/wp/2015/07/24/skipping-church-facial-recognition-software-could-be-tracking-you/>, July 2015. Accessed: 2020-12-15. 3.2.2
- [19] Rebecca Balebako, Richard Shay, and Lorrie Faith Cranor. Is your inseam a biometric? a case study on the role of usability studies in developing public policy. In *Workshop on Usable Security*, volume 23, 2014. 5.4.1
- [20] Rebecca Balebako, Florian Schaub, Idris Adjerid, Alessandro Acquisti, and Lorrie Faith Cranor. The impact of timing on the salience of smartphone app privacy notices. In *CCS Workshop on Security and Privacy in Smartphones and Mobile Devices (SPSM)*. Association for Computing Machinery, October 2015. URL <https://doi.org/10.1145/2808117.2808119>. 2.1.1, 5.4.2

- [21] Lisa Feldman Barrett and Daniel J Barrett. An introduction to computerized experience sampling in psychology. *Social Science Computer Review*, 19(2):175–185, 2001. 2.2
- [22] Douglas Bates, Martin Mächler, Ben Bolker, and Steve Walker. Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, 67(1):1–48, 2015. doi: 10.18637/jss.v067.i01. 3.4.2
- [23] Françoise Baylis and Natalie Kofler. A public health ethic should inform policies on COVID-19 immunity passports. *The Lancet Infectious Diseases*, 21(4):456, apr 2021. ISSN 1473-3099. doi: 10.1016/S1473-3099(20)30918-X. URL [https://doi.org/10.1016/S1473-3099\(20\)30918-X](https://doi.org/10.1016/S1473-3099(20)30918-X). 1.1
- [24] Monique Beals. Florida landlord requiring proof of vaccinations from tenants. <https://thehill.com/policy/healthcare/other/572370-florida-landlord-requiring-proof-of-vaccinations-from-tenants/>, September 2021. 4.2.1
- [25] Bloomberg News. Mannequins collect data on shoppers via facial-recognition software. https://www.washingtonpost.com/business/economy/mannequins-collect-data-on-shoppers-via-facial-recognition-software/2012/11/22/0751b992-3425-11e2-9cfa-e41bac906cc9_story.html, November 2012. Accessed: 2020-12-15. 3.2.2
- [26] Bram Bonné, Sai Teja Peddinti, Igor Bilogrevic, and Nina Taft. Exploring decision making with Android’s runtime permission dialogs using in-context surveys. In *Thirteenth Symposium on Usable Privacy and Security (SOUPS 2017)*, pages 195–210, Santa Clara, CA, July 2017. USENIX Association. ISBN 978-1-931971-39-3. URL <https://www.usenix.org/conference/soups2017/technical-sessions/presentation/bonne>. 2.1.2
- [27] Jan Lauren Boyles, Aaron Smith, and Mary Madden. Privacy and data management on mobile devices. *Pew Internet & American Life Project*, 4:1–19, 2012. 1.1
- [28] Virginia Braun and Victoria Clarke. Using thematic analysis in psychology. *Qualitative research in psychology*, 3(2):77–101, 2006. 3.3.1, 4.2.3
- [29] Virginia Braun and Victoria Clarke. Thematic analysis. *APA handbook of research methods in psychology, Vol 2: Research designs: Quantitative, qualitative, neuropsychological, and biological.*, pages 57–71, 2012. doi: 10.1037/13620-004. 5.2.4
- [30] David Burrows. Facial expressions show Mars the adverts that will drive sales. <https://www.foodnavigator.com/Article/2017/03/23/Facial-expressions-show-Mars-the-adverts-that-will-drive-sales>, May 2017. Accessed: 2020-12-15. 3.2.2
- [31] Ian Carlos Campbell. Google is reinstating app permissions list on play store. *The Verge*, November 5, 2020. URL <https://techcrunch.com/2022/07/21/google-app-permissions-play-store/>. 6.5.2
- [32] Laura L Carstensen, Bulent Turan, Susanne Scheibe, Nilam Ram, Hal Ersner-Hershfield, Gregory R Samanez-Larkin, Kathryn P Brooks, and John R Nesselrode. Emotional experience improves with age: Evidence based on over 10 years of experience sampling. *Psychology and Aging*, 26(1):21, 2011. 3.2.1

- [33] Daniel Castro and McLaughlin Michael. Survey: Few Americans want government to limit use of facial recognition technology, particularly for public safety or airport screening. <https://datainnovation.org/2019/01/survey-few-americans-want-government-to-limit-use-of-facial-recognition-technology-particularly-for-public-safety-or-airport-screening/>, 2019. 3.7.7, 3.7.8
- [34] CBC News. Strategic Group mandates COVID-19 vaccination for new apartment renters in Alberta . <https://www.cbc.ca/news/canada/calgary/strategic-group-rental-covid-vaccine-mandate-calgary-1.6228809>, October 2021. 4.2.1
- [35] Rex Chen, Fei Fang, Thomas Norton, Aleecia M McDonald, and Norman Sadeh. Fighting the fog: Evaluating the clarity of privacy disclosures in the age of ccpa. In *Proceedings of the 20th Workshop on Workshop on Privacy in the Electronic Society*, pages 73–102, 2021. 2.1.1
- [36] Hanbyul Choi, Jonghwa Park, and Yoonhyuk Jung. The role of privacy fatigue in online privacy behavior. *Computers in Human Behavior*, 81:42–51, 2018. 2.1.1
- [37] Richard Chow. The last mile for IoT privacy. *IEEE Security & Privacy*, 15(6):73–76, 2017. doi: 10.1109/MSP.2017.4251118. 1.1, 2.4, 3.7.8
- [38] R. H. B. Christensen. ordinal—regression models for ordinal data, 2019. R package version 2019.12-10. <https://CRAN.R-project.org/package=ordinal>. 3.4.2, 4.2.3
- [39] Tamlin Conner Christensen, Lisa Feldman Barrett, Eliza Bliss-Moreau, Kirsten Lebo, and Cynthia Kaschub. A practical guide to experience-sampling procedures. *Journal of Happiness Studies*, 4(1):53–78, 2003. 3.2.4
- [40] Liat Clark. Mannequins are spying on shoppers for market analysis. <https://www.wired.co.uk/article/mannequin-spies-on-customers>, November 2012. Accessed: 2020-12-15. 3.2.2
- [41] Ignacio N. Cofone. Immunity passports and contact tracing surveillance. *24 STAN. TECH. L. REV. (2021)*, 24:176–236, 2021. URL <https://law.stanford.edu/publications/immunity-passports-and-contact-tracing-surveillance/>. 1.1
- [42] Jessica Colnago and Hélio Guardia. How to inform privacy agents on preferred level of user control? In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct (UbiComp '16)*, pages 1542–1547, 2016. ISBN 9781450344623. doi: 10.1145/2968219.2968546. URL <https://doi.org/10.1145/2968219.2968546>. 3.7.5
- [43] 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 *Proceedings of the 2020 CHI Conference on Human Factors in Computing System (CHI '20)*, pages 1–13, 2020. ISBN 9781450367080. doi: 10.1145/3313831.3376389. URL <https://doi.org/10.1145/3313831.3376389>. 2.1.2, 2.1.3, 3.7.4

- [44] European Commission. Covid-19: Digital green certificates. https://ec.europa.eu/info/live-work-travel-eu/coronavirus-response/safe-covid-19-vaccine-s-europeans/covid-19-digital-green-certificates_en, March 2021. Accessed: 2021-4-09. 4.1, 4.2.1
- [45] Ben Conarck. Florida court: Prosecutors had no obligation to turn over facial recognition evidence. <https://www.jacksonville.com/news/20190123/florida-court-prosecutors-had-no-obligation-to-turn-over-facial-recognition-evidence>, January 2019. Accessed: 2020-12-15. 3.2.2
- [46] Kate Conger, Richard Fausset, and Serge F. Kovalski. San Francisco bans facial recognition technology. <https://www.nytimes.com/2019/05/14/us/facial-recognition-ban-san-francisco.html>, May 2019. Accessed: 2020-12-15. 3.1
- [47] Sunny Consolvo and Miriam Walker. Using the experience sampling method to evaluate ubicomp applications. *IEEE Pervasive Computing*, 2(2):24–31, 2003. 2.2
- [48] Sylvie Corbet. France’s virus pass now required in restaurants, trains. <https://apnews.com/article/europe-business-health-france-coronavirus-pandemic-655d8451d7494f8663ce2072e64cf7a6>, August 2021. Accessed: 2021-8-11. 4.2.1
- [49] Lorrie Faith Cranor. Necessary but not sufficient: Standardized mechanisms for privacy notice and choice. *J. on Telecomm. & High Tech. L.*, 10:273, 2012. 2.1.2
- [50] Anupam Das, Martin Degeling, Xiaoyou Wang, Junjue Wang, Norman Sadeh, and Mahadev Satyanarayanan. Assisting users in a world full of cameras: A privacy-aware infrastructure for computer vision applications. In *2017 IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 1387–1396. IEEE, 2017. 2.4, 3.1, 3.6, 3.7.8
- [51] Anupam Das, Martin Degeling, Daniel Smullen, and Norman Sadeh. Personalized privacy assistants for the Internet of Things: Providing users with notice and choice. *IEEE Pervasive Computing*, 17(3):35–46, 2018. doi: 10.1109/MPRV.2018.03367733. 1.1, 2.1.3, 2.4, 3.7.4, 3.7.8
- [52] Bobby J Davidson. How your business can benefit from facial recognition technology. <https://percentotech.com/how-your-business-can-benefit-from-facial-recognition-technology/>, November 2019. Accessed: 2020-12-15. 3.2.2
- [53] Dean DeChiaro. New York City eyes regulation of facial recognition technology. <https://www.rollcall.com/news/congress/new-york-city-eyes-regulation-of-facial-recognition-technology/>, October 2019. Accessed: 2020-12-15. 3.1
- [54] Nik DeCosta-Klipa. Massachusetts expands COVID-19 vaccine requirement to more workers caring for the elderly. <https://www.boston.com/news/coronavirus/2021/09/01/massachusetts-expands-covid-19-vaccine-mandate-elderly-care/>, September 2021. 4.2.1

- [55] Martin Degeling, Christine Utz, Christopher Lentzsch, Henry Hosseini, Florian Schaub, and Thorsten Holz. We value your privacy... now take some cookies: Measuring the gdpr's impact on web privacy. In *Proceedings of Network and Distributed System Security Symposium(NDSS '19)*, 2019. 2.1.2
- [56] Ron DeSantis. Governor ron desantis signs landmark legislation to ban vaccine passports and stem government overreach. <https://www.flgov.com/2021/05/03/governor-ron-desantis-signs-landmark-legislation-to-ban-vaccine-passports-and-stem-government-overreach/>, May 2021. Accessed: 2021-9-01. 4.2.2
- [57] Jessica Dickler. More than half of employers to require Covid vaccines as omicron fears grow. <https://www.cnbc.com/2021/12/01/more-employers-will-make-vaccine-s-mandatory-as-omicron-fears-grow.html>, December 2021. 4.2.1
- [58] Benchaa Djellali, Kheira Belarbi, Abdallah Chouarfia, and Pascal Lorenz. User authentication scheme preserving anonymity for ubiquitous devices. *Security and Communication Networks*, 8(17):3131–3141, 2015. 2.4
- [59] Zak Doffman. Facebook Tracks Your iPhone Location—This Is How To Stop It. *Forbes*, May May 22, 2021. URL <https://www.forbes.com/sites/zakdoffman/2021/05/22/apple-user-warning-how-to-stop-facebook-secretly-tracking-your-iphone-ipad/>. 5.4.3
- [60] Yitao Duan and John Canny. Protecting user data in ubiquitous computing: Towards trustworthy environments. In *International Workshop on Privacy Enhancing Technologies*, pages 167–185. Springer, 2004. 2.4
- [61] Chris Duckett. Western australia finally thinks about quarantining covid check-in info from cops. <https://www.zdnet.com/article/western-australia-finally-thinks-about-quarantining-covid-check-in-info-from-cops/>, 2021. 2.3, 4.2.1
- [62] Nico Ebert, Kurt Alexander Ackermann, and Björn Schepler. Bolder is better: Raising user awareness through salient and concise privacy notices. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, pages 1–12, 2021. 2.1.1, 5.4.2
- [63] Serge Egelman, Janice Tsai, Lorrie Faith Cranor, and Alessandro Acquisti. Timing is everything? the effects of timing and placement of online privacy indicators. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 319–328, 2009. 2.1.1, 5.4.2
- [64] Melanie Ehrenkranz. Burger joint teams up with surveillance giant to scan your face for loyalty points. <https://gizmodo.com/burger-joint-teams-up-with-surveillance-giant-to-scan-y-1821498988>, December 2017. Accessed: 2020-12-15. 3.2.2
- [65] Mahmoud Elkhodr, Seyed Shahrestani, and Hon Cheung. A contextual-adaptive location disclosure agent for general devices in the internet of things. In *38th Annual IEEE Conference on Local Computer Networks - Workshops*, pages 848–855. IEEE, 2013. 2.1.3

- [66] Pardis Emami-Naeini, Henry Dixon, Yuvraj Agarwal, and Lorrie Faith Cranor. Exploring how privacy and security factor into IoT device purchase behavior. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, CHI '19, page 1–12, New York, NY, USA, 2019. Association for Computing Machinery. ISBN 9781450359702. doi: 10.1145/3290605.3300764. URL <https://doi.org/10.1145/3290605.3300764>. 5.2.4
- [67] Darrell Etherington. Baidu and KFC's new smart restaurant suggests what to order based on your face. <https://techcrunch.com/2016/12/23/baidu-and-kfcs-new-smart-restaurant-suggests-what-to-order-based-on-your-face/>, December 2016. Accessed: 2020-12-15. 3.2.2, 3.5.4, 3.6.3
- [68] Benjamin Fabian, Tatiana Ermakova, and Tino Lentz. Large-scale readability analysis of privacy policies. In *Proceedings of the International Conference on Web Intelligence*, WI '17, page 18–25, New York, NY, USA, 2017. Association for Computing Machinery. ISBN 9781450349512. doi: 10.1145/3106426.3106427. URL <https://doi.org/10.1145/3106426.3106427>. 2.1.1
- [69] Ingrid Fadelli. Analyzing spoken language and 3-D facial expressions to measure depression severity. <https://techxplore.com/news/2018-11-spoken-language-d-facial-d-expression.html>, December 2019. Accessed: 2020-12-15. 3.2.2
- [70] Caitlin Fairchild. Hertz is now using facial recognition to check out cars, December 2018. URL <https://www.nextgov.com/emerging-tech/2018/12/hertz-now-using-facial-recognition-check-out-cars/153479/>. 3.1
- [71] Ryan Faircloth. University of minnesota regents approve covid vaccination requirement for students. <https://www.startribune.com/university-of-minnesota-regents-approve-covid-vaccination-requirement-for-students/600087678/>, August 2021. 4.2.1
- [72] Hao-Shu Fang, Shuqin Xie, Yu-Wing Tai, and Cewu Lu. RMPE: Regional multi-person pose estimation. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2017. 2.4
- [73] Federal Trade Commission. Mobile privacy disclosures: Building trust through transparency (FTC staff report). <https://www.ftc.gov/reports/mobile-privacy-disclosures-building-trust-through-transparency-federal-trade-commission>, February 2013. 2.1.1
- [74] Federal Trade Commission. Internet of things: Privacy & security in a connected world (FTC staff report). <https://www.ftc.gov/reports/federal-trade-commission-staff-report-november-2013-workshop-entitled-internet-things>, January 2015. 2.1.1
- [75] Yuanyuan Feng, Yaxing Yao, and Norman Sadeh. A design space for privacy choices: Towards meaningful privacy control in the internet of things. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, pages 1–16, 2021. 2.1.2

- [76] Denzil Ferreira, Jorge Goncalves, Vassilis Kostakos, Louise Barkhuus, and Anind K Dey. Contextual experience sampling of mobile application micro-usage. In *Proceedings of the 16th International Conference on Human-computer Interaction with Mobile Devices & Services (MobileHCI '14)*, pages 91–100, 2014. 3.2.1
- [77] The COVID-19 Consortium for Understanding the Public’s Policy Preferences Across States. The COVID states project. URL <https://www.covidstates.org/>. 4.6, 4.6.1
- [78] The Commons Project Foundation. CommonPass. <https://commonpass.org>, 2020. Accessed: 2021-4-09. 4.2.1
- [79] Chris Frey. Revealed: how facial recognition has invaded shops—and your privacy. <https://www.theguardian.com/cities/2016/mar/03/revealed-facial-recognition-software-infiltrating-cities-saks-toronto>, March 2016. Accessed: 2020-12-15. 3.2.2, 3.5.4, 3.6.3
- [80] Cary Funk and Tyson Alec. Growing share of americans say they plan to get a covid-19 vaccine – or already have. <https://www.pewresearch.org/science/2021/03/05/growing-share-of-americans-say-they-plan-to-get-a-covid-19-vaccine-or-already-have/>, 2021. 4.2.1
- [81] Sarah Fister Gale. Employers turn to biometric technology to track attendance. <https://www.workforce.com/news/employers-turn-to-biometric-technology-to-track-attendance>, March 2013. Accessed: 2020-12-15. 3.2.2
- [82] Mirta Galesic and Michael Bosnjak. Effects of Questionnaire Length on Participation and Indicators of Response Quality in a Web Survey. *Public Opinion Quarterly*, 73(2):349–360, 05 2009. ISSN 0033-362X. doi: 10.1093/poq/nfp031. URL <https://doi.org/10.1093/poq/nfp031>. 4.2.1
- [83] Jack Gardner, Yuanyuan Feng, Kayla Reiman, Zhi Lin, Akshath Jain, and Norman Sadeh. Helping mobile application developers create accurate privacy labels. *International Workshop on Privacy Engineering (IWPE'22)*, 2022. URL <https://privacyassistant.org/media/publications/IWPE2022.pdf>. 5.4.1
- [84] Frederic Gerdon, Helen Nissenbaum, Ruben L. Bach, Frauke Kreuter, and Stefan Zins. Individual acceptance of using health data for private and public benefit: Changes during the covid-19 pandemic. *Harvard Data Science Review*, (Special Issue 1), 4 2021. doi: 10.1162/99608f92.edf2fc97. URL <https://hdsr.mitpress.mit.edu/pub/3uc0p5dq>. 2.3
- [85] Shirin Ghaffary and Rani Molla. Here’s where the us government is using facial recognition technology to surveil americans, 2019. URL <https://www.vox.com/recode/2019/7/18/20698307/facial-recognition-technology-us-government-fight-for-the-future>. 3.1
- [86] Samuel Gibbs. Google has been tracking Android users even with location services turned off. *The Guardian*, Nov 22 2017. URL <https://www.theguardian.com/technology/2017/nov/22/google-track-android-users-location-services-turned-off-sim>. 5.4.3

- [87] Ella Glover. Vaccine passports ‘could be mandatory in pubs, bars and restaurants in bid to boost jab rates in young’. <https://www.independent.co.uk/news/uk/home-news/vaccine-passports-hospitality-coronavirus-restrictions-b1881618.html>, July 2021. Accessed: 2021-8-11. 4.1, 4.2.1
- [88] Joshua Gluck, Florian Schaub, Amy Friedman, Hana Habib, Norman Sadeh, Lorrie Faith Cranor, and Yuvraj Agarwal. How short is too short? implications of length and framing on the effectiveness of privacy notices. In *Twelfth Symposium on Usable Privacy and Security (SOUPS 2016)*, pages 321–340, 2016. 2.1.1
- [89] Google. Provide information for Google Play’s Data safety section. <https://support.google.com/googleplay/android-developer/answer/10787469>, August 2022. 6.3, 6.4.2
- [90] Colin M Gray, Yubo Kou, Bryan Battles, Joseph Hoggatt, and Austin L Toombs. The dark (patterns) side of ux design. In *Proceedings of the 2018 CHI conference on human factors in computing systems*, pages 1–14, 2018. 2.1.2
- [91] Colin M. Gray, Cristiana Santos, Nataliia Bielova, Michael Toth, and Damian Clifford. Dark patterns and the legal requirements of consent banners: An interaction criticism perspective. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, CHI ’21, New York, NY, USA, 2021. Association for Computing Machinery. ISBN 9781450380966. doi: 10.1145/3411764.3445779. URL <https://doi.org/10.1145/3411764.3445779>. 2.1.2
- [92] Nielsen Norman Group. Top 10 design mistakes in the unsubscribe experience. <https://www.nngroup.com/articles/unsubscribe-mistakes/>, April 29th, 2018. 2.1.2
- [93] Yaron Gurovich, Yair Hanani, Omri Bar, Guy Nadav, Nicole Fleischer, Dekel Gelbman, Lina Basel-Salmon, Peter M Krawitz, Susanne B Kamphausen, Martin Zenker, Lynne M Bird, and Karen W Gripp. Identifying facial phenotypes of genetic disorders using deep learning. *Nature Medicine*, 25(1):60–64, 2019. ISSN 1546-170X. doi: 10.1038/s41591-018-0279-0. URL <https://doi.org/10.1038/s41591-018-0279-0>. 3.2.2
- [94] Hana Habib and Lorrie Faith Cranor. Evaluating the usability of privacy choice mechanisms. In *Eighteenth Symposium on Usable Privacy and Security (SOUPS 2022)*, pages 273–289, Boston, MA, August 2022. USENIX Association. ISBN 978-1-939133-30-4. URL <https://www.usenix.org/conference/soups2022/presentation/habib>. 2.1.2
- [95] 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 *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, pages 1–12, 2020. 5.2.4

- [96] Hana Habib, Yixin Zou, Yaxing Yao, Alessandro Acquisti, Lorrie 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 *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, CHI '21, New York, NY, USA, 2021. Association for Computing Machinery. ISBN 9781450380966. doi: 10.1145/3411764.3445387. URL <https://doi.org/10.1145/3411764.3445387>. 2.1.1, 2.1.2, 5.4.2
- [97] Hana Habib, Megan Li, Ellie Young, and Lorrie Cranor. “okay, whatever”: An evaluation of cookie consent interfaces. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, CHI '22, New York, NY, USA, 2022. Association for Computing Machinery. ISBN 9781450391573. doi: 10.1145/3491102.3501985. URL <https://doi.org/10.1145/3491102.3501985>. 2.1.2
- [98] Foad Hamidi, Morgan Klaus Scheuerman, and Stacy M. Branham. *Gender Recognition or Gender Reductionism? The Social Implications of Embedded Gender Recognition Systems*, page 1–13. ACM, New York, NY, USA, 2018. ISBN 9781450356206. URL <https://doi.org/10.1145/3173574.3173582>. 3.7.6
- [99] Alexis Hancock and Karen Gullo. Immunity passports are a threat to our privacy and information security. <https://www.eff.org/deeplinks/2020/05/immunity-passports-are-threat-our-privacy-and-information-security>, May 2020. Accessed: 2021-4-09. 1.1
- [100] Joel M Hektner, Jennifer A Schmidt, and Mihaly Csikszentmihalyi. *Experience sampling method: Measuring the quality of everyday life*. Sage, 2007. 2.2, 3.2.1
- [101] Kashmir Hill. Wrongfully accused by an algorithm. <https://www.nytimes.com/2020/06/24/technology/facial-recognition-arrest.html>, 2020. 3.7.6
- [102] Mariko Hirose. Privacy in public spaces: The reasonable expectation of privacy against the dragnet use of facial recognition technology. *Connecticut Law Review*, (377), 2017. URL https://opencommons.uconn.edu/law_review/377. 3.7.7
- [103] Wilhelm Hofmann, Roy F Baumeister, Georg Förster, and Kathleen D Vohs. Everyday temptations: An experience sampling study of desire, conflict, and self-control. *Journal of Personality and Social Psychology*, 102(6):1318, 2012. 3.2.1
- [104] Jason Hong. The privacy landscape of pervasive computing. *IEEE Pervasive Computing*, 16(3):40–48, 2017. doi: 10.1109/MPRV.2017.2940957. 1.1
- [105] Jason I Hong and James A Landay. An architecture for privacy-sensitive ubiquitous computing. In *Proceedings of the 2nd International Conference on Mobile Systems, Applications, and Services (MobiSys '04)*, pages 177–189, 2004. 2.1.3
- [106] Jacqueline Howard. A third dose of covid-19 vaccine is now authorized for some. here’s what you need to know about boosters for all. <https://www.cnn.com/2021/08/14/health/covid-19-vaccine-boosters-explainer-fda-wellness/index.html>, August 2021. 4.2.1

- [107] Timothy Johnson. Shoplifters meet their match as retailers deploy facial recognition cameras. <https://www.mcclatchydc.com/news/nation-world/national/article211455924.html>, May 2018. Accessed: 2020-12-15. 3.2.2
- [108] Norman Sadeh Jonathan Mugan, Tarun Sharma. Understandable learning of privacy preferences through default personas and suggestions. Technical Report CMU-ISR-11-112, Carnegie Mellon University, School of Computer Science, August 2011. 2.2
- [109] Justin Henry. State launches contact tracing app. <https://www.cpbj.com/pennsylvania-launches-contact-tracing-app/>, Sep 2020. 1.1
- [110] Eiman Kanjo, Luluah Al-Husain, and Alan Chamberlain. Emotions in context: examining pervasive affective sensing systems, applications, and analyses. *Personal and Ubiquitous Computing*, 19(7):1197–1212, 2015. ISSN 1617-4917. doi: 10.1007/s00779-015-0842-3. URL <https://doi.org/10.1007/s00779-015-0842-3>. 1.1, 2.4
- [111] Maitrik Kataria. App Usage Statistics 2021 that’ll Surprise You. <https://www.simform.com/blog/the-state-of-mobile-app-usage/>, Jan 2021. 1.1
- [112] Suzanne Rowan Kelleher. U.s. to require foreign visitors to be fully vaccinated. <https://www.forbes.com/sites/suzannerowankelleher/2021/08/05/us-to-require-foreign-visitors-to-be-fully-vaccinated/?sh=452979155cb2>, August 2021. 4.2.1
- [113] Patrick Gage Kelley, Joanna Bresee, Lorrie Faith Cranor, and Robert W Reeder. A “nutrition label” for privacy. In *Proceedings of the 5th Symposium on Usable Privacy and Security*, pages 1–12, 2009. 2.1.2, 5.1
- [114] Patrick Gage Kelley, Lucian Cesca, Joanna Bresee, and Lorrie Faith Cranor. Standardizing privacy notices: An online study of the nutrition label approach. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI ’10, page 1573–1582, New York, NY, USA, 2010. Association for Computing Machinery. ISBN 9781605589299. doi: 10.1145/1753326.1753561. URL <https://doi.org/10.1145/1753326.1753561>. 2.1.2, 5.4.1
- [115] Patrick Gage Kelley, Sunny Consolvo, Lorrie Faith Cranor, Jaeyeon Jung, Norman Sadeh, and David Wetherall. A conundrum of permissions: Installing applications on an android smartphone. In Jim Blyth, Sven Dietrich, and L. Jean Camp, editors, *Financial Cryptography and Data Security*, pages 68–79, Berlin, Heidelberg, 2012. Springer Berlin Heidelberg. 5.1
- [116] Patrick Gage Kelley, Lorrie Faith Cranor, and Norman Sadeh. Privacy as part of the app decision-making process. In *Proceedings of the SIGCHI conference on human factors in computing systems*, pages 3393–3402, 2013. 5.1
- [117] Mehreen Khan. Eu plans sweeping regulation of facial recognition, August 2019. URL <https://www.ft.com/content/90ce2dce-c413-11e9-a8e9-296ca66511c9>. 3.1
- [118] Kenneth Kiesnoski. Vaccine passports gain traction as delta variant threatens travel rebound. <https://www.cNBC.com/2021/08/11/vaccine-passports-gain-traction-as-delta-variant-threatens-travel-boom.html>, August 2021. 4.2.1

- [119] Michelle E. Kiger and Lara Varpio. Thematic analysis of qualitative data: Amee guide no. 131. *Medical Teacher*, 42(8):846–854, 2020. doi: 10.1080/0142159X.2020.1755030. URL <https://doi.org/10.1080/0142159X.2020.1755030>. PMID: 32356468. 6.3
- [120] Simon Koch, Malte Wessels, Benjamin Altpeter, Madita Olvermann, and Martin Johns. Keeping privacy labels honest. *Proc. Priv. Enhancing Technol.*, 2022(4):486–506, 2022. doi: 10.2478/popets-2022-0119. URL <https://www.petsymposium.org/2022/files/papers/issue4/popets-2022-0119.pdf>. 6.5.4
- [121] Natalie Kofler and Françoise Baylis. Ten reasons why immunity passports are a bad idea. *Nature*, 580:379–381, May 2021. doi: 10.1038/d41586-020-01451-0. URL <https://doi.org/10.1038/d41586-020-01451-0>. 4.2.1
- [122] Konrad Kollnig, Anastasia Shuba, Max Van Kleek, Reuben Binns, and Nigel Shadbolt. Goodbye tracking? impact of ios app tracking transparency and privacy labels. In *2022 ACM Conference on Fairness, Accountability, and Transparency, FAccT '22*, page 508–520, New York, NY, USA, 2022. Association for Computing Machinery. ISBN 9781450393522. doi: 10.1145/3531146.3533116. URL <https://doi.org/10.1145/3531146.3533116>. 6.5.4
- [123] Douglas Korgut and Daniel Fernando Pigatto. An internet of things-based house monitoring system. In *2018 IEEE Symposium on Computers and Communications (ISCC)*, pages 01149–01152, June 2018. doi: 10.1109/ISCC.2018.8538680. 1.1, 2.4
- [124] Ingrid Kramer et al. A therapeutic application of the experience sampling method in the treatment of depression: a randomized controlled trial. *World Psychiatry*, 13(1):68–77, 2014. 3.2.1
- [125] Sarah Krouse. The new ways your boss is spying on you. <https://www.wsj.com/articles/the-new-ways-your-boss-is-spying-on-you-11563528604>, July 2019. Accessed: 2020-12-15. 3.2.2
- [126] Justin Kruger and David Dunning. Unskilled and unaware of it: How difficulties in recognizing one’s own incompetence lead to inflated self-assessments. *Journal of Personality and Social Psychology*, 77(6):1121–1134, 1999. doi: 10.1037/0022-3514.77.6.1121. 3.4.3
- [127] Paul Lavrakas. *Encyclopedia of Survey Research Methods*, 2008. URL <https://methods.sagepub.com/reference/encyclopedia-of-survey-research-methods>. 4.2.1
- [128] Hosub Lee and Alfred Kobsa. Understanding user privacy in internet of things environments. In *2016 IEEE 3rd World Forum on Internet of Things (WF-IoT)*, pages 407–412, 2016. 2.2
- [129] Hosub Lee and Alfred Kobsa. Privacy preference modeling and prediction in a simulated campuswide IoT environment. In *2017 IEEE International Conference on Pervasive Computing and Communications (PerCom)*, pages 276–285, 2017. 3.2.2
- [130] Stephen Lepitak. Disney’s Dumbo and Accenture Interactive collaborate for the movie poster of the future. <https://www.thedrum.com/news/2019/03/10/disneys-dumbo-and-accenture-interactive-collaborate-the-movie-poster-the-future>, March 2019. Accessed: 2020-12-15. 3.2.2

- [131] Gil Levi and Tal Hassner. Emotion recognition in the wild via convolutional neural networks and mapped binary patterns. In *Proceedings of the 2015 ACM on International Conference on Multimodal Interaction*, page 503–510, New York, NY, USA, 2015. ACM. ISBN 9781450339124. doi: 10.1145/2818346.2830587. URL <https://doi.org/10.1145/2818346.2830587>. 2.4
- [132] David Levine. What high-tech tools are available to fight depression? <https://health.usnews.com/health-care/patient-advice/articles/2017-10-06/what-high-tech-tools-are-available-to-fight-depression>, October 2017. Accessed: 2020-12-15. 3.2.2
- [133] David Levine. What your face may tell lenders about whether you’re creditworthy. <https://www.wsj.com/articles/what-your-face-may-tell-lenders-about-whether-youre-creditworthy-11560218700>, June 2019. Accessed: 2020-12-15. 3.2.2
- [134] Tianshi Li, Kayla Reiman, Yuvraj Agarwal, Lorrie Faith Cranor, and Jason I. Hong. Understanding challenges for developers to create accurate privacy nutrition labels. In *CHI Conference on Human Factors in Computing Systems, CHI ’22*, New York, NY, USA, 2022. Association for Computing Machinery. ISBN 9781450391573. doi: 10.1145/3491102.3502012. URL <https://doi.org/10.1145/3491102.3502012>. 5.4.1
- [135] Jialiu Lin, Shahriyar Amini, Jason I Hong, Norman Sadeh, Janne Lindqvist, and Joy Zhang. Expectation and purpose: Understanding users’ mental models of mobile app privacy through crowdsourcing. In *Proceedings of the 2012 ACM Conference on Ubiquitous Computing (UbiComp ’12)*, pages 501–510. ACM, 2012. 2.2, 3.4.4
- [136] Jialiu Lin, Bin Liu, Norman Sadeh, and Jason I. Hong. Modeling users’ mobile app privacy preferences: Restoring usability in a sea of permission settings. In *Proceedings of the Tenth Symposium on Usable Privacy and Security (SOUPS ’14)*, pages 199–212, 2014. ISBN 978-1-931971-13-3. URL <https://www.usenix.org/conference/soups2014/proceedings/presentation/lin>. 1.1, 2.1.3, 2.2, 3.6.1, 3.7.3
- [137] Bin Liu. Can Machine Learning Help People Configure Their Mobile App Privacy Settings? 1 2020. doi: 10.1184/R1/11591340.v1. URL https://kilthub.cmu.edu/articles/thesis/Can_Machine_Learning_Help_People_Configure_Their_Mobile_App_Privacy_Settings_/11591340. 1.1, 3.6.4
- [138] Bin Liu, Jialiu Lin, and Norman Sadeh. Reconciling mobile app privacy and usability on smartphones: Could user privacy profiles help? In *Proceedings of the 23rd International Conference on World Wide Web (WWW ’14)*, pages 201–212, New York, NY, USA, 2014. ISBN 978-1-4503-2744-2. doi: 10.1145/2566486.2568035. URL <http://doi.acm.org/10.1145/2566486.2568035>. 1.1, 2.1.3, 2.2

- [139] 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 *Twelfth Symposium on Usable Privacy and Security (SOUPS '16)*, pages 27–41, 2016. ISBN 978-1-931971-31-7. URL <https://www.usenix.org/conference/soups2016/technical-sessions/presentation/liu>. 1.1, 2.1.3, 2.2, 3.6.1, 3.6.4, 3.7.4, 3.7.5, 5.4.4
- [140] Jordan Liz. COVID-19, immunoprivilege and structural inequalities. *History and Philosophy of the Life Sciences*, 43(1):19, 2021. ISSN 1742-6316. URL <https://doi.org/10.1007/s40656-020-00356-5>. 1.1
- [141] Brain Logan. Pay-per-laugh: the comedy club that charges punters having fun. <https://www.theguardian.com/stage/2014/oct/14/standup-comedy-pay-per-laugh-charge-barcelona>, October 2014. Accessed: 2020-12-15. 3.2.2
- [142] Dominique Machuletz and Rainer Böhme. Multiple purposes, multiple problems: A user study of consent dialogs after GDPR. *Proceedings on Privacy Enhancing Technologies*, 2020 (2):481–498, apr 2020. doi: 10.2478/popets-2020-0037. URL <https://doi.org/10.2478/popets-2020-0037>. 2.1.2
- [143] Naresh K Malhotra, Sung S Kim, and James Agarwal. Internet users’ information privacy concerns (IUIPC): The construct, the scale, and a causal model. *Information Systems Research*, 15(4):336–355, 2004. 3.2.4
- [144] Leandro Y. Mano et al. Exploiting IoT technologies for enhancing health smart homes through patient identification and emotion recognition. *Computer Communications*, 89-90: 178–190, 2016. 1.1, 2.4
- [145] Alix Martichoux. California launches electronic vaccine verification system. <https://abc7news.com/california-vaccine-verification-passport-ca-electronic-vaccination-proof-digital/10805723/>, June 2021. 4.2.1
- [146] Kirsten Martin and Helen Nissenbaum. Measuring privacy: an empirical test using context to expose confounding variables. *Columbia Science and Technology Law Review*, 18:176, 2015. 2.3
- [147] Kirsten Martin and Katie Shilton. Putting mobile application privacy in context: An empirical study of user privacy expectations for mobile devices. *The Information Society*, 32(3):200–216, 2016. 2.2
- [148] 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. URL <https://doi.org/10.1080/1369118X.2018.1450432>. 1.1
- [149] Aleecia M McDonald and Lorrie Faith Cranor. The cost of reading privacy policies. *I/S: A Journal of Law and Policy for the Information Society*, 4:543, 2008. 1.1, 2.1.1

- [150] Nora McDonald, Sarita Schoenebeck, and Andrea Forte. Reliability and inter-rater reliability in qualitative research: Norms and guidelines for cscw and hci practice. *Proc. ACM Hum.-Comput. Interact.*, 3(CSCW), nov 2019. doi: 10.1145/3359174. URL <https://doi.org/10.1145/3359174>. 5.2.4, 6.3
- [151] Mia Sato. Singapore’s police now have access to contact tracing data. <https://www.technologyreview.com/2021/01/05/1015734/singapore-contact-tracing-police-data-covid/>, Jan 2021. 2.3
- [152] Kristopher Micinski, Daniel Votipka, Rock Stevens, Nikolaos Kofinas, Michelle L. Mazurek, and Jeffrey S. Foster. User Interactions and Permission Use on Android. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI)*, pages 362–373, Denver, Colorado, May 2017. 2.1.2
- [153] George R Milne, Mary J Culnan, and Henry Greene. A longitudinal assessment of online privacy notice readability. *Journal of Public Policy & Marketing*, 25(2):238–249, 2006. 2.1.1
- [154] Josh Moody. Colleges requiring a coronavirus vaccine for fall. <https://www.usnews.com/education/best-colleges/articles/colleges-requiring-a-coronavirus-vaccine-for-fall-what-to-know>, August 2021. Accessed: 2021-08-12. 4.1
- [155] Mozilla Wiki. Privacy icons. https://wiki.mozilla.org/Privacy_Icons, June 2011. 2.1.1, 2.1.2
- [156] Darren Murph. SceneTap app analyzes pubs and clubs in real-time, probably won’t score you a Jersey Shore cameo. <https://www.engadget.com/2011/06/12/scenetap-app-analyzes-pubs-and-clubs-in-real-time-probably-won/>, June 2011. Accessed: 2020-12-15. 3.2.2
- [157] Katie Nadworny. Some israeli workplaces require use of “green pass”. <https://www.shrm.org/resourcesandtools/hr-topics/global-hr/pages/coronavirus-israel-green-pass.aspx>, December 2021. Accessed: 2021-12-03. 4.1, 4.5
- [158] Pardis Emami Naeini, Sruti Bhagavatula, Hana Habib, Martin Degeling, Lujo Bauer, Lorie Faith Cranor, and Norman Sadeh. Privacy expectations and preferences in an iot world. In *Proceedings of the Thirteenth Symposium on Usable Privacy and Security (SOUPS ’17)*, pages 399–412, 2017. 2.2, 3.2.2, 3.3, 3.7.3
- [159] NEC Corporation. New biometric identification tools used in theme parks. <https://www.nec.com/en/global/about/mitatv/03/3.html>, 2002. Accessed: 2020-12-15. 3.2.2
- [160] Lindsey Van Ness. For states’ COVID contact tracing apps, privacy tops utility. Technical report, The Pew Charitable Trusts, March 2021. URL <https://www.pewtrusts.org/en/research-and-analysis/blogs/stateline/2021/03/19/for-states-covid-contact-tracing-apps-privacy-tops-utility>. 1.1
- [161] Alferd Ng. With facial recognition, shoplifting may get you banned in places you’ve never been, March 2019. URL <https://www.cnet.com/news/with-facial-recognition-shoplifting-may-get-you-banned-in-places-youve-never-been/>. 3.1

- [162] Helen Nissenbaum. Privacy as contextual integrity. *Washington Law Review*, 79(1):119, 2004. 1.2, 2.3, 2.4, 3.1, 3.4.1, 3.7.3, 6.5.2
- [163] Helen Nissenbaum. *Privacy in context: Technology, policy, and the integrity of social life*. Stanford University Press, 2009. 2.3, 2.4, 3.2.1, 3.5, 4.1
- [164] Helen Nissenbaum. *Respect for context as a benchmark for privacy online: what it is and isn't*, page 278–302. Cambridge University Press, Cambridge, UK, 2015. doi: 10.1017/CBO9781107280557.016. 4.5
- [165] 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 *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, CHI '20, page 1–13, New York, NY, USA, 2020. Association for Computing Machinery. ISBN 9781450367080. doi: 10.1145/3313831.3376321. URL <https://doi.org/10.1145/3313831.3376321>. 2.1.2
- [166] The State of California. California implements first-in-the-nation measure to encourage teachers and school staff to get vaccinated. <https://www.gov.ca.gov/2021/08/11/california-implements-first-in-the-nation-measure-to-encourage-teacher-s-and-school-staff-to-get-vaccinated/>, August 2021. 4.2.1, 4.2.1
- [167] The State of Georgia. Prohibition of COVID-19 vaccine passports. <https://gov.georgia.gov/document/2021-executive-order/05252101/download>, May 2021. Accessed: 2021-9-01. 4.2.2
- [168] Hawai'i State Department of Health. State of hawai'i safe travels hawai'i program. <https://hawaiicovid19.com/travel/travel-overview/>, 2021. Accessed: 2021-9-01. 1.1, 4.1, 4.2.2
- [169] Ministry of Health. Green pass, vaccination certificate and certificate of recovery. <https://www.gov.il/en/departments/general/corona-certificates>, 2021. Accessed: 2021-4-09. 4.2.1
- [170] The City of New York. Mayor de Blasio Announces Vaccine Mandate for Private Sector Workers, and Major Expansions to Nation-Leading “Key to NYC” Program . <https://www1.nyc.gov/office-of-the-mayor/news/807-21/mayor-de-blasio-vaccine-mandate-private-sector-workers-major-expansions-to>, December 2021. 4.2.1
- [171] The State of New York. Excelsior pass. <https://covid19vaccine.health.ny.gov/excelsior-pass>, March 2021. Accessed: 2021-4-09. 1.1, 4.1, 4.2.1, 4.2.2
- [172] PCMag Staff. NEC unveils facial-recognition system to identify shoppers. <https://www.pcmag.com/archive/nec-unveils-facial-recognition-system-to-identify-shoppers-305015>, November 2012. Accessed: 2020-12-15. 3.2.2
- [173] Sarah Pearman, Ellie Young, and Lorrie Cranor. User-friendly yet rarely read: A case study on the redesign of an online HIPAA authorization. *Proceedings on Privacy Enhancing Technologies (PoPETS)*, (3), 2022. 2.1.2

- [174] 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. ISSN 0022-1031. doi: <https://doi.org/10.1016/j.jesp.2017.01.006>. URL <https://www.sciencedirect.com/science/article/pii/S0022103116303201>. 4.2.2
- [175] Veljko Pejovic, Neal Lathia, Cecilia Mascolo, and Mirco Musolesi. *Mobile-Based Experience Sampling for Behaviour Research*, pages 141–161. Springer International Publishing, 2016. ISBN 978-3-319-31413-6. doi: 10.1007/978-3-319-31413-6_8. URL https://doi.org/10.1007/978-3-319-31413-6_8. 3.2.1
- [176] Charith Perera, Arkady Zaslavsky, Peter Christen, and Dimitrios Georgakopoulos. Context aware computing for the internet of things: A survey. *IEEE Communications Surveys & Tutorials*, 16(1):414–454, 2013. 1.1, 2.4
- [177] Charith Perera, Rajiv Ranjan, Lizhe Wang, Samee U Khan, and Albert Y Zomaya. Big data privacy in the internet of things era. *IT Professional*, 17(3):32–39, 2015. 1.1, 2.4
- [178] Pawani Porambage, Mika Ylianttila, Corinna Schmitt, Pardeep Kumar, Andrei Gurtov, and Athanasios V. Vasilakos. The quest for privacy in the internet of things. *IEEE Cloud Computing*, 3(2):36–45, 2016. 1.1, 2.4
- [179] Salil Prabhakar, Sharath Pankanti, and Anil K. Jain. Biometric recognition: Security and privacy concerns. *IEEE Security & Privacy*, 1(2):33–42, March 2003. ISSN 1558-4046. doi: 10.1109/MSECP.2003.1193209. 2.4
- [180] The Canadian Press. Vax pass comparison. a look at covid-19 vaccine certificate programs across canada. <https://www.castanet.net/news/Canada/347887/A-look-at-COV-19-vaccine-certificate-programs-across-Canada>, October 2021. Accessed: 2021-10-07. 4.1, 4.5
- [181] Press Association. Tesco’s plan to tailor adverts via facial recognition stokes privacy fears. <https://www.theguardian.com/business/2013/nov/03/privacy-tesco-scan-c-customers-faces>, November 2013. Accessed: 2020-12-15. 3.2.2, 3.5.4, 3.6.3
- [182] Prolific. Prolific. , 2021. Accessed: 2021-7-08. 4.2, 4.2.2
- [183] Qualtrics. Qualtrics, 2021. URL <https://www.qualtrics.com>. 4.2.2
- [184] Emilee Rader. Most Americans don’t realize what companies can predict from their data. <https://bigthink.com/technology-innovation/most-americans-dont-realize-what-companies-can-predict-from-their-data-2629911919>, February 2019. Accessed: 2020-12-15. 3.2.2
- [185] Edith Ramirez, Julie Brill, Maureen K Ohlhausen, Joshua D Wright, and Terrell McSweeney. Data brokers: A call for transparency and accountability. Technical report, Federal Trade Commission, May 2014. URL <https://www.ftc.gov/system/files/documents/reports/data-brokers-call-transparency-accountability-report-federal-trade-commission-may-2014/140527databrokerreport.pdf>. 2.4

- [186] Luis Felipe M. Ramos. Evaluating privacy during the covid-19 public health emergency: The case of facial recognition technologies. In *Proceedings of the 13th International Conference on Theory and Practice of Electronic Governance*, ICEGOV 2020, page 176–179, New York, NY, USA, 2020. ACM. ISBN 9781450376747. doi: 10.1145/3428502.3428526. URL <https://doi.org/10.1145/3428502.3428526>. 2.4, 3.7.8
- [187] Ashwini Rao, Florian Schaub, Norman Sadeh, Alessandro Acquisti, and Ruogu Kang. Expecting the unexpected: Understanding mismatched privacy expectations online. In *Twelfth Symposium on Usable Privacy and Security (SOUPS 2016)*, pages 77–96, Denver, CO, June 2016. USENIX Association. ISBN 978-1-931971-31-7. URL <https://www.usenix.org/conference/soups2016/technical-sessions/presentation/rao>. 2.1.1
- [188] Bahman Rashidi, Carol Fung, and Tam Vu. Dude, ask the experts!: Android resource access permission recommendation with RecDroid. In *2015 IFIP/IEEE International Symposium on Integrated Network Management (IM)*, pages 296–304, 2015. 2.1.3
- [189] Abhilasha Ravichander, Alan W Black, Shomir Wilson, Thomas Norton, and Norman Sadeh. Question answering for privacy policies: Combining computational and legal perspectives. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4949–4959, Hong Kong, China, November 2019. Association for Computational Linguistics. doi: 10.18653/v1/D19-1500. URL <https://www.aclweb.org/anthology/D19-1500>. 5.4.4, 6.1, 6.2, 6.3, 6.6, 7.3
- [190] Abhilasha Ravichander, Alan W Black, Thomas Norton, Shomir Wilson, and Norman Sadeh. Breaking down walls of text: How can NLP benefit consumer privacy? In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 4125–4140, 2021. 5.4.4, 6.1, 6.6, 7.3
- [191] Ramprasad Ravichandran, Michael Benisch, Patrick Gauge Kelley, and Norman Sadeh. Capturing social networking privacy preferences: Can default policies help alleviate tradeoffs between expressiveness and user burden? In *Proceedings of the 5th Symposium on Usable Privacy and Security*, SOUPS '09, New York, NY, USA, 2009. Association for Computing Machinery. ISBN 9781605587363. doi: 10.1145/1572532.1572587. URL <https://doi.org/10.1145/1572532.1572587>. 2.2
- [192] Robert W. Reeder, Adrienne Porter Felt, Sunny Consolvo, Nathan Malkin, Christopher Thompson, and Serge Egelman. An experience sampling study of user reactions to browser warnings in the field. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18)*, pages 1–13, 2018. 3.2.1
- [193] Joel R Reidenberg, Travis Breaux, Lorrie Faith Cranor, Brian French, Amanda Grannis, James T Graves, Fei Liu, Aleecia McDonald, Thomas B Norton, and Rohan Ramanath. Disagreeable privacy policies: Mismatches between meaning and users' understanding. *Berkeley Tech. LJ*, 30:39, 2015. 2.1.1

- [194] 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. 2.1.1
- [195] Joel R Reidenberg, N Cameron Russell, Vlad Herta, William Sierra-Rocafort, and Thomas B Norton. Trustworthy privacy indicators: Grades, labels, certifications, and dashboards. *Washington University Law Review*, 96:1409, 2018. 2.1.1, 2.1.2
- [196] Daniel Reinhardt, Johannes Borchard, and Jörn Hurtienne. Visual interactive privacy policy: The better choice? In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, CHI '21, New York, NY, USA, 2021. Association for Computing Machinery. ISBN 9781450380966. doi: 10.1145/3411764.3445465. URL <https://doi.org/10.1145/3411764.3445465>. 5.4.1
- [197] Reuters Staff. German restaurants object after police use covid data for crime-fighting. <https://www.reuters.com/article/us-health-coronavirus-germany-privacy/german-restaurants-object-after-police-use-covid-data-for-crime-fighting-idUSKCN24W2K6>, 2020. 2.3, 4.2.1
- [198] Timothy Revell. Computer vision algorithms pick out petty crime in CCTV footage. <https://www.newscientist.com/article/2116970-computer-vision-algorithms-pick-out-petty-crime-in-cctv-footage/>, January 2017. Accessed: 2020-12-15. 3.2.2
- [199] Hannah Ritchie, Esteban Ortiz-Ospina, Diana Beltekian, Edouard Mathieu, Joe Hasell, Bobbie Macdonald, Charlie Giattino, Cameron Appel, Lucas Rodés-Guirao, and Max Roser. Coronavirus pandemic (covid-19), 2020. <https://ourworldindata.org/coronavirus>. 4.2.2, 4.2.4
- [200] David Rosen. Disney is spying on you! https://www.salon.com/test/2013/01/17/disney_is_spying_on_you/, January 2013. Accessed: 2020-12-15. 3.2.2
- [201] 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. Frontiers*, 252: 181–195, 2019. 2.1.2
- [202] Norman Sadeh. Design of a privacy infrastructure for the internet of things. In *2020 USENIX Conference on Privacy Engineering Practice and Respect (PEPR 20)*. USENIX Association, 2020. URL <https://www.usenix.org/conference/pepr20/presentation/sadeh>. 2.4, 3.1, 3.7.8
- [203] T. Scott Saponas, Jonathan Lester, Carl Hartung, Sameer Agarwal, and Tadayoshi Kohno. Devices that tell on you: Privacy trends in consumer ubiquitous computing. In *16th USENIX Security Symposium (USENIX Security '07)*, pages 55–70, 2007. 2.4
- [204] Florian Schaub, Rebecca Balebako, Adam L Durity, and Lorrie Faith Cranor. A design space for effective privacy notices. In *Proceedings of the Eleventh Symposium On Usable Privacy and Security (SOUPS '15)*, pages 1–17, 2015. 2.1.3
- [205] Florian Schaub, Bastian Könings, and Michael Weber. Context-adaptive privacy: Leveraging context awareness to support privacy decision making. *IEEE Pervasive Computing*, 14(1): 34–43, 2015. 2.1.1, 2.1.2

- [206] Florian Schroff, Dmitry Kalenichenko, and James Philbin. FaceNet: A unified embedding for face recognition and clustering. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 815–823, 2015. doi: 10.1109/CVPR.2015.7298682. 2.4
- [207] Emily Schultheis and Geir Moulson. Austrian parliament approves vaccine mandate for adults. <https://apnews.com/article/austrian-parliament-covid-vaccine-mandate-8539164285f87443a8b80a213d2dacc0>, January 2022. Accessed: 2022-01-21. 4.5
- [208] E. J. Schultz. Facial-recognition lets marketers gauge consumers’ real responses to ads. <https://adage.com/article/digital/facial-recognition-lets-marketers-gauge-real-responses/298635>, May 2015. Accessed: 2020-12-15. 3.2.2
- [209] Secretary’s Advisory Committee on Automated Personal Data Systems. *Records, Computers, and the Rights of Citizens: Report*. US Department of Health, Education & Welfare, 1973. 2.1
- [210] Shawn Shan, Emily Wenger, Jiayun Zhang, Huiying Li, Haitao Zheng, and Ben Y. Zhao. Fawkes: Protecting privacy against unauthorized deep learning models. In *29th USENIX Security Symposium (USENIX Security ’20)*, pages 1589–1604, August 2020. ISBN 978-1-939133-17-5. URL <https://www.usenix.org/conference/usenixsecurity20/presentation/shan>. 2.4
- [211] Mahmood Sharif, Sruti Bhagavatula, Lujo Bauer, and Michael K. Reiter. Accessorize to a crime: Real and stealthy attacks on state-of-the-art face recognition. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security (CCS ’16)*, pages 1528–1540, 2016. ISBN 9781450341394. doi: 10.1145/2976749.2978392. URL <https://doi.org/10.1145/2976749.2978392>. 2.4
- [212] Fuming Shih, Ilaria Liccardi, and Daniel Weitzner. Privacy tipping points in smartphones privacy preferences. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI ’15)*, pages 807–816, 2015. 2.1.3, 2.2
- [213] Yan Shvartzshnaider, Schrasing Tong, Thomas Wies, Paula Kift, Helen Nissenbaum, Lakshminarayanan Subramanian, and Prateek Mittal. Learning privacy expectations by crowdsourcing contextual informational norms. *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, 4(1):209–218, Sep. 2016. URL <https://ojs.aaai.org/index.php/HCOMP/article/view/13271>. 2.3, 4.2.1, 4.2.4, 4.3.3
- [214] Ed Silverstein. New Konami casino facial recognition technology could rival reward cards. <https://www.casino.org/news/new-konami-casino-facial-recognition-technology-could-rival-reward-cards/>, October 2019. Accessed: 2020-12-15. 3.2.2
- [215] Ravi Inder Singh, Manasa Sumeeth, and James Miller. A user-centric evaluation of the readability of privacy policies in popular web sites. *Information Systems Frontiers*, 13(4): 501–514, 2011. 2.1.1

- [216] Arron Smith. More than half of U.S. adults trust law enforcement to use facial recognition responsibly. Technical report, Pew Research Center, September 2019. URL <https://www.pewresearch.org/internet/2019/09/05/more-than-half-of-u-s-adults-trust-law-enforcement-to-use-facial-recognition-responsibly/>. 3.1, 3.7.7
- [217] Daniel Smullen, Yuanyuan Feng, Shikun Zhang, and Norman M. Sadeh. The best of both worlds: Mitigating trade-offs between accuracy and user burden in capturing mobile app privacy preferences. *Proc. Priv. Enhancing Technol.*, 2020(1):195–215, 2020. doi: 10.2478/popets-2020-0011. URL <https://doi.org/10.2478/popets-2020-0011>. 2.1.3, 3.7.3
- [218] Daniel Smullen, Yuanyuan Feng, Shikun Aerin Zhang, and Norman Sadeh. The best of both worlds: Mitigating trade-offs between accuracy and user burden in capturing mobile app privacy preferences. *Proceedings on Privacy Enhancing Technologies*, 2020(1):195–215, 2020. 5.4.4
- [219] Benjamin Snyder. This beer ad only works when women pass by. <https://fortune.com/2015/05/21/astra-beer-ad/>, May 2015. Accessed: 2020-12-15. 3.2.2, 3.5.4, 3.6.3
- [220] Daniel J. Solove. A taxonomy of privacy. *University of Pennsylvania Law Review*, 154(3): 477–564, 2006. ISSN 00419907. URL <http://www.jstor.org/stable/40041279>. 3.5.5, 3.5.5, 3.5.5, 3.5.5, 3.5.5, 3.7.7
- [221] Luke Stark, Amanda Stanhaus, and Denise L. Anthony. “I Don’t Want Someone to Watch Me While I’m Working”: Gendered Views of Facial Recognition Technology in Workplace Surveillance. *Journal of the Association for Information Science and Technology*, 71(9): 1074–1088, 2020. doi: <https://doi.org/10.1002/asi.24342>. URL <https://asistdl.onlinelibrary.wiley.com/doi/abs/10.1002/asi.24342>. 3.7.7
- [222] Statista. Global smartphone penetration rate as share of population from 2016 to 2020. <https://www.statista.com/statistics/203734/global-smartphone-penetration-per-capita-since-2005/>, Dec 2021. 1.1
- [223] Statista. Time spent with nonvoice activities on mobile phones every day in the United States from 2019 to 2023. <https://www.statista.com/statistics/1045353/mobile-device-daily-usage-time-in-the-us/>, Apr 2021. 1.1
- [224] Statista. Number of Internet of Things (IoT) connected devices worldwide from 2019 to 2030, by vertical. <https://www.statista.com/statistics/1194682/iot-connected-devices-vertically/>, Mar 2022. 1.1
- [225] Statista. Average number of mobile apps opened per month on mobile devices in the United States from 2019 to 2021. <https://www.statista.com/statistics/1288955/us-apps-opened-per-month/>, Feb 2022. 1.1
- [226] Statista. Percentage of U.S. adults who own a smartphone from 2011 to 2021. <https://www.statista.com/statistics/219865/percentage-of-us-adults-who-own-a-smartphone/>, May 2022. 1.1

- [227] Statista. Share of global daily time spent online via mobile devices 2013-2021. <https://www.statista.com/statistics/1289723/share-time-spent-mobile-internet-daily/>, May 2022. 1.1
- [228] Nili Steinfeld. “I agree to the terms and conditions”: (How) do users read privacy policies online? An eye-tracking experiment. *Computers in human behavior*, 55:992–1000, 2016. 2.1.1
- [229] Steve Stemler. An overview of content analysis. *Practical assessment, research, and evaluation*, 7(1):17, 2000. 3.3.1
- [230] Francesca Street. How facial recognition is taking over airports, 2019. URL <https://www.cnn.com/travel/article/airports-facial-recognition/index.html>. 3.1
- [231] Ryota Suzuki and Hidetoshi Shimodaira. Pvclust: an R package for assessing the uncertainty in hierarchical clustering. *Bioinformatics*, 22(12):1540–1542, 04 2006. ISSN 1367-4803. doi: 10.1093/bioinformatics/btl117. URL <https://doi.org/10.1093/bioinformatics/btl117>. 4.2.1, 4.4
- [232] Zoe Tabary. Expert views-what role for vaccine passports in coronavirus fight? <https://www.reuters.com/article/health-coronavirus-global-tech-idUSL5N2KZ5GR>, 2021. Accessed: 2021-4-09. 1.1
- [233] Yaniv Taigman, Ming Yang, Marc’ Aurelio Ranzato, and Lior Wolf. Deepface: Closing the gap to human-level performance in face verification. *2014 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1701–1708, 2014. 2.4
- [234] Prolific Team. Representative samples on prolific. <https://researcher-help.prolific.co/hc/en-gb/articles/360019236753-Representative-Samples-on-Prolific>, March 2019. Accessed: 2021-7-08. 4.2.2
- [235] Welderufael B Tesfay, Peter Hofmann, Toru Nakamura, Shinsaku Kiyomoto, and Jetzabel Serna. Privacyguide: towards an implementation of the EU GDPR on internet privacy policy evaluation. In *Proceedings of the Fourth ACM International Workshop on Security and Privacy Analytics*, pages 15–21, 2018. 5.4.4
- [236] Ma. Dolores C. Tongco. Purposive sampling as a tool for informant selection. *Ethnobotany Research and Applications*, 5:147–158, 2007. 5.2.1
- [237] Christine Utz, Martin Degeling, Sascha Fahl, Florian Schaub, and Thorsten Holz. (un)informed consent: Studying gdpr consent notices in the field. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security, CCS ’19*, page 973–990, New York, NY, USA, 2019. Association for Computing Machinery. ISBN 9781450367479. doi: 10.1145/3319535.3354212. URL <https://doi.org/10.1145/3319535.3354212>. 1.1

- [238] Christine Utz, Steffen Becker, Theodor Schnitzler, Florian M. Farke, Franziska Herbert, Leonie Schaewitz, Martin Degeling, and Markus Dürmuth. Apps against the spread: Privacy implications and user acceptance of covid-19-related smartphone apps on three continents. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, CHI '21, New York, NY, USA, 2021. Association for Computing Machinery. ISBN 9781450380966. doi: 10.1145/3411764.3445517. URL <https://doi.org/10.1145/3411764.3445517>. 2.3
- [239] 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. 2.1.1
- [240] Niels Van Berkel, Denzil Ferreira, and Vassilis Kostakos. The experience sampling method on mobile devices. *ACM Computing Surveys (CSUR)*, 50(6):1–40, 2017. 3.2.1
- [241] Simone JW Verhagen, Laila Hasmi, Marjan Drukker, Jim van Os, and Philippe AEG Delespaul. Use of the experience sampling method in the context of clinical trials. *Evidence-based Mental Health*, 19(3):86–89, 2016. 3.2.1
- [242] Tony Vila, Rachel Greenstadt, and David Molnar. Why we can't be bothered to read privacy policies. In *Economics of information security*, pages 143–153. Springer, 2004. 2.1.1
- [243] Kim-Phuong L Vu, Vanessa Chambers, Fredrick P Garcia, Beth Creekmur, John Sulaitis, Deborah Nelson, Russell Pierce, and Robert W Proctor. How users read and comprehend privacy policies. In *Symposium on Human Interface and the Management of Information*, pages 802–811. Springer, 2007. 2.1.1
- [244] Samuel D. Warren and Louis D. Brandeis. The right to privacy. *Harvard Law Review*, 4(5): 193–220, 1890. URL <http://faculty.uml.edu/sgallagher/Brandeisprivacy.htm>. 3.5.5, 3.7.7
- [245] Alan F. Westin. Privacy and freedom. *Washington & Lee Law Review*, 25:166, 1968. 1.1, 3.5.5, 3.7.7
- [246] Jason Whitely. How facial recognition technology is being used, from police to a soccer museum. Online, November 2018. URL <https://www.wfaa.com/article/features/originals/how-facial-recognition-technology-is-being-used-from-police-to-a-soccer-museum/287-618278039>. Accessed: 2020-12-15. 3.2.2
- [247] Primal Wijesekera, Arjun Baokar, Lynn Tsai, Joel Reardon, Serge Egelman, David Wagner, and Konstantin Beznosov. The feasibility of dynamically granted permissions: Aligning mobile privacy with user preferences. In *2017 IEEE Symposium on Security and Privacy*, pages 1077–1093, 2017. 2.2
- [248] Primal Wijesekera, Joel Reardon, Irwin Reyes, Lynn Tsai, Jung-Wei Chen, Nathan Good, David Wagner, Konstantin Beznosov, and Serge Egelman. Contextualizing privacy decisions for better prediction (and protection). In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, CHI '18, page 1–13, New York, NY, USA, 2018. Association for Computing Machinery. ISBN 9781450356206. doi: 10.1145/3173574.3173842. URL <https://doi.org/10.1145/3173574.3173842>. 5.4.4

- [249] Primal Wijesekera, Joel Reardon, Irwin Reyes, Lynn Tsai, Jung-Wei Chen, Nathan Good, David Wagner, Konstantin Beznosov, and Serge Egelman. Contextualizing privacy decisions for better prediction (and protection). In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18)*, pages 1–13, 2018. 2.1.3, 2.2
- [250] Zuxuan Wu, Ser-Nam Lim, Larry S. Davis, and Tom Goldstein. Making an invisibility cloak: Real world adversarial attacks on object detectors. In Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm, editors, *Computer Vision – ECCV 2020*, pages 1–17. Springer International Publishing, 2020. ISBN 978-3-030-58548-8. 2.4
- [251] Huijuan Xu, Abir Das, and Kate Saenko. R-C3D: Region convolutional 3d network for temporal activity detection. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2017. 2.4
- [252] Baobao Zhang, Laurin Weissinger, Johannes Himmelreich, Nina McMurry, Tiffany Li, Naomi Schinerman, and Sarah Kreps. Building robust and ethical vaccination verification systems. <https://www.brookings.edu/techstream/building-robust-and-ethical-vaccination-verification-systems/>, January 2021. 4.2.2
- [253] Shikun Zhang and Norman Sadeh. Do privacy labels answer users’ privacy questions? In *Symposium on Usable Security and Privacy (USEC) 2023*, Feb 2023. ISBN 1-891562-91-6. URL <https://dx.doi.org/10.14722/usec.2023.232482>. 6.7
- [254] Shikun Zhang, Yuanyuan Feng, Lujo Bauer, Lorrie Faith Cranor, Anupam Das, and Norman M Sadeh. " did you know this camera tracks your mood?": Understanding privacy expectations and preferences in the age of video analytics. *Proc. Priv. Enhancing Technol.*, 2021(2):282–304, 2021. 4.2.4
- [255] Shikun Zhang, Yuanyuan Feng, and Norman Sadeh. Facial recognition: Understanding privacy concerns and attitudes across increasingly diverse deployment scenarios. In *Seventeenth Symposium on Usable Privacy and Security (SOUPS 2021)*, pages 243–262. USENIX Association, August 2021. ISBN 978-1-939133-25-0. URL <https://www.usenix.org/conference/soups2021/presentation/zhang-shikun>. 3.1, 3.8
- [256] Shikun Zhang, Yuanyuan Feng, Yaxing Yao, Lorrie Faith Cranor, and Norman Sadeh. How usable are ios app privacy labels. *Proc. Priv. Enhancing Technol.*, 2022(4):204–228, 2022. doi: 10.2478/popets-2022-0106. URL https://privacyassistant.org/media/publications/privacy_labels.pdf. 5.1, 6.5.4, 6.7
- [257] Shikun Zhang, Yan Shvartzshnaider, Yuanyuan Feng, Helen Nissenbaum, and Norman Sadeh. Stop the spread: A contextual integrity perspective on the appropriateness of covid-19 vaccination certificates. In *2022 ACM Conference on Fairness, Accountability, and Transparency, FAccT '22*, page 1657–1670, New York, NY, USA, 2022. Association for Computing Machinery. ISBN 9781450393522. doi: 10.1145/3531146.3533222. URL <https://doi.org/10.1145/3531146.3533222>. 4.1, 4.7

- [258] Shikun Aerin Zhang, Yuanyuan Feng, Anupam Das, Lujio Bauer, Lorrie Cranor, and Norman Sadeh. Understanding people's privacy attitudes towards video analytics technologies. Technical Report CMU-ISR-20-114, Carnegie Mellon University, School of Computer Science, December 2020. 3.7.4, 3.8
- [259] Shikun Aerin Zhang, Yuanyuan Feng, Anupam Das, Lujio Bauer, Lorrie Cranor, and Norman Sadeh. "Did you know this camera tracks your mood?": Understanding privacy expectations and preferences in the age of video analytics. *Proc. Priv. Enhancing Technol.*, 2021(2): 282–304, 2021. doi: 10.2478/popets-2021-0028. URL <https://doi.org/10.2478/popets-2021-0028>. 1.1, 3.1, 3.8
- [260] Sebastian Zimmeck, Peter Story, Daniel Smullen, Abhilasha Ravichander, Ziqi Wang, Joel Reidenberg, N Cameron Russell, and Norman Sadeh. Maps: Scaling privacy compliance analysis to a million apps. *Proceedings on Privacy Enhancing Technologies*, 2019(3):66–86, 2019. 5.4.4