

Syracuse University

SURFACE

Dissertations - ALL

SURFACE

May 2020

Privacy For Whom? A Multi-Stakeholder Exploration of Privacy Designs

Yaxing Yao
Syracuse University

Follow this and additional works at: <https://surface.syr.edu/etd>



Part of the [Social and Behavioral Sciences Commons](#)

Recommended Citation

Yao, Yaxing, "Privacy For Whom? A Multi-Stakeholder Exploration of Privacy Designs" (2020).
Dissertations - ALL. 1181.
<https://surface.syr.edu/etd/1181>

This Dissertation is brought to you for free and open access by the SURFACE at SURFACE. It has been accepted for inclusion in Dissertations - ALL by an authorized administrator of SURFACE. For more information, please contact surface@syr.edu.

Abstract

Privacy is considered one of the fundamental human rights. Researchers have been investigating privacy issues in various domains, such as our physical privacy, data privacy, privacy as a legal right, and privacy designs. In the Human-Computer Interaction field, privacy researchers have been focusing on understanding people's privacy concerns when they interact with computing systems, designing and building privacy-enhancing technologies to help people mitigate these concerns, and investigating how people's privacy perceptions and the privacy designs influence people's behaviors.

Existing privacy research has been overwhelmingly focusing on the privacy needs of end-users, i.e., people who use a system or a product, such as Internet users and smartphone users. However, as our computing systems are becoming more and more complex, privacy issues within these systems have started to impact not only the end-users but also other stakeholders, and privacy-enhancing mechanisms designed for the end-users can also affect multiple stakeholders beyond the users.

In this dissertation, I examine how different stakeholders perceive privacy-related issues and expect privacy designs to function across three application domains: online behavioral advertising, drones, and smart homes. I choose these three domains because they represent different multi-stakeholder environments with varying nature of complexity. In particular, these environments

present the opportunities to study technology-mediated interpersonal relationships, i.e., the relationship between primary users (owners, end-users) and secondary users (bystanders), and to investigate how these relationships influence people's privacy perceptions and their desired ways of privacy protection.

Through a combination of qualitative, quantitative, and design methods, including interviews, surveys, participatory designs, and speculative designs, I present how multi-stakeholder considerations change our understandings of privacy and influence privacy designs. I draw design implications from the study results and guide future privacy designs to consider the needs of different stakeholders, e.g., cooperative mechanisms that aim to enhance the communication between primary and secondary users.

In addition, this methodological approach allows researchers to directly and proactively engage with multiple stakeholders and explore their privacy perceptions and expected privacy designs. This is different from what has been commonly used in privacy literature and as such, points to a methodological contribution.

Finally, this dissertation shows that when applying the theory of Contextual Integrity in a multi-stakeholder environment, there are hidden contextual factors that may alter the contextual informational norms. I present three examples from the study results and argue that it is necessary to carefully examine such factors in order to clearly identify the contextual norms. I propose a research agenda to explore best practices of applying the theory of Contextual Integrity in a multi-stakeholder environment.

PRIVACY FOR WHOM? A MULTI-STAKEHOLDER EXPLORATION
OF PRIVACY DESIGNS

by

Yaxing Yao

B.A., Harbin Institute of Technology, 2012

M.S., University of Washington, 2014

Dissertation

Submitted in partial fulfillment of the requirements for the degree of
Doctor of Philosophy in Information Science & Technology.

Syracuse University
May 2020

Copyright

Portion of Chapter 3[®] 2017 ACM

Portion of Chapter 5[®] 2016 De Gruyter

Portion of Chapter 6[®] 2017 ACM

Portion of Chapter 7[®] 2017 ACM

Portion of Chapter 8[®] 2019 ACM

Portion of Chapter 9[®] 2019 ACM

All other materials[®] 2020 Yaxing Yao

Acknowledgments

My experience in the past four years and a half at Syracuse has been invaluable and unforgettable. The process of attempting to get a Ph.D., the same as for many other people, has not been easy for me. When I first started my Ph.D., I barely knew research as well as the field of human-computer interaction and privacy. Fortunately, I have gained tremendous support from many people over the years, who have made my journey joyful, fruitful, and full of memories. I feel deeply indebted to all of them.

The first and great thanks go to my advisor, Dr. Yang Wang. Since the beginning, he has been supporting me in everything that I can think of within and beyond academics. As my advisor, his insights and wisdom have been an invaluable source for me. He understands the problem, the field, and the community, and is always able to point me in the right direction when I am confused. Only after a few years into my Ph.D. did I realize how privileged I am to be able to work with him. More importantly, he always stands by me as a mentor. He is extremely caring and supportive and has been mentoring me through many life-changing transitions beyond academics, such as being a father. He is also a trustworthy friend, and I can talk to him about everything. I appreciate all that he has done for me. I feel very fortunate to have worked with him during my Ph.D., and I will miss these days.

I thank all my committee members. I have known Dr. Jason Dedrick, Dr. Joon Park, and Dr. Bryan Semaan since I started my Ph.D. They have been mentoring me as instructors, collaborators, and committee members throughout my doctoral study. Particularly in the past year, the guidance and instructions they have given me have helped me through the development of this dissertation, and I am grateful to that. I thank Dr. Joshua Introne, who spends so much time to help me with the document development and the dissertation defense. I am very honored to have Dr. Florian Schaub on the committee. My past research, including the work presented in this dissertation, is very relevant to his research interest, and I look forward to working with him in the years ahead. I also thank Dr. Jamie Winders for agreeing to be the chair of my defense out of her busy schedule.

I thank all my teachers and professors throughout my life. In particular, I am forever grateful to the mentorship during my master's program by Dr. Hans J Scholl (UW), Dr. Robert Mason (UW), Dr. Ricardo Gomez (UW), Dr. Joshua Blumenstock (UC Berkeley), and Dr. Emma Spiro (UW). They have led me through the door of research and supported me during my Ph.D. application. Without their help, I would not even start my life as a researcher.

Over the years, I have received tremendous help from many faculty at Syracuse. I cannot thank them enough for all the guidance and mentorship they have been giving me. I thank Richel Clark, Kevin Crowston, Ingrid Erickson, Caroline Haythornthwaite, Jeff Hemsley, Yun Huang (now at UIUC), Jeff Saltz, Steve Sawyer, David Seaman, Arthur Thomas, Lu Xiao, and Bei Yu. I would also like to thank the staff members who have made my student life much more manageable. I thank Jennifer Barclay, Susan Nemier, Lois Elmore, Meghan Macblane, Mariann Major, Roger Merrill, Susan Nemier, Maureen O'Connor Kicak, Jennifer Pulver, Kevin Shults, and Jose Tavaréz.

I have also been very fortunate to work with many brilliant students at Syracuse and publish with them. They have taught me valuable lessons, particularly in how to be a mentor myself. I

thank my co-authors Qunfang Wu, Huichuan Xia, Sen Huo, Chao Niu, Jordan Hayes, Nata Barbosa, Yisi Sang, Charlotte Price, Justin Basedo, Oriana Mcdonough, and Smirity Kaushik. In particular, Justin Basedo, Oriana Mcdonough, and Smirity Kaushik are three amazing students at Syracuse University who have been working with me for more than two years. I am deeply touched and motivated by their diligence and hard work. I wish them all the best in their new careers.

I want to give special thanks to all the teachers, staff, and interns at the Syracuse University Early Education & Child Care Center for taking care of my son, Christopher. Their caring and kindness have made Christopher's early education experience colorful and enjoyable. I am very grateful to them.

Collaboration is one of the most exciting and rewarding experiences in academia and one of the most important lessons I have learned as a graduate student. Outside of Syracuse, I have the opportunity to work with amazing scholars from over 20 universities and institutions. They have been teaching me different values and opinions in research, giving me advice and feedback on career development and job hunting, and guiding me to be a better scholar. I look forward to continuing working with them in the years ahead. I thank Richmond Wong, David Lo Re, Karla Badillo-Urquiola, Xinru Page, Pamela Wisniewski, Ashley Walker, Christine Geeng, Roberto Hoyle, Yuhang Zhao, Zhicong Lu, Nick Merrill, Pardis Emami-Naeini, Shruti Sannon, Xinning Gui, Lu Zhang, Ruojia Wang, YuanYuan Feng, Peter Story, Daniel Smullen, Justin Donnell, Toby Jia-Jun Li, Haojian Jin, Justin Weisz, and Jill Woelfer. I also thank Dr. Norman Sadeh for offering me a fantastic opportunity as a Post Doctoral Fellow in the School of Computer Science at Carnegie Mellon University. I look forward to working with him in the next couple of years.

Besides academics, I was surrounded by many, many great friends who have been an integral part of my life. I want to express my greatest gratitude to them. I thank Erin Bartolo, Sarah

Bratt, Mahboobeh Harandi, Yingya Li, Erica Mitchell, Jean Philippe Rancy, Ehsan Sabaghian, Ivan Shamshurin, Sarika Kumari Sharma, Alain Shema, Ellen Simpson, Jennifer Sonne, Feifei Zhang, Han Zhuang, Brian Dobreski, Bryan Dosono, Corey Jackson, Sikana Tanupabrungsun, Jerry Robinson, and all other students in the Ph.D. program in the iSchool at Syracuse University. I also thank all my friends whom I have met during my time at UW and back in China for being my friends and supporting me all the time.

I give my deep-hearted and most sincere thanks to my father, Gewen Yao, and my mother, Xiaowei Feng. Like most Chinese people who generally do not express their emotions and feelings explicitly, my parents rarely express their love verbally. However, everything they have done is a statement of their love, which has been the strongest motivation for me to keep pushing myself and moving forward because I know they are always behind me. They have always believed in me from the very beginning and witnessed my growth along the way. They make me who I am today, even though I can only stay with them for a few weeks each year since I was 14 (I went to high school in a different city). Their sacrifices are beyond imagination and my debts to them are beyond measure. There is certainly no way for me to pay them back, but I hope that what I have been through so far can at least make them somewhat proud.

Although she is no longer with us, I know my grandma is watching me. She passed away during my Ph.D. For many reasons, I was not able to go back to China and attend her funeral. Still, I know she has given me courage and integrity, which have been motivating me ever since to achieve my goal. I thank her for all the years she has been with me, along with all my other family members who have always been by my side.

Finally, I would like to thank my best friend in the world, the rock of our family, and the love of my life, Xinyi Zhang. We have known each other for almost fifteen years and been married for

nearly seven years. Ever since we met at high school, she has been the most significant part of my life as a friend, lover, life partner, as well as my unyielding supporter in every way possible. She gives me advice when I lost; she hugs me when I am down; she listens to me when I need to talk; and after all that, she cheers with me when I succeed. She has also been an amazing mom since we had our son, Christopher Yao, and daughter, Iris Yao (they were both born during my Ph.D. study). Her dedication in taking care of the children is beyond my words and I thank her for everything she has done for our family. I admire her for her perseverance, sacrifice, caring, and love throughout. Without her, I will not stand where I am today.

To my parents, my wife, my son, and my daughter, for their unconditional love.

Contents

Abstract	i
Copyright	iv
Acknowledgments	v
I Introduction	1
1 Introduction	2
1.1 Document Organization	5
1.2 Overview	6
1.3 Major Contributions	6
2 Literature Review	8
2.1 What is Privacy?	8
2.2 Theory of Contextual Integrity	11
2.3 Stakeholder Theory	13
2.4 Multi-Stakeholder Privacy Research	15

2.5	Research Question	16
II	Case 1: Online Behavior Advertising (OBA)	18
3	Online Privacy: Users' Understandings of OBA	19
3.1	Introduction	20
3.2	Related Work	23
3.2.1	People's Attitudes and Perceptions of OBA	23
3.2.2	People's Mental Models of Privacy and Security	25
3.3	Method	27
3.3.1	Pilot Study	27
3.3.2	First-Round Interviews	28
3.3.3	Second-Round Interviews	30
3.3.4	Participant Recruitment	31
3.3.5	Data Analysis	32
3.4	Results	33
3.4.1	Participants	33
3.4.2	Folk Models of OBA	34
3.4.3	Browser-Pull Model	35
3.4.4	First-Party-Pull Model	36
3.4.5	Connected-First-Party Model	39
3.4.6	Third-Party Model	42
3.4.7	Misconceptions and Speculations of OBA	45

3.4.8	Privacy-Enhancing Tools for OBA	47
3.5	Discussion	50
3.5.1	Why Folk Models of OBA Matter	51
3.5.2	Implications for Design and Policy	53
3.5.3	Limitations and Future Research	55
3.5.4	Conclusion	56
3.5.5	Acknowledgement	56
4	Online Privacy: Implications Beyond Users' Privacy	57
4.1	Introduction	57
4.2	Technical Background	59
4.3	Related Work	60
4.3.1	Socio-technical Countermeasures for Web Tracking	60
4.3.2	Speculative Design for Privacy Research	63
4.4	Iterative Processes of Speculation	65
4.4.1	Step 1 - Explore The Tracking Blocker Design Space	65
4.4.2	Step 2 - Initial Interface Designs	68
4.4.3	Step 3 - Crafting Scenarios Around the Interfaces	69
4.4.4	Step 4 - Critically Examine the Speculations	71
4.5	Speculative Scenarios	73
4.5.1	Scenario 1: Invisible Blocker	73
4.5.2	Scenario 2: Balanced Web Browser	74
4.5.3	Scenario 3: Greedy Web Tracker	75

4.6	Reflections of Our Speculative Scenarios	76
4.6.1	Commentary: Technological Reflections	77
4.6.2	Commentary: Legal Reflections	78
4.6.3	Commentary: Social Reflections	79
4.6.4	Commentary: Economic Reflections	80
4.6.5	Commentary: Users' Digital Literacy	80
4.6.6	Implications and Questions for Blocking Technologies Adoption	81
4.7	Discussion	83
4.7.1	Reflecting Our Speculative Exploration Process	83
4.7.2	Speculative Design of PETS	84
4.8	Conclusion	86
 III Case 2: Drones		87
 5 Drone Privacy: Bystanders' Perspectives		88
5.1	Study Design	88
5.2	Data Analysis	89
5.3	Results Summary	89
 6 Drone Privacy: Controllers' Perspectives		92
6.1	Introduction	92
6.2	Related Work	93
6.3	Methodology	94
6.4	Findings	96

6.4.1	Privacy Perceptions	97
6.5	Discussion	101
6.6	Conclusion	104
7	Privacy Mechanisms for Drones: Perceptions of Drone Controllers and Bystanders	105
7.1	Introduction	105
7.2	Related Work	107
7.2.1	Perceptions of Tracking and Recording Technologies	107
7.2.2	Privacy Issues of Drones	108
7.2.3	Privacy Mechanisms for Drones	109
7.3	Methodology	109
7.3.1	Survey Flow	110
7.3.2	Survey One	111
7.3.3	Survey Two	112
7.3.4	Data Analysis	117
7.4	Results	117
7.4.1	Results of Survey One	117
7.4.2	Results of Survey Two	119
7.5	Discussion	129
7.5.1	Perceptions and Preferences of Privacy Mechanisms	129
7.5.2	Important Questions for Addressing Drone Privacy Issues	130
7.5.3	Study Limitations	135
7.6	Conclusion	136

7.7	Acknowledgements	136
IV	Case 3: Smart Home	137
8	Smart Home Privacy: Users' Perspectives	138
8.1	Introduction	138
8.2	Related Work	140
8.2.1	Smart Home Privacy Concerns and Risks	140
8.2.2	Smart Home Privacy Mechanisms	141
8.3	Method	143
8.3.1	Participants	143
8.3.2	Session 1	144
8.3.3	Session 2	146
8.3.4	Data Analysis	146
8.4	Results	147
8.4.1	Data Transparency and Control	148
8.4.2	Security	152
8.4.3	Safety	154
8.4.4	Usability and User Experiences	155
8.4.5	System Intelligence	156
8.4.6	System Modality	157
8.5	Discussion	158
8.5.1	Smart Home Privacy	158

8.5.2	Design Implications	159
8.5.3	Policy Implications	162
8.5.4	Reflections on Participants' Privacy Designs	162
8.5.5	Limitations of Our Research	163
8.5.6	Future Directions	164
8.6	Conclusion	165
8.7	Acknowledgement	165
9	Smart Home Privacy: Bystanders' Perspectives	166
9.1	Introduction	166
9.2	Related Work	168
9.2.1	Smart Home Privacy Risks and Concerns	169
9.2.2	Bystanders' Privacy Concerns	170
9.2.3	Smart Home Privacy Mechanisms	171
9.2.4	Gap in the Literature	172
9.3	Method	173
9.3.1	Study Settings	173
9.3.2	Study Flow	175
9.3.3	Data Analysis	178
9.4	Results	179
9.4.1	Participants' General Perceptions	179
9.4.2	Three Aspects of Bystanders' Perceptions	180
9.4.3	Privacy-Seeking Behaviors	185

9.4.4	Privacy Designs	186
9.4.5	Cooperative Mechanisms	187
9.4.6	Bystander-Centric Mechanisms	189
9.5	Discussion	193
9.5.1	Comparing Bystanders' and Users' Privacy Perceptions	194
9.5.2	Unpacking Bystanders' Privacy Designs	196
9.5.3	Design Implications	200
9.5.4	Limitations and Future Work	202
9.6	Conclusion	203
9.7	Acknowledgement	204

V Discussion 205

10 Summary of Results 206

10.1	Synthesis of Results	207
10.1.1	Commonalities among users in three domains	208

11 Implications 211

11.1	Design Implications	211
11.2	Methodological Contribution	213
11.3	Theoretical Contribution	214
11.3.1	When is CI theory promising?	216
11.3.2	When is CI theory challenging?	218
11.3.3	Research Agenda	225

11.4 Limitations	227
12 Conclusion	228
Appendix A	230
Appendix B	250
Bibliography	256
Curriculum Vitae	289

List of Tables

3.1	Participants used three factors in reasoning about OBA and constructing their folk models.	32
3.2	Participants' folk models and attitudes of trackers and OBA.	34
4.1	Different factors in the three speculative scenarios	71
5.1	Summary of the privacy perceptions of drone bystanders	91
7.1	The pros and cons of each mechanism suggested by controllers and bystanders. . .	120
8.1	The six factors identified from our participants' designs of smart home privacy controls	149
9.1	Summary of participants' demographics	174
9.2	Summary of bystanders' privacy design factors, organized based on design purposes and large categories.	187
10.1	Synthesis of results across all three study domains	210
11.1	Examples through the lens of the CI theory	217
11.2	An example from the drone study	223

11.3 An example from the airbnb scenario 223

11.4 An example from the friends visiting scenario 224

12.1 Full list of participants' perceived benefits of smart homes 250

List of Figures

3.1	Browser-pull model: an example from P9.	35
3.2	First-party-pull model: an example from P10.	37
3.3	First-party-pull model: an example from P16.	38
3.4	Connected-first-party model: an example from P6.	40
3.5	Third-party model: an example from P4.	42
3.6	Third-party model: an example from P19.	44
4.1	Our three interactive prototypes.	64
4.2	The icon of ProtectionPlus in Scenario 1 (a), and the two interfaces of Limestone in Scenario 2 (b,c).	72
4.3	The main interfaces of StopInfo in Scenario 3.	74
7.1	Survey one results.	119
7.2	Survey two results.	122
8.1	A photo was taken during one study session.	139
8.2	The flow of our co-design study, including its various components.	142
8.3	The online and offline modes of security cameras (P3)	153

9.1	The app design by P4.	189
9.2	The signal blocker designed by P10.	192

Part I

Introduction

Chapter 1

Introduction

Privacy research has been primarily centered around end-users, trying to understand users' privacy perceptions, concerns, needs, and expectations. As a result, privacy-enhancing mechanisms that aim to mitigate people's privacy concerns have also been designed around end-users. However, as the existing socio-technical systems are becoming more and more complex, the privacy issues associated with these systems have impacted people beyond the end-users. For example, in the domain of the Internet of Things (IoT), when IoT devices collect data from their end-users, they may also collect data of other people accidentally, such as a security camera capturing the images of passersby and Amazon Echo recording the voice of people who are in the background.

More critically, privacy-enhancing mechanisms that are designed to address people's privacy concerns may potentially create an unintended impact on other stakeholders, making privacy designs less desirable or usable. For example, the Platform for Privacy Preferences (P3P), a computer-readable language for privacy policies, provided a tangible way for Internet users to manage their privacy. However, it eventually became ill-fated due to several reasons, such as insufficient enforcement by web organizations [60, 186], ad-hoc derivation of privacy policies [199], and

discrepancies with their natural language counterparts [179]. In a sense, the lack of consideration about how P3P may have impacted or have been impacted by other stakeholders (web organizations, websites, etc.) contributes to its fate. Should these stakeholders be accounted for, P3P might have a different outcome.

In this dissertation, building on the notion of multiple stakeholders, I investigate in the following overarching research question, **how does the multi-stakeholder perspective change our understandings of privacy and inform privacy designs?** The goal is to understand in a socio-technical system, 1) how the consideration of multiple stakeholders may change people's privacy perceptions and expectations, and 2) how privacy designs may influence or are influenced by different stakeholders. To be more specific, in the scope of this dissertation, I do not focus on all stakeholders in a socio-technical system. Instead, I primarily focus on the perspective of technology-mediated interpersonal relationships (e.g., primary users and secondary users, owners and guests, etc).

To explore the above questions, I conducted a series of studies across three domains in which multiple stakeholders were involved, including Online Behavior Advertising (OBA), drones, and smart homes. I studied OBA because OBA focuses on individual users and relies on the data collected from these users. At the same time, OBA is significantly constrained by the needs and resources of multiple stakeholders (e.g., legal policies, tracking technology, etc.). I studied drones because, as an emerging technology, drones represent a relatively simple multi-stakeholder environment that primarily contains drone controllers and drone bystanders. Finally, I chose to study smart homes because smart homes represent a complex social environment with many variables, such as different social relationships, power dynamics, and potential confrontations.

I started the research by investigating average Internet users' understandings of how OBA works through an interview study. The results of the interview study informed the design and development

of a web tracking blocker. To explore the impact of the blockers on different stakeholders who are involved in the web tracking eco-system (e.g., legal stakeholder, economic stakeholder, technological stakeholders, and societal stakeholder), I followed up with a speculative design exploration. The speculation surfaced many potential issues beyond the usability and efficacy of a blocker, such as the tension between protecting users' privacy and tracking companies making profits. These issues could potentially inform the design of a web tracking blocker in the current time.

In the drone domain, building upon my prior research on understanding the privacy perceptions of drone bystanders (i.e., people who can be captured by the drone footage), I conducted an interview study to investigate the privacy perceptions of drone controllers. The results suggested the mismatched perceptions between controllers and bystanders. For example, when bystanders were concerned about their photos being taken by the drones, controllers thought the concerns were exaggerated since the drone cameras could not capture their images. Based on these perceptions, I proposed several privacy-enhancing mechanisms for drones. I ran a survey study to examine how controllers and bystanders perceive these mechanisms and whether they would like to use them. The results suggest that mechanisms that considered the needs of both bystanders and controllers received more recognition from both groups, such as the controller-bystander app and the automatic face blurring.

In the smart home domain, I started with discovering smart home users' privacy perceptions and their desired ways to protect their own privacy through a co-design study. I then conducted similar research with smart home bystanders (i.e., people who are not the users nor the owners of smart home devices but are subject to the data collection of these devices, such as guests, visitors, passersby, etc.), and tried to understand their privacy perceptions, needs, and desired ways to address their concerns. The results suggest that bystanders, whose privacy needs are often ignored, also have

privacy needs in smart homes. Through comparing the results with the two types of stakeholders in smart homes, I also discovered that cooperative mechanisms (i.e., mechanisms that bridge the needs of users and bystanders) are promising to fulfill the privacy needs of both stakeholders.

In the end, I synthesized the results from all three domains and summarized the design implications for privacy designs in general, considering the needs of multiple stakeholders. Besides, I also argued how the results of this dissertation complement the theory of Contextual Integrity and called out a research agenda to facilitate the discovery of contextual informational norms in the multi-stakeholder environment.

1.1 Document Organization

This dissertation starts by reviewing the literature on various conceptualizations of privacy, the theoretical foundation of this dissertation, and privacy research from the multi-stakeholder perspective. The literature review leads to the primary research question, **how does the multi-stakeholder perspective change our understandings of privacy and inform privacy designs?**

To answer this question, I present three case studies: online behavioral advertising, drone, and smart homes. In each case study, I first present the privacy issues and designs from the end user's perspective. Then, I present the exploration of other stakeholders' perspectives. In each case, I discuss the design implications drawing from this case regarding how considering the needs of multi-stakeholder informs privacy designs.

In the discussion, I synthesize the results from all three cases and analyze the commonalities, then draw design implications to reflect the lesson learned across all three domains. I then discuss how this dissertation contributes to the theory of Contextual Integrity. At last, I present a research agenda

to facilitate the application of the theory of Contextual Integrity in multi-stakeholder environments.

This dissertation contains text from several previous papers [231, 233, 232, 229, 230]. Additional text written for this dissertation will be used for future publication.

1.2 Overview

This dissertation investigates the privacy perceptions and expectations of two different stakeholders (e.g., technology users, bystanders, and external stakeholders) in three domains (i.e., online behavioral advertising, drones, and smart homes), as well as their desired controls to manage their privacy. The dissertation includes three projects, which were conducted from April 2017 through August 2019. The first project on online behavioral advertising includes an interview study with 21 participants, followed by a speculative design exploration. The second project on drones includes an interview study with 12 drone controllers, a second interview study with 16 participants, and a survey with 169 drone controllers and 717 drone bystanders. The third project on smart homes includes a co-design study with 25 smart home users and another co-design study with 18 smart home bystanders. The series of studies yield the following contributions.

1.3 Major Contributions

This dissertation makes four main contributions.

First, by examining three technological domains, this dissertation explores the privacy issues in each domain from the multi-stakeholder perspective with a primary focus on technology-mediated interpersonal relationships (e.g., drone controllers and bystanders). We find that conflicting interests

and social confrontations may change people's privacy perceptions, which further change how privacy mechanisms are designed.

Second, through the Contextual Integrity (CI) theory lens, this dissertation shows that in an environment that involves multiple stakeholders, there are some hidden factors that may influence the contextual informational norms (e.g., social relationships). As a result, this dissertation advocates that future research is needed to apply the CI theory to multi-stakeholder contexts.

Third, methodologically, this dissertation, unlike other methods in the literature that involve multiple stakeholders indirectly (e.g., Privacy Impact Assessment), uses a combination of qualitative, quantitative and design methods to directly engage different stakeholders in the research and allow researchers to dive deeply in the privacy perceptions of different stakeholders.

Fourth, this dissertation lays out design implications for practitioners and provides a research agenda for the application of the CI theory, particularly in multi-stakeholder contexts.

Chapter 2

Literature Review

In this section, I first review the literature related to various conceptualizations of privacy. Then, I draw the definition of privacy from the Theory of Contextual Integrity and the definition of stakeholders from the Stakeholder Theory to form the theoretical foundation of this dissertation. Finally, I review the literature related to the multi-stakeholder perspective of privacy research, which motivates the overarching research question for the remaining dissertation.

2.1 What is Privacy?

Solove defined privacy in the taxonomy, that privacy “is in disarray [and n]obody can articulate what it means” [193]. Although researchers from different science fields have attempted to define privacy numerous times, it is widely recognized that privacy is hard to define [192]. As such, in reviewing the literature on privacy conceptualizations, I will not provide a concrete definition of privacy. Instead, I will focus on several different conceptualizations of privacy from a social and legal perspective, and summarize how researchers from different realms define privacy. This

approach is largely inspired by Smith et al.'s work about information privacy research [192].

Privacy as a human right.

In Harvard Law Reviews (1890), Warren and Brandeis' defined general privacy as "the right to be left alone" [212]. They considered privacy as a fundamental human right. This definition has been found to influence many court cases that were related to general privacy [192]. In my opinion, this is a very initial attempt to define general privacy and is a very valuable benchmark. However, this definition does not consider the role of information privacy versus physical privacy, and I perceive this definition as more leaning towards physical privacy.

Privacy as a commodity.

Bennett (1995) and Cohen (2001) defined general privacy from an economic perspective. They both believed that privacy is on the contrary side of the information market economy [28, 53]. To them, rather than a fundamental human right, privacy is the economic principles of cost-benefits and trade-off [28, 53]. The idea of this definition is the fact that consumers shift their minds of privacy from a human right to a commodity. They can sell their information or cooperate the gathering of their data in exchange for some benefits, or perceived benefits [45, 65, 8]. This definition was leaning towards information privacy rather than physical privacy.

Privacy as a state.

Westin (1967) introduced the concept of "privacy as a state" [215]. He defined privacy from four psychological states: anonymity, solitude, reserve, and intimacy [215]. It is worth noting that Westin's conceptualization of privacy is largely in the "individual" level rather than the "collective" or "group" level, as Margulis later commented, "(Westin) models the organization on an individual who acts rather than on a collective" [145]. Weinstein (1971) defined privacy as a state of "being apart from others" [214]. In this definition, it is important to understand that unlike other states such

as alienation, loneliness, ostracism, and isolation which are generally initiated by the larger society, privacy is normally pursued by individual [214].

Privacy as a control.

Altman (1975) introduced the concept of privacy as a control. Altman defined privacy as “an interpersonal boundary-control process which paces and regulates the interaction with others”, and “the selective control of access to the self or one’s group” [13]. Based on this definition, Margulis [145] further defined, “Privacy, as a whole or in part, represents the control of transactions between a person(s) and other(s), the ultimate aim of which is to enhance autonomy and/or to minimize vulnerability”. Such a definition has been used in mainstream research in the privacy field and becomes the origin of many other definitions of privacy in other fields, such as information system and marketing [62]. For example, in the information system field, Culnan defined privacy as “the ability of individuals to control the terms under which their personal information is acquired and use.” [62]

Privacy as a situational concept.

Laufer and Wolfe (1977) conceptualized general privacy as a situational concept which tied to concrete situations [130]. These situations generally have three dimensions: self-ego, environmental, and interpersonal. Information systems, economics, and marketing scholars narrowed these definitions of general privacy so that they addressed information-based issues. The state of limited access was translated to the state of limited access to information. Furthermore, in the theory of Contextual Integrity, Nissenbaum introduced two information norms that serve as the benchmark of privacy: the appropriateness of the information to be revealed in a certain context’ and the flow or distribution of information from one party to another [165, 164]. She further noted that privacy is maintained if both of the two norms are held, and privacy is breached if either of the two norms is

violated.

Privacy taxonomy.

Privacy is not a singular concept but resembles a family taxonomy with different categories of privacy-related problems and the relations among these problems. These problems include information collection, processing, dissemination, and invasions. It is more philosophical than practical [193]. For instance, how to operationalize different categories of privacy-related problems in this taxonomy is an open question [193]. Also, it may fall into a recursive situation to conceptualize privacy with privacy-related problems.

Aside from the above conceptualizations of privacy, Nissenbaum defined privacy as a *contextual integrity*. Next, I will focus on the theory of Contextual Integrity, from which I draw the definition of privacy for the remaining of this dissertation.

2.2 Theory of Contextual Integrity

Since privacy research is very interdisciplinary, many theories from other fields (e.g., psychology, education, etc.) have contributed to the evolvement of privacy theories. In this dissertation, I choose the theory of Contextual Integrity (CI) as the theoretical foundation, primarily because CI is not conditioned to a specific time, location, or a specific situation [164]. Instead, it can be applied to different contexts “sweepingly defined as spheres of life such as education, politics, and the marketplace or as finely drawn as the conventional routine of visiting the dentist. . .” [164]. As a multi-stakeholder perspective can potentially introduce new angle of the privacy issues in various contexts, the CI theory provides a solid foundation to examine the emerging privacy issues. A more in-depth discussion of the CI theory will be presented in the discussion section. Here I briefly

review the theory formation and its application.

Nissenbaum proposed the theory of Contextual Integrity (CI) in 2004. The CI theory introduced the appropriate information to be collected under a certain context, and the appropriate information flow or distribution as the benchmark of privacy integrity [164]. As such, privacy is considered as breached if either one of the two norms is violated, and maintained if both of the norms are honored [164]. Another main idea in the CI theory is such norms are always associated with a specific context, thus are not universally applicable [164, 165]. In the further development of the CI theory, Nissenbaum adapted it to the online environment [165].

The development of this theory was originally motivated by the ambiguity of public surveillance practices, which consist of various ways of monitoring individuals, such as public records online, consumer profiling and data mining, and RFID usage [164]. It was previously assumed that people have little privacy expectations in public spaces and that aggregations of information are not intrusive if individual pieces of the information were not intrusive [54]. Nissenbaum argues that public records online violate the norm of information flow because certain locally kept records of an individual had been moved to the web and this placement had altered the information accessibility from local to global [165, 164]. RFID usage can also violate customers' privacy because RFID can make customers' information available to retailers, manufacturers, and other entities. As such, it transfers the discretion of using the information from the customers to the information gatherers. Therefore, it is also a violation of the information flow [164].

Several projects have leveraged the theory of contextual integrity in unpacking privacy violations in different contexts [127, 33, 239, 169, 205, 191, 23, 202]. Similarly, the theory of contextual integrity can guide our inquiry towards what information might be appropriate or inappropriate to collect by various technologies, and how the information should be distributed or shared among

different stakeholders. Discovering these norms can help us design privacy mechanisms to address potential privacy issues in various contexts.

2.3 Stakeholder Theory

Situated in the organizational context, Freeman proposed the Stakeholder Theory of organizational management and business ethics [88]. In the original work that detailed the Stakeholder Theory framework, Freeman first defined a “stakeholder” as “any group or individual who can affect or is affected by the achievement of an organization’s objectives” [88]. The theory then stresses the relationship between a business and its various stakeholders, such as customer, suppliers, investors, and others who may assist or hinder the achievement of the organization objectives [88], and addresses the business morals and values explicitly as a central feature in organization management [88]. For example, when making decisions in the organization context, sometimes managers have to ignore some existing relationships and obligations among different stakeholders in the organization, because they need to consider other stronger relationships and obligations [102]. The key here is that this entire process should be exposed and examined, and the conversation among different stakeholders should be facilitated [118]. In a sense, it calls for the needs to explicitly consider different stakeholders in the decision-making process of an organization. Besides, it also implies that the interests of these stakeholders are eventually joint to create values for each stakeholder [89].

As such, the Stakeholder Theory is fundamentally a theory about how to manage a business effectively and how the business works at its best [89]. Since the origin of the Stakeholder Theory in 1984, it has become a key theoretical foundation when studying business ethics and management [21, 113, 184, 51, 155, 90, 189, 181] and has challenged the traditional view that

business should treat the shareholders' wealth as the primary goal [89].

Inspired by the Stakeholder Theory, privacy researchers, when trying to understand the privacy issues or designing for protection, it is critical to consider not only the direct users but also other stakeholders who might impact or be influenced by the direct users. Research should also look into the interconnected relationship among all the stakeholders to understand privacy more thoroughly. Privacy designs, on the other hand, should also create values for all stakeholders rather than only for direct users. Drawing from the Stakeholder Theory [88], in the context of this dissertation, I define a "stakeholder" as "an individual or entity who can facilitate or hinder the use of technology".

Under this definition, many existing technologies have multiple stakeholders. For example, in the context of a smart home, homeowners are the primary stakeholders since they are the users of the smart home devices. In addition, there are many other stakeholders as well, such as the guests who are visiting the home, the passersby and the neighbors who may be recorded by the outdoor security cameras in the house, and so on. More broadly speaking, policymakers, device manufacturers, and advertising companies are all stakeholders of the smart home since they can impact or can be impacted by the use of smart home devices.

However, the existing privacy research has been overwhelmingly focusing on the privacy of the primary users and rarely touched on the privacy of and the potential implications on other stakeholders. In the next section, I will review the literature, though limited, on multi-stakeholder privacy research, from which I will draw the main research question for this dissertation.

2.4 Multi-Stakeholder Privacy Research

Privacy research has been primarily focusing on the privacy of the end-users, aiming to understand their privacy perceptions and concerns. Privacy designs and mechanisms are also built to address the needs of end-users. However, socio-technical systems may have influenced many stakeholders other than the end-users.

To this end, only a few prior research has explicitly discussed and addressed the multi-stakeholder aspect of privacy issues. For example, in wireless communication, the owner/primary users of cognitive radio require detailed usage information from the secondary users to calculate how much they should be paid [176]. However, the detailed information provided by the secondary users may invade the privacy of the secondary users. In this context, Zamkov et al. developed a privacy-preserving mechanism that allows the primary users to only know how much they should be paid without having all information from the secondary users [176]. Similarly, as a key component in the cognitive radio networks, cooperative spectrum sensing also poses great privacy risks to the secondary users since the sensing requires precious location information from the secondary users [100]. To address this issue, Hamdaoui et al. developed a privacy-preserving framework based on cryptography to protect the privacy of secondary users [100].

Besides, some other research has also loosely touched on the privacy of different stakeholders in various contexts. For example, from the view of bystanders, Denning et al. find that people expected to be asked for permission before they are recorded by Virtual Reality devices [66]. Lifelogging devices, such as SenseCam [105], have the risks to capture bystanders' faces and behaviors and put their privacy at risk [108]. In the context of a smart home, Lau et al. found that secondary

users' privacy might be breached if their voices were accidentally recorded by the voice assistant [129]. Stone and Stone-Remero discussed the privacy from the multi-stakeholder perspective in organization employment contexts and expressed the concerns on conflicts between companies' needs to enhance their authority and power through collecting employee's data and individual employee's needs for privacy [198]. On the legal front, the US National Telecommunication and Information Administration (NTIA) published the multi-stakeholder process for the Internet of Things Security [4], the unmanned aircraft systems [3], and cybersecurity vulnerabilities [2]. Such a process aims to provide opportunities to discuss and come to a consensus across all stakeholders, and negotiate potential solutions when there is difficulty to achieve consensus [5].

However, one core issue in a multi-stakeholder environment is the potential conflict of interests among different stakeholders. Such conflict of interests could potentially change how different stakeholders perceived privacy, how privacy-enhancing mechanisms are designed, and whether the privacy designs will be adopted. In this dissertation, I aim to unpack how the multi-stakeholder perspective informs privacy designs by examining three cases, each of which involves at least two different stakeholders. The goal is to understand the privacy perceptions of different stakeholders, their desired way of privacy protection, how they are impacted by other stakeholders in the ecosystem, and then draw design implications to guide future privacy designs.

2.5 Research Question

Drawing from the above literature as well as the theory of Contextual Integrity and the Stakeholder Theory, I aim to investigate the privacy perceptions and designs from a multi-stakeholder perspective. As such, in this dissertation, the overarching research question is:

How does the multi-stakeholder perspective change the understandings of privacy and inform privacy designs?

To explore this research question, I turn to three different domains, i.e., online behavioral advertising, drones, and smart homes, to understand how considering multiple stakeholders in these specific domains changes people's privacy perceptions and designs. In each case, I start with the immediate stakeholder, i.e., the end-users, then expand the scope to other stakeholders.

Part II

Case 1: Online Behavior Advertising (OBA)

Chapter 3

Online Privacy: Users' Understandings of OBA

Online Behavioral Advertising (OBA) is pervasive on the Internet. While there is a line of empirical research that studies Internet users' attitudes and privacy preferences of OBA, little is known about their actual understandings of how OBA works. This is an important question to answer because people often draw on their understanding to make decisions. Through a qualitative study conducted in an iterative manner, we identify four "folk models" held by our participants about how OBA works and show how these models are either incomplete or inaccurate in representing common OBA practices. We discuss how privacy tools can be designed to consider these folk models. In addition, most of our participants felt that the information being tracked is more important than who the web trackers are. This suggests the potential for an information-based blocking scheme rather than a tracker-based blocking scheme used by most existing ad-blocking tools.

3.1 Introduction

Online Behavioral Advertising (OBA), or targeted advertising, is prevalent on today's Internet [196]. OBA is "the practice of tracking an individual's online activities in order to deliver advertising tailored to the individual's interests" [55]. A common practice of OBA is that first-party sites (i.e., sites that a user visits voluntarily) rely on third-party entities (e.g., ad networks) to track a user's browsing activities across websites and to serve ads targeted at the user [146]. OBA can benefit both advertising companies (e.g., increasing click-through rates and prices of ads [27]) and Internet users (e.g., providing ads that better match their potential interests [147, 205]). However, since OBA involves online tracking and profiling of users, it has raised significant privacy issues [204, 147, 205].

Prior studies have found various user attitudes and perceptions of OBA (e.g., [204, 147, 205, 133, 177]). For instance, Ur et al. note that people find OBA "creepy and scary" because of its online tracking practices, but sometimes people also find OBA "smart and useful" [205]. As such, individual users seem to have varying acceptance of OBA depending on the context [48, 211, 150]. However, most of these studies either (1) did not study people's understandings of how OBA works (e.g., [204, 132]) or (2) investigated people's perceptions of OBA after the researchers explained OBA (e.g., [133, 177, 48, 211, 150]), therefore it is not clear to what extent ordinary Internet users actually understand how OBA works nor what their understandings are.

Drawing from the literature on mental models, we examine people's understandings of how OBA works. Psychologist Kenneth Craik pioneered the concept of mental models, describing "the mind constructs 'small-scale models' of reality that it uses to anticipate events, to reason, and to underlie explanation" [57]. Since then, the notion of mental models has been further developed. For instance, Phil Johnson-Laird, an influential scholar of mental models defines them as "psychological

representations of real, hypothetical, or imaginary situations” [116]. Mental models have also been studied extensively to understand how people comprehend various things such as language and music [96]. In addition, “mental models affect people’s reasoning” [116] and people draw from their mental models to make various decisions [116, 117]. For instance, people’s mental models of how thermostats work influence the ways in which they control these devices [217].

The mental model approach has also been applied in the domain of privacy and security (e.g., [9, 20, 44, 213, 35, 138, 158, 75]), but has not been systematically used in the context of OBA. Rick Wash conducted an interview study to examine people’s mental models of home computer security [213]. He notes, “to understand the rationale for people’s behavior, it’s important to understand the decision model that people use” [213]. Drawing from prior literature (e.g., [185, 10]), he uses the term *folk models* to denote mental models that can be incorrect representations of reality but are used by people in practice [213].

Our work was in part inspired by Wash’s study [213]. We aim to uncover people’s folk models of OBA, regardless of whether these models accurately represent the reality of OBA. We note that mental models can encompass more than a picture of how things work [96], but here we use folk models to denote people’s understanding of how OBA works. There is little work that touches on this question, and our study aims to fill the gap. We believe that understanding people’s folk models of OBA is important because these models can influence people’s behavior or decisions regarding OBA, for instance, how they control or limit OBA. Furthermore, privacy tools for OBA can be more effective when they incorporate people’s folk models, for example, by helping people recognize privacy risks (e.g., third-party tracking) and adopt countermeasures (e.g., blocking third-party trackers).

We inductively developed four folk models of OBA held by our participants through a qualitative

study conducted in an iterative manner. In addition to a pilot study with eight people, we conducted two rounds of semi-structured interviews with another 21 Internet users from different U.S. states and cities. These models differ in terms of the following: who tracks Internet users' information; where the tracked information is stored; and how targeted ads are selected or served.

Similar to Wash's study [213], our qualitative research does not support claims that can be generalized to all Internet users, but it instead aims to uncover folk models that people have about OBA and that can inform future privacy-enhancing designs for OBA. In the sense of theoretical sampling [142], the discovered folk models are held by real people but the study says little about how common or statistically representative these models are in the general population.

To guide future privacy tools for OBA, we also asked participants' opinions about what tools or features they desire in order to help them protect their privacy in the context of OBA. While most OBA tools focus on trackers, most of our interviewees felt that the information being tracked is more important than trackers (i.e., who is tracking the information). This result suggests the potential for an information-based blocking scheme rather than a tracker-based blocking scheme used by most existing ad-blocking tools such as Ghostery.

This chapter makes two main contributions. First, we uncover different folk models of OBA that ordinary Internet users have. These models have implications for privacy designs and public policies of OBA. Second, we identify people's desired features of privacy-enhancing tools for OBA. These features should be explored by future privacy tools.

3.2 Related Work

Our work was mainly inspired by prior research on people’s attitudes and perceptions of online tracking and OBA, as well as people’s mental models of privacy and security.

3.2.1 People’s Attitudes and Perceptions of OBA

There is a line of empirical research that examines people’s attitudes towards OBA mostly via surveys. Several surveys have shown people’s objection of online tracking and OBA. For instance, Turow et al. polled 1000 Internet users in the U.S. and found that 87% of them did not want advertisers to track them online [204]. Similarly, McDonald and Cranor found that 64% of their survey respondents considered targeted ads to be “invasive” [147]. Another survey found that one major reason why the respondents disliked OBA was because of online tracking and subsequent analyses of the tracked data [175].

However, Ur et al.’s interview study has painted a more nuanced picture. They found that many of their interviewees considered OBA “creepy and scary” because of its online tracking practices, but sometimes people also found targeted ads “smart and useful” [205]. This study also suggested people’s acceptance of OBA may vary depending on the context.

A number of subsequent survey studies focused on people’s context-based preferences of OBA [48, 211, 150]. Leon et al. found that the data retention period and scope of data use significantly affected their respondents’ willingness to share data for OBA [133]. Chanchary and Chiasso found that people’s OBA preferences differ by the first-party sites they visit [48]. Melicher et al. combined their participants’ browsing histories and interview data in identifying additional situational factors such as the types of information being tracked and the frequency of visiting

first-party sites that can affect people's attitudes towards online tracking [150]. Wang et al. surveyed both American and Chinese Internet users and found that both user groups had different OBA preferences based on the type of first-party sites, despite the fact that the former had more privacy concerns over OBA than the latter [211].

While these prior studies offer invaluable insights into people's perceptions of OBA, most of these studies (e.g., [133, 177, 48, 211, 150]) provided a detailed explanation of OBA before examining people's preferences of it. In contrast, four prior studies asked people's perceptions of OBA before explaining OBA [204, 147, 205, 132] and two of them did not ask about people's understandings of how OBA works [204, 132]. The other two studies touched on this question but did not yield mental models that represent people's understandings of OBA [147, 205]. Ur et al.'s study focused on people's attitudes towards OBA rather than their understandings of OBA [205]. The remaining study investigated people's beliefs about OBA [147] but differed significantly from our study.

More specifically, McDonald and Cranor provided their survey respondents four diagrams depicting different configurations of first- and third-party cookies in OBA and asked the respondents to select the configuration which was not possible [147]. Unlike their approach, we sought to discover people's folk models of OBA without providing any a priori models or pictures to constrain or influence their thinking. We have discovered folk models (e.g., browser-based models) that differ from the models they provided in their study. We will present our folk models in the results section.

In addition, few studies have touched on people's understanding of online tracking and OBA. Rader conducted an online experiment and found that most participants were aware that sites like Google or Facebook can collect information about their users' activities on them (e.g., what pages they visit or what links they click) [177]. This is a case of first-party tracking. Ur et al. asked their

interviewees the ways in which ads are tailored to them. The two most common methods mentioned were based on users' browsing histories and web searches [205]. Another survey study found that people have various understandings of the type of data (e.g., personal information or location) web trackers can track online [48]. Some of these perceptions were incorrect, e.g., people thought online tracking was malware and online tracking directly involved local browsing history [150]. Our work differs from these studies in that we focus on people's folk models of OBA rather than exploring them in passing.

Overall, the extant literature does not provide a clear picture of the folk models people have about how OBA works. Our study aims to fill this gap.

3.2.2 People's Mental Models of Privacy and Security

The mental model approach has been employed by a number of researchers to investigate people's understandings of the Internet [200, 119]. Thatcher and Grey's work utilized drawing as a means of understanding people's mental models [200]. Their work revealed several typical understandings of how the Internet works, such as considering the Internet as a central database, or as a modular structure network [200]. Our study adopted a similar drawing task to solicit people's understandings of OBA.

Kang et al. observed that people's mental models of how the Internet works can be very different, and these models were partially influenced by people's technical knowledge [119]. The researchers suggested that users with more technical knowledge tend to have a more sophisticated mental model, but the level of technical knowledge barely affects users' security and privacy practices [119].

Researchers have also used the mental model approach to investigate users' perceptions related

to their privacy and security. Camp proposed five possible mental models that can be used to explain people's understandings of computer risks, including models of physical safety, medical infections, criminal behavior, warfare activities, and market failures [44]. Asgharpour et al. conducted a card sorting study and found that computer security experts and non-experts have different mental models of computer security [20]. For instance, experts associated passwords with a criminal model whereas non-experts thought of a physical safety model [20]. Wash's work on people's mental models of threats towards their home computers suggested eight folk models, including four virus-centered models and four hacker-centered models [213]. Bravo-Lillo et al. used a mental model approach to understand computer users' psychological processes and reactions toward computer warnings [35]. They were able to identify different perceptions of novice and advanced users and to obtain insights in improving computer warnings [35]. Most recently, Naiakshina et al. studied people's mental models of the security of mobile messaging tools and found that people overestimated the capabilities of attackers [158].

The above studies shed light on people's mental models of the Internet as well as privacy and security risks. However, people's mental models of OBA still remain unclear. Our study aims to address this gap by inductively analyzing people's understandings of how OBA works.

Our primary research question is what folk models people employ in practice about OBA, for instance, regarding the information flow in OBA. This was in part inspired by Helen Nissenbaum's theory of contextual integrity which presents a framework to determine privacy violations based on the norms and appropriateness of information flow in a particular context [164]. A secondary research question is what privacy-enhancing features or tools people desire for OBA. Answers to both questions will inform future privacy designs for OBA.

3.3 Method

We designed and conducted a qualitative study in an iterative manner to understand people's folk models of OBA. This study was approved by the IRB. We started with a pilot study to test the interview script and explore people's understandings of OBA. We then conducted a first-round of interviews to develop initial folk models, followed by a second-round of interviews to further verify the models.

3.3.1 Pilot Study

Drawing from prior research examining people's attitudes and perceptions of OBA [204, 147, 205, 133, 177, 48, 211, 150], we developed a list of interview questions that investigate people's understandings, attitudes, and experiences of OBA. To assess the quality of these questions, we pilot tested this interview protocol with eight family members and friends during January and February, 2016. The pilot results suggest that they understood the questions albeit most of them did not understand how OBA works. For instance, most of them did not know that third-party entities (e.g., ad networks) are likely involved in OBA. These pilot study participants' understandings of OBA were covered by the four folk models developed in the subsequent two rounds of interviews. For instance, many of them held the connected-first-party model.

The pilot results also suggest that they varied in their opinions of OBA after we explained OBA and that they differed in their interests in learning more about OBA and/or using tools to control OBA. In order to further identify people's understanding of OBA (i.e., mental models) and their preferences of OBA, we added a drawing task and a card sorting task.

3.3.2 First-Round Interviews

We revised the interview protocol based on the feedback from the pilot study. Next, we describe the updated protocol.

Questions about Internet usage

We began our interviews with questions about interviewees' demographics such as age, gender, and occupation. We then asked about their background in using computers and the Internet, e.g.,

“What do you usually do when you browse the web? What devices do you use to browse the web?”

We also asked about their usage of web browsers, e.g., “Do you know that you can change your browser settings? Do you know what a browser extension/add-on is? Do you save any of your account information in your browser? What kind of information do you save?”

We then asked them to sort 18 cards, each containing an information item (e.g., name or home address), based on their comfortableness with saving the data into their browser. This card sorting task was designed to assess their perceived sensitivity of different information. Most interviewees put the information items into two or three clusters based on their perceived sensitivity, for instance, social security numbers as highly sensitive and religion as mildly or moderately sensitive. Since these card sorting results mostly corroborate the findings reported in the prior literature (e.g., [133, 122, 211]), we removed this task from the second-round of interviews.

Mental models of OBA

Next, we asked about interviewees' attitudes toward and interactions with online ads, e.g., “Do you notice that there are ads on websites? Do you generally click on ads?”

Similar to the use of hypothetical scenarios in Wash's mental model study [213], we presented a

hypothetical ad scenario in which a user first looks for shoes on Amazon.com and a few hours later he or she visits Facebook and sees other shoe ads there. This scenario was designed to represent common OBA practices that interviewees can easily understand since Amazon and Facebook are popular sites that people visit. We then asked them to draw what they think happened in this scenario on a piece of paper and to explain their drawing. This drawing with think-aloud task explored interviewees' own understandings of how OBA works before we offered our definition and explanation of OBA. These drawings visualized the interviewees' folk models of OBA (i.e., their own theories of how OBA works).

We followed up with additional questions about their knowledge and understanding of OBA and web trackers, e.g., "Have you heard of targeted ads? Do you know how targeted ads work? Have you ever heard of web trackers? What do you think web trackers are, who they are and what they do?"

After the interviewees answered the above questions about OBA, we offered the same explanation of web trackers to each interviewee. Specifically, we explained that the sites they visit voluntarily are first-party entities, and that web trackers are typically third-party entities which track user information and can provide ads targeted to the user based on the collected user data (e.g., browsing activities, page visits). We then answered any questions that interviewees had about web trackers. We also asked them "What do you think trackers are collecting when they are tracking you? What's more important to you, the trackers or the data is being tracked?"

Privacy-enhancing tools for OBA

Finally, to help inform future privacy design for OBA, we asked interviewees questions about their desired features in helping them deal with web trackers, e.g., "If there was a magic tool that can do

anything, what types of features would you like this tool to have pertaining to web trackers?”

We asked these questions after explaining OBA with the rationale that if interviewees did not have a correct understanding of OBA, they may miss features that they would have wanted. For instance, similar to what we found in the pilot study, many participants in this round of interviews were not aware that web trackers are often third-party entities. These participants requested the privacy tools to provide more information about OBA, including the third-party trackers involved. If we asked these tool-related questions before explaining OBA, these participants would not know the existence of third-party trackers and thus are unlikely to ask for corresponding tool support. However, asking these tool-related questions before explaining OBA might discover that people having different folk models desire different privacy features. Therefore, we asked these tool questions both before and after explaining OBA in the second-round interviews.

3.3.3 Second-Round Interviews

We analyzed the first-round interviews and developed four folk models that our participants had about how OBA works. Similar to the iterative methodology used in Wash’s mental model study [213], we conducted a second-round of interviews with new participants to check the validity of these models by seeking “negative” examples [168] that are not covered by these models.

There were two major updates of the interview protocol in this round. First, we removed the card sorting task. Second, we asked the questions related to privacy tools both before and after explaining OBA. The updated study thus included the following sequence of components: questions about Internet usage, questions about mental models (with the same hypothetical scenario), questions about privacy tools, our explanation of OBA, and the questions about privacy tools (second time).

3.3.4 Participant Recruitment

We recruited prospective participants from Syracuse University campus, shopping malls, public libraries, and online communities (e.g., Craigslist). We also used snowball sampling, i.e., asking participants to refer our study to their contacts [30]. We deliberately selected participants in order to create a diversified sample in which participants have various demographic characteristics and occupational backgrounds.

From March to May 2016, we recruited and conducted our 1st-round of interviews with 14 participants from the Syracuse area. These interviews were face-to-face. From July to August 2016, we recruited and conducted our second-round of interviews with seven additional participants from Pennsylvania, Seattle, and Los Angeles. These interviews were conducted online using services such as Skype. Participants showed and explained their drawings in the interviews and sent their drawings to the researchers afterwards. Each interview took about one to two hours and was compensated \$10.

It is worth noting that our sample is not statistically representative of the general Internet user population, but it is diverse in terms of participants' age, geographic locations and occupations. Similar to Wash's study [213], we do not believe our sample is particularly special. There are probably other people similar to our participants in the general population. In addition, we did not observe any significantly new findings, particularly regarding people's understandings of how OBA works, from our second-round of interviews. This suggests theoretical saturation [99] and thus we did not conduct any additional interviews.

Table 3.1: Participants used three factors in reasoning about OBA and constructing their folk models.

Folk model	Who info	tracks	Where stored	info	How selected or served	ads or
Browser-pull	Browser		Browser		Browser	pulls ads
1st-party-pull	Browser		Browser		1st-party	sites pulls ads
Connected 1st-party	1st-party		1st-party		1st-party	sites share data di- rectly and pull ads
3rd-party	1st-party		3rd-party		1st-party	shares data with 3rd-party, 3rd-party pulls ads

3.3.5 Data Analysis

We audio recorded all interviews upon participants’ permission, and transcribed the recordings. We then conducted a *thematic analysis* [32], a common approach for analyzing qualitative data.

First, we read through all the interview transcriptions multiple times to immerse ourselves in the data. Second, two co-authors coded one interview together at the sentence level to develop a code book.

Then, the two coders coded the same subset of interviews independently using the code book. When they encountered concepts not covered by the existing code book, they added new codes accordingly. Once finished, the two coders compared, discussed and converged the codes into an updated code book of 210 unique codes, such as, “Internet experience,” “attitudes toward OBA,” and “PETs features.” We wrote the codes on post-it notes and created an affinity diagram to group these codes into nine themes: background, misconception, advertisement, specific information

concerns, privacy-enhancing technologies, mental models, privacy and security practices, privacy expectations, and web trackers.

Finally, we read the associated interview quotes to ensure the coherence within each theme. Based on our review, we adjusted the inappropriately grouped codes and the affinity diagram accordingly.

3.4 Results

In this section, we report the results from the 21 interviews, focusing on our participants' folk models of how OBA works and their preferences of privacy tools for OBA.

3.4.1 Participants

The ages of the 21 participants ranged from 19 to 67, with an average of 34. Six participants were female and 15 were male. They were from a wide range of locations, including large and small cities in the states of New York, Pennsylvania, California and Washington. Our participants also had various occupations such as university staff, college students, software engineers, business professionals, retired workers, a mechanical engineer and a waitress.

All of our interviewees use computers and the Internet on a daily basis. Two of them use the Internet less than 2 hours a day, the rest of them use the Internet more than 7 hours a day. The primary purposes of using the Internet include checking emails, using social media, doing research for their jobs, contacting friends and families, and reading news. In addition, 19 of our interviewees had heard about targeted ads. Some of them voluntarily talked about their experiences of targeted ads. Four interviewees said that they have heard of web trackers, but only one understood what a web tracker is.

3.4.2 Folk Models of OBA

We provided our interviewees a detailed scenario to understand their thoughts about how OBA works and how information flows. The interview results suggested that our participants’ understandings of how OBA works mainly differed by three factors: who tracks users’ information; where the information is stored; and how ads are selected or served. Based on these three factors, we identified four major folk models. Table 3.1 summarizes the factors that our participants used to reason about OBA and construct their folk models. Table 9.1 summarizes participants’ folk models as well as their attitudes toward web trackers and OBA.

Table 3.2: Participants’ folk models and attitudes of trackers and OBA.

ID	Folk model	Accept trackers	Accept OBA
P1	3rd-party	Yes	Yes
P2	Connected 1st-party	No	Yes
P3	3rd-party	Yes	Yes
P4	3rd-party	No	No
P5	Browser-pull	No	Yes
P6	Connected 1st-party	No	No
P7	Connected 1st-party	No	No
P8	3rd-party	No	No
P9	Browser-pull	No	No
P10	1st-party-pull	Yes	Yes
P11	Browser-pull	Yes	Yes
P12	1st-party-pull	Yes	Yes
P13	Browser-pull	Yes	Yes
P14	Connected 1st-party	No	No
P15	Browser-pull	No	No
P16	1st-party-pull	Yes	Yes
P17	3rd-party	Yes	Yes
P18	Connected 1st-party	No	Yes
P19	3rd-party	Yes	Yes
P20	3rd-party	Yes	Yes
P21	1st-party-pull	Yes	Yes

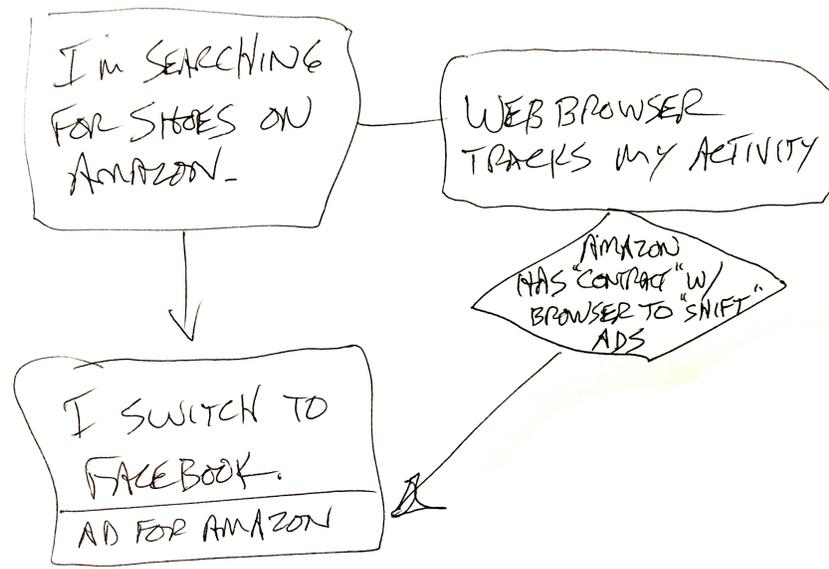


Figure 3.1: Browser-pull model: an example from P9.

When a user searches for a pair of shoes on Amazon, the web browser will save the search information. The web browser has contracted with Amazon. When the user visits Facebook, the browser will pull the saved information and display ads for Amazon on the user’s Facebook page.

3.4.3 Browser-Pull Model

Five interviewees held this model. They believed that all tracking is done by the browser, which would pull from advertisers relevant ads that target user data/profiles stored locally by the browser. In this model, the web browser plays the primary role in OBA. For instance, P5 thought that the web browser monitors and detects his browsing patterns and pulls ads based on those patterns. He also believed that all tracked information is saved in his local computer.

“The system is set up to notice your patterns and to pull information that seems relevant to you...I’m just thinking [the information] is [transmitted to] my computer.” (P5)

P9 had a similar view as illustrated in his drawing (see Figure 3.1) in which the browser tracks his online activities and has contracted with Amazon to ship their ads.

“I’m searching on Amazon and looking for shoes, web browser tracks my activity, and, you know, I’m just thinking that Amazon and ads are contracted with web browser, and browser just ships ads. There’s when I’m on Facebook, the ads just pops up.” (P9)

P15 also held this model but felt that he can control the browser’s tracking through the browser settings.

“I think it is all based on your Internet options what you allow. I think it is the browser that allows this...No matter whatever browser I’m on...I can go to the Internet options and mess around the way it looks into my information.” (P15)

The essence of this model is that the browser is the key – the browser tracks users’ activities, saves their information on the local computer, and selects and displays the relevant ads. Because the browser is on users’ computers, some participants holding this model (e.g., P15) also had the perceived agency to limit or control OBA through the browser settings.

3.4.4 First-Party-Pull Model

Four interviewees held the first-party-pull model. Similar to the browser-pull model, participants of this model also believed that all tracking is done by the browser. However, unlike the browser-pull model, people of this model thought that first-party sites (e.g., Amazon or Facebook) rather than the browser pull relevant ads based on the user’s data/profile stored in the browser. In this model, both the web browser and first-party websites play active roles in OBA.

For instance, P10 explained the use of cookies and the retrieval of targeted ads by the first-party sites, as shown in his drawing (see Figure 3.2). In this example, when he searches for kayaks on Amazon, the browser will save the search information in a browser cookie. Then, the browser will find other sites that also sell kayaks. When he visits Facebook later, Facebook will pull these

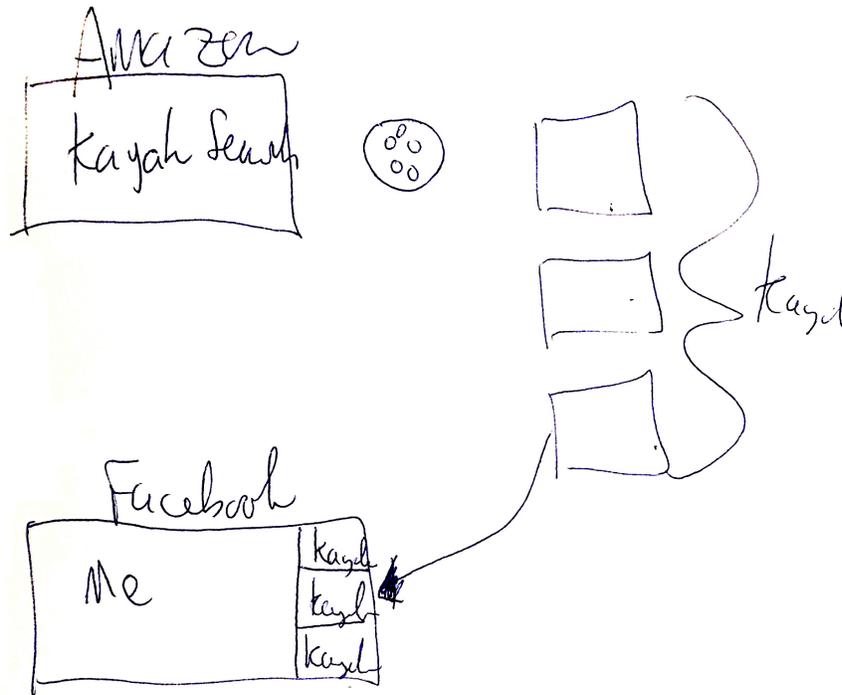


Figure 3.2: First-party-pull model: an example from P10.

When a user searches for a kayak on Amazon, the browser will save the search information in a browser cookie. The browser will find other sites that also sell kayaks. Later, the user visits Facebook, which will pull these sites from the browser and display them on the user's Facebook page.

sites that sell kayaks from the browser and display them on his Facebook page. Note again, the first-party site (Facebook) pulls the relevant ads. In addition, P10 believed that first-party websites can only access the cookies from the last website that the user visited. P12 shared a similar model but described his theory in a more technically sophisticated way, highlighting the use of the HTML meta tag on first-party sites (e.g., eBay or Facebook). He believed that these websites are designed in a way (with similar meta tag structures) so that they can directly access all of the user's browsing/searching history and cookies in order to select targeted ads.

“So this is the eBay webpage, and in your meta-tag you're gonna have embedded information

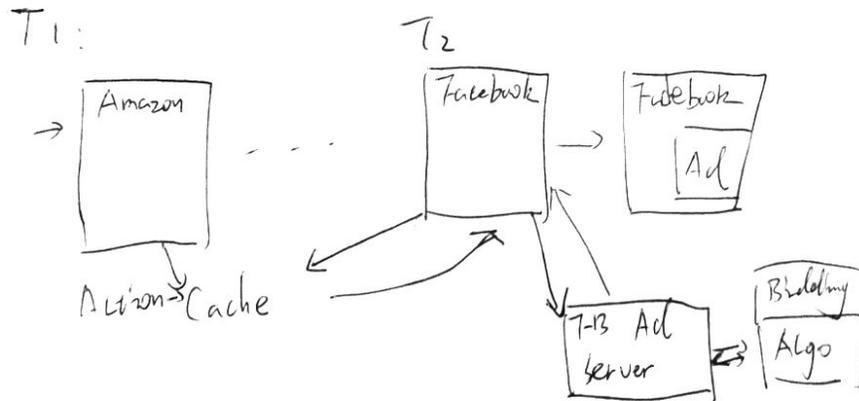


Figure 3.3: First-party-pull model: an example from P16.

He searches shoes on Amazon, then the browser stores the action in the cache; later, he visits Facebook, which then goes to its advertising server, and bids ads with the cached user information using a bidding algorithm. Facebook then displays the ads on his Facebook page.

that not only pulls up the information from your cookie and consent your account to automatically login...but it also contains advertising tracking data...And then if you log into, for instance, Facebook, if they have a similar meta-tag structure they can access the search data from this tracking cookie, so that this controls the same search criteria.” (P12)

P16 is a web developer with technical knowledge of the Internet. His drawing (see Figure 3.3) illustrated that the browser stores the user’s Amazon activities in its local cache; then Facebook pulls that user’s information from the cache, bids ads with that user’s information, and finally displays the targeted ads on the user’s Facebook page. He was our only participant who mentioned ad bidding, which suggests that he had more knowledge about the online ad ecosystem than other participants. However, he was not aware of third-party tracking in OBA. After he described his mental model, we explained OBA to him and asked about his feelings of OBA again. He was surprised to learn that often third-party entities (e.g., ad networks) track users’ online activities for OBA purposes but his concerns about his information and privacy remained the same.

“I never thought it is third party. It matters to me since I’m not sure how they can use my

information legally and the purposes. But I think I feel the same, because before I know this, I'm still worried about my information and privacy.” (P16)

P16's example is rather telling because we would normally assume that technically savvy users would know how OBA works. But, that was not the case. This is somewhat surprising because even web developers like him who seemed to have knowledge about ad bidding (a rather advanced understanding of online ad systems) did not know about common online tracking done by third-party entities.

Other less technically savvy participants also held this model albeit with less details. For instance, P21 thought that the browser records his activities on Amazon and then Facebook pulls his information from the browser. However, unlike P10 and P16, he did not know the technical specifics of how Facebook can actually access his information stored by the browser.

“Chrome gets all my transaction from Amazon, and for some reason, Facebook can access this information.” (P21)

While these participants provided different levels of technical details in their explanations, the underlying theory is the same. In this model, the browser still plays an important role in tracking users' browsing activities and storing this information locally. But, the first-party sites (e.g., Facebook) rather than the browser select relevant ads based on the user profile stored in the browser. This means, unlike the browser-pull model, the browser cannot single-handedly deliver the targeted ads. Instead, first-party sites select ads that they think are relevant to users.

3.4.5 Connected-First-Party Model

Five interviewees believed that first-party websites (e.g., Amazon and Facebook) are directly connected and collaborate with each other. In this model, users' data is tracked and stored by each

first-party site that they visit. Different first-party websites are connected, exchanging the user data that each of them tracks and saves. As such, first-party websites form a kind of a collaborative network and play the main role for delivering targeted ads.

For instance, P18 believed that first-party sites have shared resources between them so they can share their user information stored in their databases.

“It definitely goes into the database on Amazon, and then it will probably, I guess there is some kind of shared resource between them, so it will basically go into the database on Facebook, then shows on my page.” (P18)

P7 shared the idea and also explicitly mentioned a partnership between Amazon and Facebook, which makes the information sharing possible.

“Amazon and Facebook have some type of partnership and so Amazon gives them certain information and are able to locate certain people for specific products and advertise certain

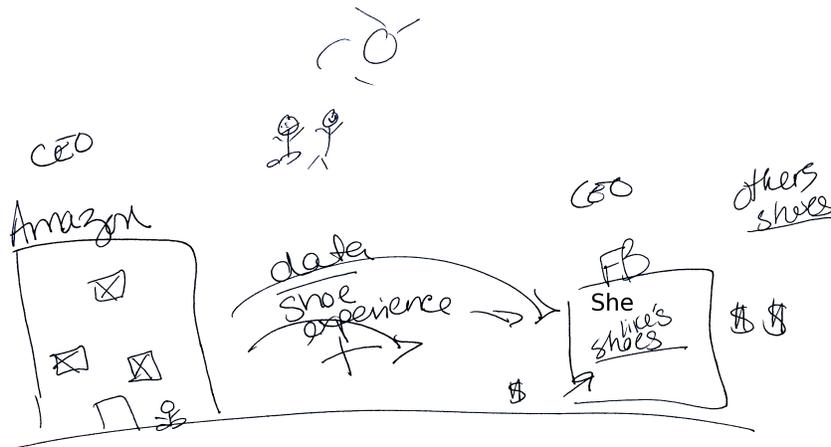


Figure 3.4: Connected-first-party model: an example from P6.

When she searches shoes on Amazon, Amazon saves the search information. Amazon sells the data to Facebook (indicated by a dollar sign at the bottom left of the Facebook box). Facebook then gets more money from “other shoe” companies (two dollar signs) to show the shoe ads on her Facebook. The CEOs of Amazon and Facebook under the sun would not share her data if it is not for money.

products to certain people for whatever they're looking for.”

P6 went further and suggested that Amazon sells her data to Facebook, as shown in her drawing (see Figure 3.4).

“Amazon has decided to work in conjunction with Facebook, this is my personal belief and transferred all of my data about my shoe experience and gave it to Facebook, to say 'Hey, she likes shoes.' And now Facebook is going to get bulk dollars from other shoe companies because now I know about other places beside or other shoe types beside the one I just bought...I think amazon sells that to Facebook.” (P6)

P6 emphasized the economic or business model in her understanding. Her drawing shows that Facebook gives money to Amazon (with one dollar sign) for her data that Amazon shares and Facebook receives more money from other shoe companies (with two dollar signs) to serve their ads on her Facebook page. As such, P6 believed that money drives the connection between Amazon and Facebook. She disapproved an alternative explanation in a witted fashion and articulated money as the driving force behind this connection.

“I don't see why Amazon would do this because I don't see like the CEO of Amazon and the CEO of Facebook hanging out under the sun as best friends smiling...so there's got to be a reason...the biggest lubricant I ever come across is money, or at least some kind of gain of some sorts.” (P6)

The key of this model is that first-party sites are directly connected and they share user data with each other in order to select targeted ads. According to this model, the connected first-party websites enable OBA, regardless of the reasons for their connections (e.g., a partnership or user data purchases).

3.4.6 Third-Party Model

Seven interviewees held this model. In this model, people believed that first-party sites track and collect user data then contribute the data to a third-party entity, and then the third-party entity leverages the user data it has (presumably from different first-party sites) to select relevant ads for users. As such, various first-party websites and third-party entities are involved in OBA, according to this model.

Some participants believed that there are third-party entities involved but they knew almost nothing else about these third-party entities, for instance, who they are or whom they belong to. For example, P4 drew a big bubble that she called an “Internet space” that stores and provides user data to different sites such as Facebook (see Figure 3.5). But, she cannot tell what this Internet space is or who controls it.

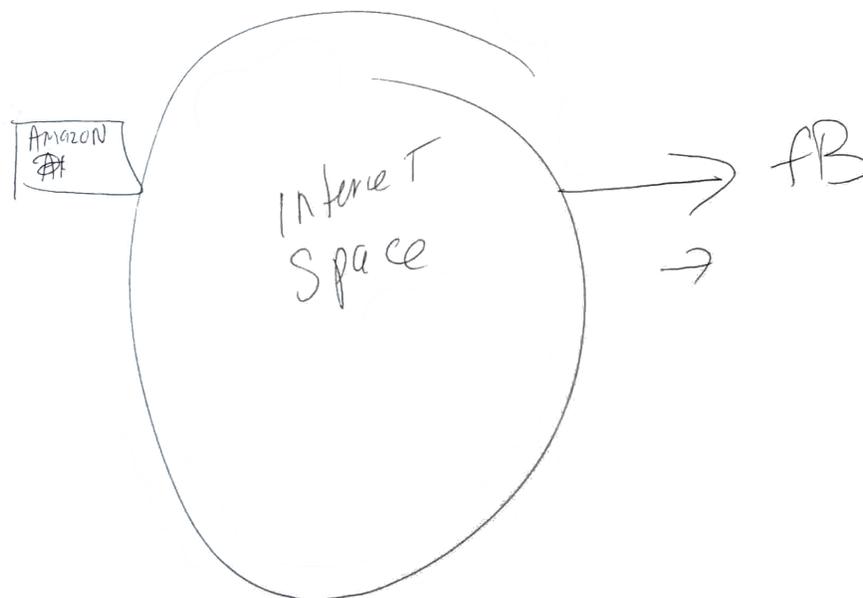


Figure 3.5: Third-party model: an example from P4.

When a user searches shoes on Amazon, Amazon collects the user’s data and then transmits the data to an Internet space. This Internet space sends ads to Facebook, which will display the shoe ads on the user’s Facebook.

“I don’t know, it must be like some Internet thing, Internet space I don’t know, and somehow it just goes to like Facebook and whatever else there is out there.” (P4)

P19 drew a more detailed graph, illustrating the existence of some database that all companies such as Amazon and Google share (see Figure 3.6). But, he knew nothing else about this database.

“I don’t really know. I guess there should be some sort of database in the middle, then not only Amazon and Facebook, but also other companies, have access to it, keep injecting new information to it. It’s more of a shared space, or common space for all companies who are involved in this ecosystem.”

P19 also questioned how Amazon and Facebook match the same user. He hypothesized the use of cookies, which include a user’s IP address. He also doubted first-party sites (Amazon and Facebook in our scenario) are directly connected. This is an important difference from the connected-first-party model in which first-party sites are directly connected.

P17 also believed there is a central database and it is likely dominated by Google.

“There must be some central database or data center...I think it’s like Google. Google has something like this, like many big companies have this kind of data center. But it is dominated by Google. In the example you mentioned, there is no Google involved, so I guess it is third party.”
(P17)

In this case, he suggested that large companies like Google represent the third-party entities. This understanding was fairly accurate since Google indeed represents a major web tracker and serves targeted ads across the Internet [196].

Like P6 who had a connected-first-party model, P1 also focused her understanding on the economic aspects. However, P1 believed that those third-party entities rather than the connected

first-party sites make the ads ecosystem work.

“These guys [third parties] have an agreement with Amazon, they are like, ‘Oh, I’m just going to take information from this guy’. Facebook gets money by displaying the ads sent by these guys [third parties]...this branch [third parties] allows that to happen. So in a way it is a neutral third party.” (P1)

In her view, the third-party entities connect Amazon and Facebook, collect and store users’ data,

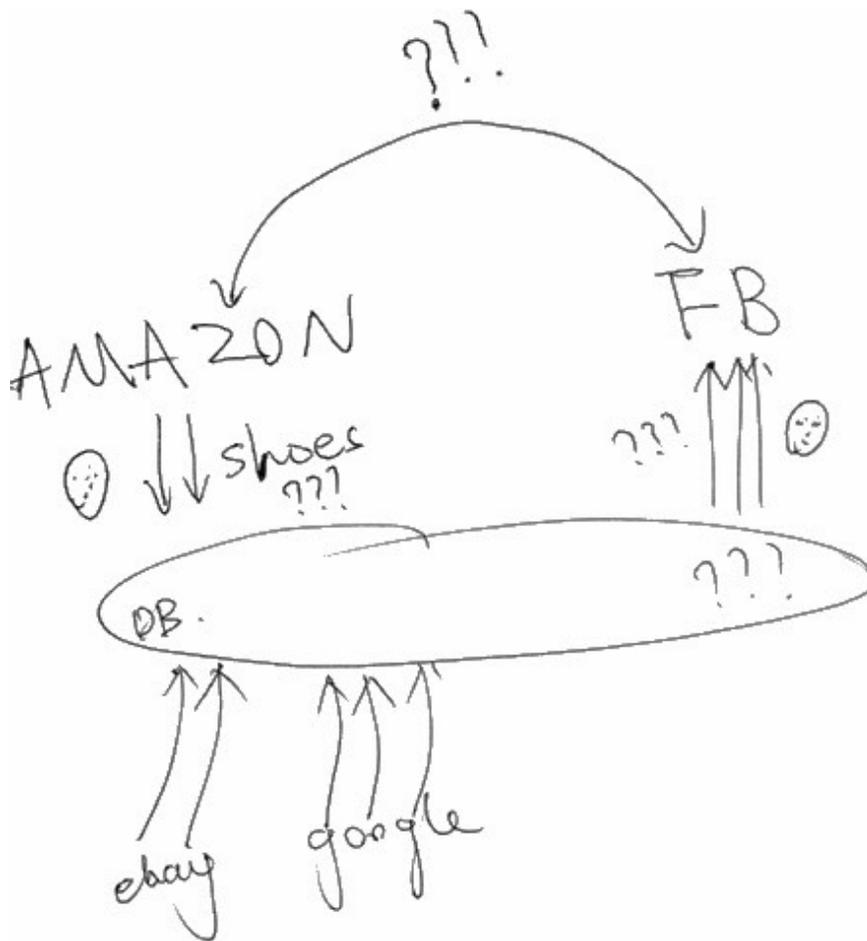


Figure 3.6: Third-party model: an example from P19.

All companies share a common database. When a user searches shoes on Amazon, Amazon sends that information to the shared database. Other companies such as eBay and Google also contribute user information to this database. When the user visits Facebook, the site obtains user data from this database to select relevant ads.

and then send targeted ads to Facebook.

Regardless of whether these participants knew who the third-party entities are or represent, they shared the key understanding that these third-party entities rather than the first-party sites that users visit voluntarily make the OBA work. Both user tracking and selection of targeted ads are done by these third-party entities. According to this model, first-party sites are not connected directly but are bridged through third-party entities.

3.4.7 Misconceptions and Speculations of OBA

During our interviews, we also observed participants' recurring misconceptions and speculations about OBA. We use the word "misconceptions" to denote our participants' inaccurate understandings of web trackers and what trackers collect, mainly from a technical standpoint. Typical misconceptions our participants had include: trackers are hackers, and trackers are viruses. Wash's home computer security study has uncovered several hacker-based and virus-based mental models [213], however, his participants did not report considering web trackers as hackers or viruses. Furthermore, we use the word "speculations" to represent our participants' views that are technically possible but their applications for OBA are not clear. Common speculations our participants made include: trackers access local files on a user's computers, and trackers resides locally on users' computers.

Misconception: trackers are hackers

Some interviewees identified trackers as hackers, people with malicious intentions. For instance, before getting our explanation of OBA, P2 believed that web trackers can hack into his online accounts.

"They say it's a secure site and you got to login, but of course I login with the same password that I always use...I'm sure that those web trackers can hack in there too." (P2)

P2 seemed to confuse web trackers with hackers that aim to break into people's accounts and steal their personal data.

Misconception: trackers are viruses

Considering trackers as computer viruses was another common misconception among our participants. For instance, P4 expressed her belief that her anti-virus software will protect her from trackers.

“Thank God for Norton because sometimes it comes up oh so and so just attacked you or something, so I don't even pay attention because I figure that will save me.” (P4)

These participants seemed to misconstrue web trackers as computer viruses designed to attack their computers.

Speculation: trackers access local files

Some participants thought that trackers can access files stored on their local computers. For instance, when asked whether he would be interested in a tool that can block trackers, P5 expressed his lack of interest because there is very little on his computer that he worries about. He mentioned, *“Even things that are around my desktop, besides my resume and cover letters and that's about it.”* His explanation reflected his overestimation of the capabilities of trackers in which they can access (arbitrary) files stored locally on his computer. P5 also believed that trackers can log his typing, saying *“everything you type in can technically be downloaded.”* While tracking users' typing is technically possible, we are not aware of any reports of this kind of tracker behavior in practice.

Speculation: trackers reside locally on user computers

Some participants indicated that trackers can not only be something in the browser but also reside locally on their computers. For instance, P5 said *“I think it’s in the web browser. I also think there’s something on your computer.”* But he could not elaborate what he meant by “something” on his computer.

3.4.8 Privacy-Enhancing Tools for OBA

To inform future design of privacy tools for OBA, we asked our interviewees questions about tools or features that can help protect their privacy in the context of OBA.

Trackers vs. the information being tracked

Existing ad blockers such as Ghostery are structured by trackers. When a user visits a website, the ad blocker shows a list of trackers on the site that the user can selectively block. However, these tools do not show what type of information each tracker tracks. In addition, prior research has shown that ordinary Internet users do not recognize the names of most trackers (e.g., BlueKay) with few exceptions being household names such as Google [132]. Furthermore, our card sorting results support the prior literature (e.g., [133, 122, 211]) that people perceive different levels of sensitivity for different information items (e.g., home address is perceived more sensitive than educational level). Given these observations, we wondered whether the information being tracked is more recognizable and thus more useful to users than the trackers. Therefore, we asked our participants “what is more important to you, the tracker or the information being tracked?”

All but one interviewee answered that the information being tracked is more important. For instance, P1 cared more about the information being tracked because this information can be used to make assumptions about her.

“I would say what is being tracked. I guess they use the information to build out their profile, I guess it is a little strange using the information they collect to make assumptions about me, what type of person or Internet user.” (P1)

P7 provided a different justification, arguing that the tracked information can be used to identify individuals.

“I mean the biggest thing is the information. I mean trackers are replaceable, but information is not because that’s a specific set of info per person.” (P7)

P8 was the only participant that did not perceive the information being tracked to be more important than the trackers because he wanted to know both.

“What information is being collected for sure, but I also want know who is collecting it. I want to say both, because, you know, I would want to know who that person, or the entity is, how they are gonna use that information.” (P8)

Desired privacy features for OBA

When asked about their expectations of a magic tool that can help protect their privacy regarding OBA, our participants suggested many features.

Block tracking. A commonly desired feature is to block tracking. For instance, P17 would like to automatically block trackers based on his preferences. P16 desired a feature that allows him to select the type(s) of information that he wants the trackers to track or not to track.

Interestingly, when we asked participants’ experiences with online ads, some participants reported using ad blockers to block ads but they did not relate these ad blockers to web trackers. This might be because they are called ad blockers rather than tracker blockers.

Transparency. Several interviewees were also interested in knowing more about trackers and

their behaviors.

“It is a scary technology, but maybe if I have a better understanding of how it connects with companies or something like that. Maybe I can see like the scope of web trackers? Like how many people it affects, how many places my information is going.” (P1)

P1 hoped to know detailed information about the scope and effect of tracking. In addition, P19 was interested in knowing what data is being tracked by whom and for what purposes.

In our second-round interviews, we asked this privacy tool-related question both before and after we explained OBA. P18 held a connected-first-party model and requested additional privacy tool support after our OBA explanation, which made him realize the existence of third-party trackers. He then suggested the tool to provide detailed information about third-party trackers and their behaviors.

Effortless to use. In addition to concrete features, many interviewees emphasized the tool should be effortless to use. P10, for instance, expressed that he would only use such a tool if it only needs a one-time setup for all websites.

“This is per website or do I do it one time and it does it for every website? That was my first thing cause I don’t want to have to do it per website.” (P10)

This is understandable because privacy protection is often not people’s primary or direct task. Therefore, they would not want to divert from their main task to spend too much time in using a privacy tool. For example, automatic blocking of tracking as suggested by P17 would satisfy this need.

3.5 Discussion

Drawing from the literature on mental models and particularly Rick Wash's work on folk models of home computer security [213], we examined Internet users' understandings of how OBA works through a qualitative study including a pilot study and two rounds of semi-structured interviews.

We discovered four folk models of how OBA works. The *browser-pull* model assumes that all tracking is done by the browser, which would pull from advertisers relevant ads that tailor to the user data/profile the browser stores locally. In this case, the browser is the "middleman" between the first-party site and advertisers. The *first-party-pull* model presumes that all tracking is still done by the browser, but first-party sites pull relevant ads based on the user data/profile stored in the browser (e.g., cookies). In this case, first-party sites decide which ads to show. The *connected-first-party* model posits that different first-party sites directly share and even sell user data that they collect with each other and that one first-party site can use another first-party site's user data to pull relevant ads directly from advertisers. In this model, first-party sites directly interact with each other and with the advertisers. Lastly, the *third-party* model assumes that first-party sites first track and collect user data then contribute the data to a third-party entity, then this third-party entity uses the user data it has (presumably from different first-party sites) to select relevant ads. This model is closer to common OBA practices than other models but it is still not detailed enough, e.g., some participants hardly knew anything about the third-party entities.

As discussed in the related work section, our work is one of the first studies that investigate people's mental models of OBA. The body of literature on mental models of privacy and security rarely touches on the topic of online tracking or OBA, for instance, Wash's study focuses on home computer security [213]. The extant research on people's privacy perceptions of OBA also does not

focus on people's understanding and mental models of OBA. The notable exception is the work of McDonald and Cranor in which they provided their survey respondents four diagrams of OBA, focusing on who have access to users' cookies [147]. They then asked their respondents which diagram is unlikely to happen [147]. In comparison, our folk models emerged from our interviews rather than pre-defined by us. Our folk models differ from their cookie-centered models [147] because ours are based on three factors: who tracks user information, where the tracked information is stored, and how the targeted ads are selected or served.

3.5.1 Why Folk Models of OBA Matter

The folk models uncovered by our study are novel, but why do they matter? There are several reasons why they matter.

User education

All four folk models are either inaccurate or incomplete. Similar to Camp's suggestions that risk communication should be designed based on non-expert mental models [44], we believe that it would be useful to customize user education of OBA based on the folk models.

In our second round of interviews, some of our interviewees changed their attitudes towards OBA because of our explanation of OBA. For instance, some participants of the connected-first-party model were surprised to learn that their information can be tracked or even sold by third-party entities. Therefore, their attitudes towards OBA were changed from neutral to negative. Knowing a user's current folk model can tailor the education to reduce the knowledge asymmetry between the user and the OBA practitioners.

Previous studies have suggested that technically savvy users have more accurate or sophisticated mental models than their less technically savvy counterparts (e.g., [119]). However, we did not

observe a clear relationship between technical knowledge and folk models. Somewhat surprisingly, our arguably most technically savvy participants P16 and P18, two web developers, held the 1st-party-pull model and the connected-first-party model, respectively. Both of them were not aware of third-party trackers. This is important because even technically savvy users can have inaccurate models and need user education to gain a more accurate picture of OBA.

Attitudes towards OBA

Capturing people's folk models can help understand people's attitudes towards OBA. We observed some associations between the two.

Interviewees of the browser-pull model had different attitudes toward tracking and OBA. These participants believed that the browser tracks and stores their data. Interviewees who were aware of different browser settings (e.g., clear browser history and cookies) tended to be positive about OBA because of their perceived ability to control tracking by setting the browser options. In contrast, those who did not know about browser settings tended to be more critical of OBA.

Interviewees of the first-party-pull model generally accept OBA because they only expected first-party sites to access their information in order to select relevant ads. They had little concern because they generally trusted the sites that they visit voluntarily. However, they were unaware of the existence and impact of third-party tracking.

For interviewees holding the connected-first-party model, they were generally not against online tracking because they thought their data is only shared between first-party sites that they trust. However, they did not appreciate the idea of first-party sites selling their information between each other. They understood that this is one of the main business models of the Internet, but they still disliked it.

Participants of the third-party model all included third-party entities in their explanations. However, their descriptions of third-party entities varied significantly, ranging from a clear idea of a specific organization to a vague notion of an “*Internet space*.” Their attitudes toward OBA also varied. We did not observe any significant patterns in this group.

User behavior

The literature of mental models suggest that these models can influence people’s reasoning and decision making (e.g., [116, 117]). We also encountered some examples of certain folk models affecting people’s behavior in our study. For example, some participants of the browser-pull model rely on browser settings to control online tracking because they believed that web tracking and OBA are carried out by the browser. Another example is that P18 of the connected-first-party model requested a transparency feature that provides detailed information about third-party trackers only after we explained OBA. This suggests that people of different folk models may need different privacy features (particularly educational features) that tailor to their (lack of) understanding of OBA.

3.5.2 Implications for Design and Policy

Our results have a number of implications for privacy designs and public policies of OBA.

First, as mentioned before, future privacy tools for OBA could highlight different information to cater to people with different folk models. For example, for people having the connected-first-party model, the tools can emphasize that third-party entities can be tracking and sharing their online activities. In addition, governmental policies or industry best practices could require or encourage privacy policies of web tracker companies to include simple but visual representations of how they work in the OBA ecosystem, similar to the way that our interviewees drew their folk models.

Wash argued that technologies should be designed to work with people’s mental models even if these models are incorrect because it is more difficult to educate users about the correct mental model [213]. We agree with this viewpoint to some extent. For instance, while the browser-pull model does not capture the common OBA practice, researchers have proposed privacy-preserving, client-based OBA systems, resembling the browser-pull model [31]. However, we still believe there are benefits to educate people about OBA practices that are common on the Internet. For instance, people holding the browser-pull model might think they can control or stop OBA by just setting their browser options. Therefore, that folk model could discourage them from adopting more effective privacy tools such as ad blockers that can block trackers.

Second, popular tools such as Ghostery and Adblock are capable of blocking third-party trackers. These tools list the trackers on a site and allow people to block them selectively. However, most of our interviewees felt that the information being tracked is more important than the trackers themselves. This is a significant finding because it suggests that a completely different blocking scheme, one based on the type of information being tracked, might be perceived more useful by Internet users than the status quo, a tracker-based blocking scheme. In other words, the tools can be structured by the information being tracked rather than by a list of trackers. In addition, these tools can allow users to block tracking of certain types of information. Alternatively, future tools can support both schemes.

Emerging technologies, such as OpenWPM [81], Sunlight [131] and ReCon [180] are promising in identifying or inferring what information is being tracked by a tracker and the purpose of tracking to some extent. They pave the way for information-based blocking tools. On the policy front, we advocate that web trackers and ad networks should clearly explain what information they collect and why they collect the information in their privacy policies and preferably in a machine-readable

format. This could enable future privacy tools to automatically analyze and compare the behaviors of different trackers and the OBA practices of different sites.

3.5.3 Limitations and Future Research

We outline our study limitations and directions for future research. First, we did not have a particularly large sample. But our study was conducted in an iterative manner including a pilot study with eight people and two rounds of interviews with a total of 21 participants. The results from the pilot study and the actual interview study were consistent. In fact, we did not learn any significantly new things from our second-round interviews, suggesting theoretical saturation. Our sample is also diversified in that our participants came from various age groups and geographical areas, representing different occupations. Therefore, we are confident our results are valid.

Second, our qualitative study aims to examine people's folk models of OBA in depth rather than assess how statistically representative these models are in the generic population. In future work, we plan to conduct a large-scale survey to further examine how common these models are.

Third, when we asked our interviewees to draw their mental models of OBA and web tracking, we only used one hypothetical scenario. This may prevent us from discovering additional models. Future work can include multiple scenarios and ideally ones that people have experienced themselves.

Fourth, we asked participants to do the card sorting task before the drawing task in our first-round interviews. The card sorting task asked about participants' comfortableness with saving their data into their browser. This might prime people to think more about browsers. However, we believe the priming is minimum because we removed the card sorting task in our second-round interviews and there were participants having the browser-pull model and the first-party-pull model. In both

models, the browser is responsible for tracking users.

Finally, our interviews are self-reported data and thus do not include participants' actual behavioral data. To further examine the impact of these folk models on people's behavior, future work can consider collecting and analyzing user behavior data, for instance, through experiments and/or log analyses.

3.5.4 Conclusion

Online Behavior Advertising is pervasive on the Internet. We interviewed 21 people from the US to investigate their understandings of how OBA works. We identified four folk models held by our interviewees. These models are either inaccurate or incomplete in representing common OBA practices. User education tailoring to people's folk models of OBA is likely to be more effective. In addition, most of our interviewees felt that the information being tracked is more important than the trackers. Future privacy tools should consider these folk models and user preferences of OBA.

3.5.5 Acknowledgement

We thank our interviewees for sharing their insights. We also thank Alexander Krapf, Satoko Mii, Eduardo Nunez-Amador, Huichuan Xia for their assistance as well as Jason Dedrick and anonymous reviewers for their thoughtful comments on earlier versions of this chapter. This work was supported in part by NSF Grant CNS-1464347.

Chapter 4

Online Privacy: Implications Beyond Users' Privacy

Tools to limit web tracking often operate in either a tracker-based model or an information-based model. Prior research has been focused on the effectiveness and usability of these tools on the individual users' level. How these tools impact the larger space beyond individual users over longer time periods remains understudied. In this chapter, we use speculative design scenarios to surface questions about how the adoption of web tracking blockers impacts not only users, but also a broader set of societal, economic, legal, and technical stakeholders. Our speculative explorations, building on a Wizard-of-Oz study, surface many potential issues beyond the usability and efficacy of blockers, e.g., the tension between protecting users' privacy and tracking companies' need to make profits. Our exploration also broadens the understanding of privacy beyond individual users and argues that future tools to limit online tracking should consider the needs from a broader set of stakeholders.

4.1 Introduction

Online tracking is a pervasive practice on the Internet. One common use of online tracking is online behavioral advertising (OBA), which most Americans felt was an invasion of their privacy [204, 147]. People's attitudes towards OBA are also contextual based on the type of information being shared

with OBA [133] and online activities people undertake [211].

To help users limit online tracking, researchers and practitioners have proposed and implemented many blockers. These blockers can be divided roughly into two types, i.e., the *tracker-based* blocker and the *information-based* blocker. A tracker in the tracker-based model consists of a piece of software (usually javascript) that tracks a user's information. Most existing tools and blockers use the tracker-based model, in that they list the trackers that are present on a website and allow users to block the list of trackers collectively or selectively (e.g., Ghostery [98], Privacy Badger [25], and Disconnect.me [67]). However, the usability of this model has been challenged by a few prior studies (e.g., users only recognize trackers from household names, such as Google [132]). The information in the information-based model refers to the types of data that are collected by trackers (e.g., location or IP address). Research has shown that most users are more concerned about what information is collected rather than who is tracking that information [231]. As such, researchers have proposed and implemented tools based on this model in mobile systems (e.g., users can choose to block their location data or contacts from being tracked by installed applications) and on the web (e.g., P3P [59]).

This research expands prior work that focuses on the usability issues of these tools. We aim to investigate multiple facets of blocking tools in two ways. First, we explore the multi-dimensional web tracking blocker design space by examining existing practices and literature, and a Wizard-of-Oz study. Second, grounded in the design space and inspired by speculative design research approaches [74], we explore the implications for adoption [140] for blockers through a set of speculative scenarios in order to understand the technical, economic, political, and legal contexts that could affect or be affected by the widespread use of blockers. We deliberately choose this approach over a traditional user study because while the traditional user study allows to understand

user reactions to the blockers, the speculative scenarios allow us to take a step back and situate end user privacy as one value among a more complex landscape of additional stakeholders and other social values. This approach aims to surface broader issues and impact that are beyond what would normally emerge from traditional user studies of privacy enhancing technologies.

This chapter makes two contributions. First, through a speculative exploration complemented by a Wizard-of-Oz study, we explore the potential technical, economic, political, and legal spaces that situate or impact the wide adoption of blockers and understand the implications beyond individual users. Second, we discuss how our approach using a Wizard-of-Oz study and a speculative exploration broadens our understanding of usable privacy research and argue for adopting similar methodology in studying privacy in other domains.

4.2 Technical Background

In this section, we provide some background knowledge related to online tracking and details of the two blocking models.

Broadly speaking, online tracking can be categorized into first-party tracking and third-party tracking. **First-party tracking** is conducted by the websites that users visit and directly interact with [11]. For instance, if Amazon tracks what a user does on its own website, this is considered first-party tracking. **Third-party tracking** refers to the practice where third parties (e.g., advertising companies) embed their tracking technologies (e.g., cookies) across the first-party websites (e.g., CNN) to track individuals' activities on these websites [182]. This is third-party tracking because users are tracked by someone different than who they are directly interacting with.

In this chapter, we adopt the definition of online tracking from Melicher et al.'s work, "online trackers partner with websites to track visitors' activities" [150], and focus on third-party tracking.

We will use tracking/trackers to denote third-party tracking/trackers unless specified otherwise.

The tracker-based blocking model is the mechanism that has been adopted by the majority of the available products on the market, e.g., Ghostery [98]. In a tracker-based blocker, a list of trackers on a particular website is displayed in real-time as the users browse through the website. Users can block online tracking by blocking trackers collectively or selectively. As mentioned before, existing blockers have adopted this model and some of these tools have suffered from several usability and effectiveness issues, such as the unfamiliar names of the trackers [132].

The information-based blocking model is a blocking mechanism that enumerates what types of information are potentially tracked on the website, and allow the users to block the types of information collectively or selectively. Prior research has proposed several mechanisms based on the information-based model, such as P3P [59]. In the real world, the information-based blocking model has been partially implemented in the smartphone eco-system (i.e., users are asked to share their location and Bluetooth usage with some apps) and on the web (i.e., when users visit some websites, they are asked whether to allow or block the tracking of their locations with the websites [121]). It is worth noting that tools that operate in the information-based blocking model primarily target the first-party tracking.

4.3 Related Work

4.3.1 Socio-technical Countermeasures for Web Tracking

Numerous countermeasures have been proposed to combat online tracking. These countermeasures include Do-Not-Track (DNT), an HTTP head field that once enabled signifies that the user wants to be opt-out of being tracked [146]; the Platform for Privacy Preferences (P3P), a machine-readable privacy disclosure that could be retrieved automatically by web browsers to make users aware of

website privacy practices [58]; the proxy server, an intermediary between users and the destination server to preserve users' information [186]; alternative services that do not track users or provide extra security measures, such as the search engine DuckDuckGo [71], the Tor web browser [39], and the Brave web browser [38]. It is worth noting that both DNT and P3P are not just technical projects but also web standards and organizational projects carried out by the World Wide Web Consortium (W3C). They later became ill-fated, particularly P3P, due to several reasons, such as insufficient enforcement by web organizations [60, 186], ad-hoc derivation of privacy policies [199], and discrepancies with their natural language counterparts [179].

Another type of countermeasures includes several web tracking blocking tools. For example, Ghostery [98] and Disconnect.me [67] can detect and block various types of trackers such as online behavioral advertising (OBA), social media, and site analyzers, and are classifying and blacklisting trackers by examining their service domains [203]. Privacy Badger [25] leverages a slightly different mechanism that requires less "custom configuration to block non-consensual trackers." It judges and prevents a third-party by detecting its behavior across different sites, hence it could be regarded as "behavior-based" instead of "blacklist based" [203]. There are also more severe blockers which inhibit all third-parties, such as Request Policy [172], or ad-oriented blockers such as Adblock Plus [171]. Recently, some web browsers embedded features to limit web tracking as well, such as the Intelligent Tracking Protection in Safari to prevent cross-site tracking [226] and the Tracking Protection in Firefox to block web trackers [6].

These blockers have been well studied in prior research regarding their effectiveness and usability. For example, Traverso et al. compared seven of the most popular blockers and found none of them could deliver their promised protection [203], even though Ghostery outperformed the other blockers [203, 153]. Roesner et al. also found that third-party cookie blocking is ineffective

and plugins would remove potentially useful features for users (e.g., social media buttons) [182]. Wills and Uzunoglu also found that the effectiveness of the current blockers varied. For example, Disconnect.me could only provide moderate protection and Ghostery, Adblock Plus, and Adguard have to be manually configured to enable better protection because their default protection would be minimal [218]. Leon et al. tested nine online ads blocking tools and identified a number of serious usability flaws: the interfaces were hard to understand, the opt-out options were hard to set up, and the settings were complex to configure [132].

One noticeable trend in the above research is that they all consider individuals' privacy as the primary goal when designing and implementing the blockers. However, very limited research has expanded the scope to understand the impact of these blockers beyond individual users. This is an important topic mostly because the online tracking system represents a complex socio-technical environment. In this system, web tracking blockers impact and are also influenced by many other stakeholders. For example, partly in response to blocking, some websites and third parties are trying to track people through new ways that are less easily detectable [77]. In some other cases, websites may limit the user's experience or access to content if they detect blockers being used. These cases demonstrate the impact of adopting the blockers and potential push back from other stakeholders other than the end-users.

Our goal in this chapter is to explore the implications of adopting various online tracking blockers on and beyond individual users. Such implications can potentially illuminate the designs of future blocking technologies.

4.3.2 Speculative Design for Privacy Research

Proposed by Dunne and Raby, speculative design uses design to create artifacts that exist in the future or an alternate present, and brings those artifacts to the present, demanding responses from the present [74]. Speculative design uses design artifacts and the process of design to imagine alternate socio-technical configurations of the world as a way to interrogate questions about values and politics through design [219]. Rather than focus on solving an immediate problem, it uses design to ask questions and open up or explore a design space [74, 73, 22]. By creating an imagined world surrounding a speculative artifact, it also allows the investigation of how technologies are entangled with social, legal, and economic forces.

As privacy researchers, we have seen prior research in the design community beginning to explore issues related to privacy using speculative design practices. For example, Lindley et al. used design fiction to articulate how an Internet of Things door lock could communicate with voice user interfaces (e.g., Amazon Echo) and surfaced potential data collection and sharing issues under the umbrella of the EU General Data Protection Regulation (GDPR) [141]. Using design packets, Pierce et al.'s work imagined creative alternatives to the existing forms of privacy policies and discovered a number of new insights for future privacy policies from a design lens, such as the need for active user engagement [170]. Inspired by the science fiction novel *The Circle*, Wong et al. crafted a set of design fictions to explore privacy and surveillance issues in emerging sensing and tracking technologies [222, 221].

However, while technical privacy research has adopted user-centered design techniques, speculative design approaches are rare in the privacy research community [220]. Speculative design can be useful to help probe issues related to privacy that exist beyond interactions between a human

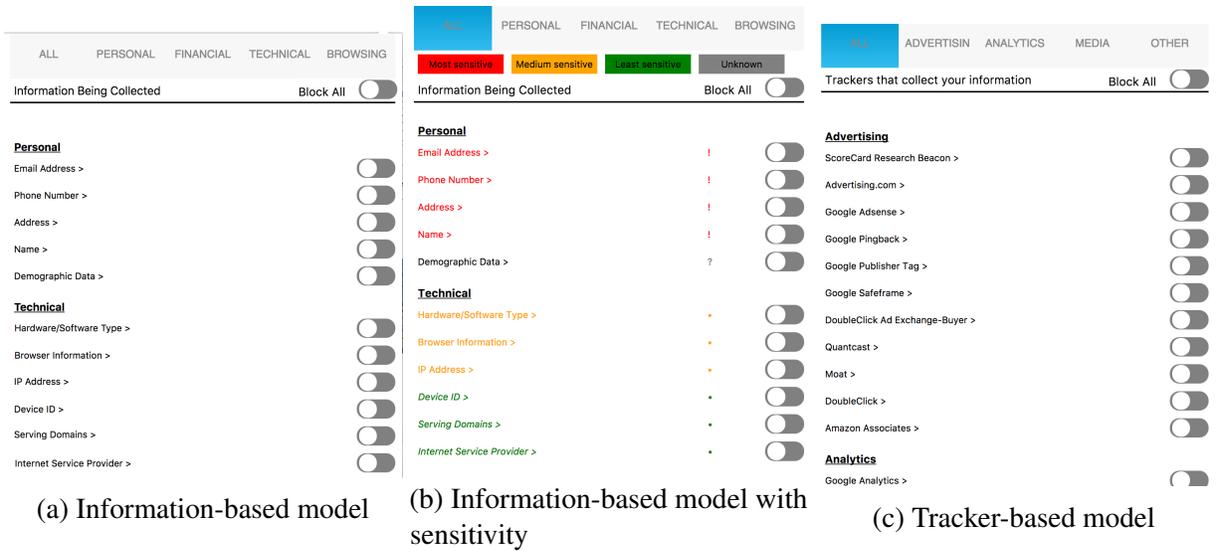


Figure 4.1: Our three interactive prototypes.

(a) An information-based model; (b) an information-based model with information sensitivity; (c) a tracker-based model, which mimics the Ghostery interfaces [98]. Participants can toggle the button to allow or block the collection of specific information / tracker.

and system at the interface level. At a broader level, speculative design can also help us explore the problem space of “privacy” itself, allowing us to see privacy as situated in a specific context or situation. When looking at privacy in online tracking contexts, speculative design also begins to view privacy as a social construct, interrogating privacy for who, from whom is privacy is protecting, and who is responsible for providing privacy? How might the implications of the answers these questions differ for different stakeholders in the online environment? These are in line with recent conceptualizations of privacy as situated in different social contexts and different subject positions [157, 164]. Moreover, speculative design allows us to investigate where privacy might come into conflict or tension with other values held by end-users or other stakeholders. This follows calls by Lindley et al. for HCI researchers to consider “implications for adoption” to explore and discuss technologies beyond their prototype stage using design practices such as speculative design or design fiction [140].

4.4 Iterative Processes of Speculation

In this section, we describe the process of how we came up with our scenario-based speculations.

The decision to use scenarios was made for several reasons. Usable privacy research has a tradition of using user-centered design techniques, so creating scenarios that resemble traditional user scenarios may help this work to be more accessible to other privacy researchers. Our work is also inspired by Nathan et. al's value scenarios, which similarly repurposed traditional scenario-based design with a critical and speculative design perspective, in order to envision long term systemic effects of new technologies [160, 159]. Our detailed steps are presented below.

4.4.1 Step 1 - Explore The Tracking Blocker Design Space

To explore the tracking blocker design space, we first looked at existing blockers and literature to understand common design elements in their designs. We reviewed popular blockers (e.g., Ghostery [98], Disconnect.me [67], AdBlock Plus [171]), blocking features embedded in various browsers (e.g., Safari [226], Brave [38], Firefox [6]), and prior research on the effectiveness and usability of these tools (e.g., [203, 182, 218, 132]). We then summarized five dimensions that have been used in designing blockers. These dimensions include Information Dimension (i.e., what information is presented to the users), Mode Dimension (i.e., information-based or tracker-based blocking), Automation Dimension (i.e., whether the blocker can run automatically without human intervention), Preferences Dimension (i.e., whether users are allowed to set their blocking preferences), and Categories Dimension (i.e., whether the information is categorized for better readability).

These dimensions were identified based on our observations and reviews of the existing blockers and research, particularly in terms of how these blocker designs could differ significantly (e.g., levels

of human involvement rather than graphical details of the UIs). For instance, Ghostery provides a list of trackers to the users [98] whereas Adblock Plus runs in the background without providing users much information [171]. This observation led us to create the Information Dimension. As another example, Ghostery runs in a tracker-based model [98] whereas Google Chrome allows users to block certain information from being tracked, e.g., location [121]. Thus, we created the Mode Dimension to capture this difference in design. The Automation Dimension was drawn from the observation that Adblock Plus automatically identify and block trackers [171], whereas Disconnect.me relies on a tracker list to identify trackers [67].

To ensure what these dimensions covered align with end-users' expectations of an ideal blocker, we designed three interactive prototypes in the form of browser plugins (as shown in Figure 4.1) using different combinations of the above five dimensions. We then used them as probes in a Wizard of Oz study with 15 participants to explore whether we have missed any other dimensions. We recruited our participants through Craigslist, flyers, and word of mouth from a medium-size city in the east coast of US. Their ages ranged from 20 to 40 (Mean=28, Median=29, STD=6.2). Nine were male and six were female. They represented a wide range of occupations (e.g., a librarian, a social worker, an accountant, etc.) and different levels of experiences with the Internet and web tracking blockers.

In the Wizard of Oz study, we asked our participants to complete a set of tasks with the three interfaces (e.g., setting their preferences, and blocking necessary information based on their preferences) in the lab. We observed them when they were working on their tasks, then conducted an exit interview to understand whether they encountered any challenges, how they hoped to improve the design and any other feedback. Given that this study is not the focus on this chapter, we only discuss what we have learned from the process.

Through the Wizard of Oz study, we did not identify any new dimension in the blocker design space. However, our participants' feedback helped us to form a deeper and more thorough understanding of what factors should be included in these dimensions, which further guided our subsequent speculative exploration. We revised the description of each dimension and present them below. For each dimension, we also include examples to illustrate where we identified that dimension from.

Information Dimension refers to the amount of information provided to users. The blocker can provide more information to end-users (about the data each tracker collects, sensitivity, why the data is collected, and who the trackers are), or it might provide less information to end-users (operating more in the background).

Mode Dimension refers to which mode the blocker is based on. The mode of blocking can be based on the type of data being collected (i.e., the information-based model), or who the trackers are (i.e., the tracker-based model), or a hybrid approach of both aforementioned models (i.e., the hybrid model).

Automation Dimension refers to whether the blocker is automatic or dependent on human intervention. If it is the former, then it can automatically figure out who the trackers are and what data and information the trackers are collecting. However, if the blocker is dependent, it requires the cooperation of people and organizations who create these online trackers to disclose what data they are collecting.

Preferences Dimension refers to whether the blocker needs to collect users' preferences. The blocker can allow more user-guided input or be more paternalistic and make decisions without user input.

Categories Dimension refers to the extent to which the information presented to end-users

about what gets blocked is categorized.

4.4.2 Step 2 - Initial Interface Designs

Practically, creating interfaces helped us explore multiple design opportunities and possibilities inspired by the above five dimensions. As a way to explore implications for adoption, we followed Wong and Merrill's work [219] by starting with a purposely mundane design – a series of web tracking settings pages. This mundane form, however, was helpful in getting us to think about how these technologies may get adopted as part of people's everyday practices. As privacy researchers, we also thought that the form of a privacy settings page might make the scenarios more familiar to others in the usable privacy and privacy-enhancing technology communities, as these types of settings pages are common in their work.

We created three scenarios, varying the factors given for each dimension, as summarized in Table 1. Our goal was to produce a diverse set of blocker scenarios to explore different potential effects, rather than to exhaustively explore every possible combination.

We started by going through the dimension list. To reflect the Information Dimension, we first designed the first blocker (S1) in which the blocker is running in the background once it is installed and requires limited interactions with the users, thus has no tangible interface. We then designed the second (S2) and third (S3) blockers to have a user interface.

To reflect the Mode Dimension, given that the tracker-based model has been already widely adopted in real-world, we focused on the information-based model and the hybrid model. We designed S1 and S2 to be the information-based model, and S3 to be the hybrid model.

To reflect the Automation Dimension, we designed S1 to be fully automated since it would run in the background. S2 is also fully automated, and S3 is more dependent on other companies to

reveal their data collection practices.

To reflect the Preferences Dimension and Categories Dimension, we created two permutations. We only required S2 to categorize all the information and S3 to collect users' preferences. As S1 did not have an interface, these two categories were not applicable.

4.4.3 Step 3 - Crafting Scenarios Around the Interfaces

It is worth noting that, while we created and shared a set of speculative interfaces to represent what the blockers might look like in each scenario, we found that thinking about the text of a scenario is helpful in allowing us to sketch out the speculative world that exists beyond the blocker interface—for instance, documenting the organizations and institutions who develop and maintain the blockers, or describing laws that promote the use of blockers.

Before crafting the broader scenarios surrounding the interfaces, we began to ask ourselves and discuss what would be necessary to implement these blockers, what effects and consequences might that have, and how would the effects and consequences influence people's behaviors. Because we were in part interested in the ways individuals' behaviors and other aspects of the world might interact and be affected by the adoption of blockers, we looked to Lessig's framework which includes four regulating forces [134], i.e., law, technical design, economic markets, and social norms that are all ways to affect, or "regulate", human behavior. We used this framework as a basis for thinking about aspects in our scenarios that went beyond an interaction by a human and a computer. We took a dynamic view using Lessig's forces; i.e., people's experience of the mass adoption and use of blockers could affect and be affected by the law, technical design choices, markets, and social norms.

Based on this framework, our speculative exploration focused on how legal changes or technological changes would influence other aspects of the scenarios which eventually may influence people's behaviors. One drawback of using Lessig's model, however, was that the model did not account for individual differences which could impact people's behavior as well, i.e., people in different subject positions or people with different skills and resources may be affected differently by the same technologies or by the same laws. As such, based on what we found from the Wizard of Oz study, we added an additional force, users' digital literacy, in order to capture individual differences in people's understandings of online tracking which might affect their behavior.

When conducting design research, beyond the outcome of creating design artifacts, moments of critical reflection during the *process* of design can provide insight into a phenomenon being studied [178, 190]. As such, attuned to these forces (technology, law, economic markets, social norms, and users' digital literacy), we reflected on different ways that blockers might be *implemented and deployed*, leading us to reflect on three additional design dimensions for the scenarios, including:

Format Dimension refers to the format of the blocker. For instance, is the blocker a third-party browser plugin or an installed browser feature?

Implementation Dimension refers to who is responsible for implementing, updating, and maintaining the blocker?

Interests Dimension refers to in whose interests is the blocker created for? For instance, is the blocker created in the interest of end-users, advertisers, or the developers?

So far, we have iteratively speculated three interfaces and the scenarios around them. Table 4.1 summarizes how these three speculative scenarios differ from each other.

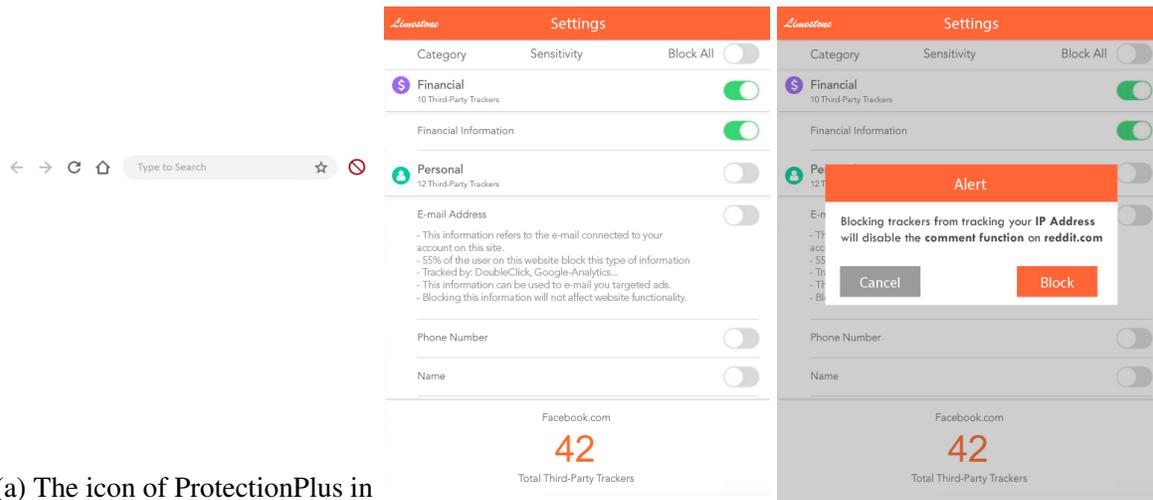
Table 4.1: Different factors in the three speculative scenarios

Dimensions	S1	S2	S3
Information	Background	UI	UI
Mode	Info	Info	Hybrid
Automation	Auto	Auto	Dependent
Preferences	NA	N	Y
Categories	NA	Y	N
Format	Plugin	Feature	Settings
Implementation	Non-profit	Browser	Tracker
Interests	Users	Company	Advertiser

4.4.4 Step 4 - Critically Examine the Speculations

When we finished creating the scenarios and interfaces in this fictional world, we then used the five dimensions (four dimensions from Lessig’s model [134] plus the users’ digital literacy we observed from our prior Wizard of Oz study) to critically examine the scenarios. The goal was to understand what the consequences and implications of adoption might be—particularly if there might be possible consequences that we might not have initially expected. To guide us through the process, we asked ourselves the following questions:

- **Technical design:** Are there any design changes that can lead to the successful or unsuccessful implementation of the blocker?
- **Law:** What will happen if there is a lack of law enforcement or policy regulation to limit the power of industry?
- **Social norm:** What is the intention of the blocker and what are the expectations of the public?
- **Economic markets:** Is there a balance between the needs of users and industry stakeholders?
- **User literacy:** Will the end-users have the knowledge and the ability to use the blocker



(a) The icon of ProtectionPlus in S1. The red icon to the right indicates the plugin has been installed and is currently in use.

(b) Main interface.

(c) Notification on reddit.com.

Figure 4.2: The icon of ProtectionPlus in Scenario 1 (a), and the two interfaces of Limestone in Scenario 2 (b,c).

It is worth noting that, in this chapter, we do not claim the exhaustive or complete representations of scenarios, interfaces, or questions. Traditionally, research in the usable privacy community focuses on designing interfaces, developing functional prototypes, and evaluating them through users studies or field studies to understand the potential user experiences issues or effectiveness in helping users mitigating privacy and security concerns. Our goal in this research is to broaden our understanding of blockers beyond the level of the user interface, to instead surface potential technological, legal, social, and economic impacts that implementing a blocker widely may cause from privacy researchers' perspective. Thus, our scenarios and interfaces, as well as the questions we ask, represent only a subset of all possibilities as a way to begin to surface these questions. Future research can take other routes and understand new issues from other perspectives. Next, we present the three speculative scenarios and interfaces.

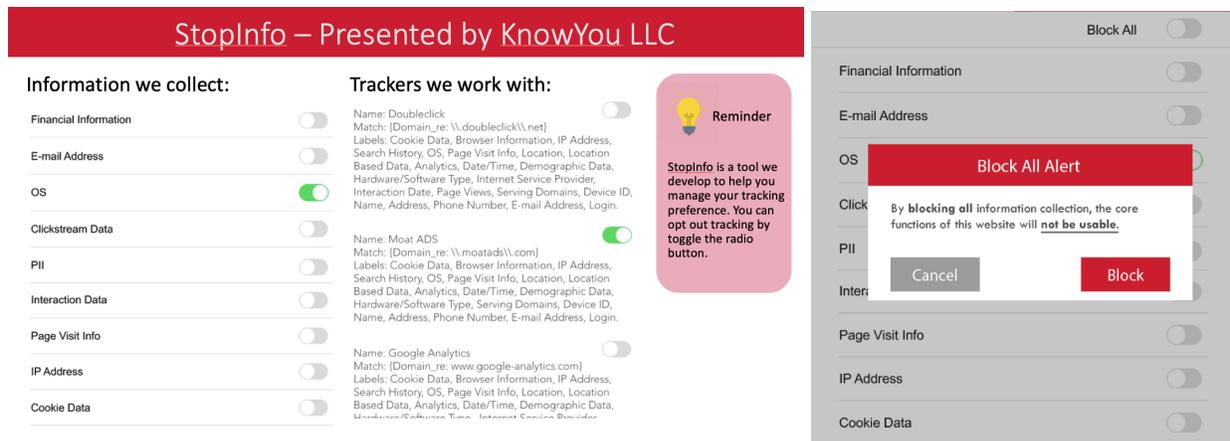
4.5 Speculative Scenarios

Each speculative scenario contains three elements: (1) a background introducing the technical or legal environment, (2) a detailed description of the scenario, and (3) an interface design of the blocker. Below are the details.

4.5.1 Scenario 1: Invisible Blocker

Background. Computer science researchers at EFDN University develop an algorithm that can accurately infer what types of information are being tracked when a user is browsing websites on a mobile or computer device. The work comes to the attention of privacy researchers working in partnership with NoTracking.org, a non-profit organization that is dedicated to limiting web tracking for online users and promoting a more private, equitable web. Together, they develop a user-oriented blocker that has automatic agency in detecting what types of data are being tracked when a user is browsing the web, then blocks sending certain types of data to the tracker. They have also collaborated with the researchers to understand what types of information are sensitive to users in different contexts.

Description. NoTracking.org builds a new web tracking blocker, ProtectionPlus, to help their users limit online tracking. It is a browser plugin available on all the mainstream web browsers. Compared to other existing blockers (e.g., Ghostery [98]) that require users to set up the blocking rules, ProtectionPlus minimizes users' effort by running in the background. No settings or configuration are required; once the users install the blocker, they do not need to worry about it anymore. When users visit a website, ProtectionPlus will show a red icon next to the browser search bar, indicating that it is working (Figure 4.2a). Then it will be able to automatically detect the types of information being collected on that website and block the types of information that



(a) Information and Tracker view.

(b) Notification View.

Figure 4.3: The main interfaces of StopInfo in Scenario 3.

are sensitive to general users based on the organization’s research. When users want to learn more about the information that is blocked, ProtectionPlus will provide more details, such as who collect and why are they collecting the information, on their website.

4.5.2 Scenario 2: Balanced Web Browser

Background. Computer science researchers at EFDN University develop an algorithm that can accurately infer what types of information are being tracked when a user is browsing websites on a mobile or computer device. The work comes to the attention of Limestone, a mainstream web browser company that aims to develop the next-generation web browser to give users more controls of their data. It helps the company to build a positive public image of the company and attract more users. Together, they develop a new feature in the Limestone browser which enables users to perform fine-grained control of their data.

Description. This new feature (Figure 4.2b) allows the users to block the tracking of certain types of information based on their preferences. Once a user blocks on a certain type of information, the browser will pop out a notification telling the user the potential consequence of blocking this

type of information. For example, in Figure 4.2c, when the user tries to block the “IP address” from being tracked, the notification reads, “Blocking trackers from tracking your IP Address will disable the comment function on reddit.com”. The user will then make the blocking decision based on the notification. At the same time, the browser company also needs to work with different third-party tracking companies and advertisers to make a profit, thus they do not expect that the users block everything through this new feature. In fact, the notification tries to nudge users to sometimes not block information from being tracked.

4.5.3 Scenario 3: Greedy Web Tracker

Background. Following the European Union’s General Data Protection Regulation (GDPR), individual states in the United States begin passing data privacy and protection laws, starting with the California Consumer Privacy Act. In order to standardize a dozen different state rules into a single set of laws, the U.S. government faces enough pressure by technology companies and consumer advocates to pass the national Data Accountability and Trust Act (DATA), 15 U.S.C §§9012-9021. Several subsections are of particular importance for online tracking. §9014 c.1.E affirms individuals’ right to have “meaningful choices” in how data about them are collected and used. §9015 a.2 mandates that data collectors and data processors provide machine-readable privacy notices (in addition to existing text-based notices) that disclose what types of information are collected. Liability rules established in §9020 state that a website can be liable if they use a third-party tracker that is not compliant. This leads to a growing marketplace of tracking companies that are DATA-compliant, providing machine-readable data policies and providing options for users to block certain types of data uses.

Description. KnowYou is a web tracking company that collects users’ data when they interact

with the website. The primary goal of the company is to maximize its profit by collecting as much information as possible, then monetize it by analyzing it or directly selling it to other advertising companies. They collect users' data by embedding their web tracker in the mainstream news websites, social media websites as well as shopping websites. Based on the DATA law, all web tracking companies should provide the users the options where users can specify their preferences of what types of information they allow to be collected. KnowYou, therefore, develops StopInfo, a tool to collect users' preferences for information collected through web tracking. They embed the tool in their privacy policy and provide an information view and a tracker view (Figure 4.3a) where users can specify their preferences. In the information view, users can block all types of information collectively or certain types of information selectively. When users click on *block all*, an interface (Figure 4.3b) will show up to notify users of the consequences of blocking all types of information. For the best interest of the company, the fewer users use this feature, the better it will be.

4.6 Reflections of Our Speculative Scenarios

As mentioned before, the speculative scenarios were designed not only to reflect the design dimensions that we have identified from the literature and current practices but also concerns and risks that were not particularly present in the user studies, such as the blocker being designed in the interest of different stakeholders and the consequences it might bring, and the tension between the users and the web trackers. Some of these concerns and risks were ones that we have seen from prior research on web tracking and blocking, other concerns and risks surfaced as we began to think about ways in which multiple types of stakeholders (including different types of users, secondary users, and corporate stakeholders) might adopt and respond to the use of blockers while we were creating the scenarios.

Situating ourselves in the fictional world with the three scenarios we created and thinking about the five questions we posed in the prior section, we unpack the potential issues for the adoption and implementation of a blocker in different scenarios and discuss the implications from legal, technological, social, economic, and users' perspectives. It is worth noting that these issues and implications are not mutually exclusive. Like in the present world, they often intertwine with each other and thus create a complicated situation. For clarity, we discuss each perspective independently.

4.6.1 Commentary: Technological Reflections

In our speculative scenarios, technological changes across different settings may change how people use these blockers. For example, since in Scenario 3, all information presented in the interfaces are not categorized. As such, users may face a lot of usability issues, such as suffering from clotted information. Eventually, the tool may end up not being used as much due to usability problems.

In Scenario 2, the Limestone browser aims to balance the needs of both users and the tracking and advertising companies. By showing the users the consequences of blocking certain types of information, it helps users make better and more informed decisions without interfering with their experiences. The tracking and advertising companies can also collect user information in a non-intrusive way. However, considering the imbalanced power dynamics between the online tracking industry and end-users, Limestone also leaves potential opportunities for tracking and advertising companies or even the websites themselves to work around the goals of the blockers. One possible way is to tie the core functions of the website with the major types of information collection so that if the users block any information collection, they will essentially not be able to use the website. In an earlier example (as shown in Figure 4.2c), if the web trackers embed the tracking of the IP address in the comment function on an online discussion site, then blocking the IP address

tracking will essentially make the site not usable. This practice will force users to allow information collection in order to keep the website usable even when they hope for the opposite. This can lead to a situation where websites design their systems to work in such a way that forces users to have little choice in blocking, i.e., users have to allow all tracking to access a site's functions which is adversarial to the intent of the blockers.

4.6.2 Commentary: Legal Reflections

In the scenarios, we also speculate the potential consequences on users' behavior caused by legal changes. For example, in Scenario 1 and Scenario 2, we did not make any changes to the legal environment. Scenario 3 considers how legal and regulatory changes may shift economic incentives for how companies track, collect and use data, which could change the environment in which trackers operate. However, if changes happen in the law that requires trackers to reveal the information collection details, we might also see potential pushback from the tracking and advertising industry or workarounds to alternative routes given that such changes in the law will hurt that industry's current business model. One potential way to work around the law is that users may be asked to temporarily disable their blockers or disable the blocking function on their browser while they visit some sites, otherwise they will not be able to visit the sites, as this is currently happening with browser extensions that block advertisements (e.g., AdBlocker Plus [171]). Thus the plugin remains not effective, as it will be possible for the tracking and advertising industry to follow laws of disclosing their practices, while still making it difficult for users to opt-out or make choices about tracking. Furthermore, the DATA law's focus on individual choice and consent may not work in cases where data are about multiple people, such as on shared devices. On a more hopeful side, the ability to levy fines against companies for violations may create more incentives to respect users'

privacy or follow their expectations of data collection.

4.6.3 Commentary: Social Reflections

We observed potential social issues that can arise from each scenario, especially when considering whose interest the blocker is designed for and what is the intention of the blocker. For example, in Scenario 1, given that a non-profit organization such as NoTracking.org aims to protect the users' privacy, users may have more trust towards ProtectionPlus. However, since developing and maintaining a technological product requires a significant amount of resources (e.g., funding, technicians, developers to update it, etc.). If NoTracking.org does not have the resources to continue supporting ProtectionPlus and starts to look for resources outside of the organization (e.g., through fundraising), users' trust in an independent non-profit organization may decrease. Or, without resources to maintain ProtectionPlus, the blocker may not get updated, reducing its feasibility as a sustainable long term solution.

In another example, in Scenario 2, Limestone, the browser creating and implementing the blocker, aims to serve the interests of both end-users and the advertising companies. Limestone aims to provide users some level of control over their own data, which can help build a positive image of the company among consumers, and at the same time, aim to provide the advertising companies enough data for their benefits. The motivation for providing advertising companies data could be heightened if Limestone also runs their own tracking or advertising network, as some current browser providers do (e.g., Chrome and web trackers from Google). However, users may have misconceptions that Limestone's goal is to protect them wholeheartedly. This mismatch may defeat users' expectations. If Limestone's role in tracking is revealed to users in a surprising or "creepy" way, it can lead to a strong push back from the users against the company.

4.6.4 Commentary: Economic Reflections

In Scenario 1, even though ProtectionPlus seems to be a very helpful tool for users by automatically blocking trackers, it can potentially damage the web tracking and online advertising industry if adopted at scale. ProtectionPlus may automatically block many types of information that tracking companies need in order to provide advertising services. In a world where people's information is used as a commodity and advertisements are delivered to the users in exchange for free Internet services, without considering the needs of web tracking companies, a tool like ProtectionPlus may hurt the Internet economy and may eventually lose its value. Alternatively, we also considered that online tracking and advertising companies may try to develop workarounds to make blockers less useful to end-users. This can lead to a type of blocking "arms race", where blockers try to find new ways to function, and tracking and advertising companies find new workarounds.

4.6.5 Commentary: Users' Digital Literacy

When we started to think about how users' digital literacy impact their ability to use the blockers in these scenarios, we discovered several interesting insights which were also related to the four dimensions discussed already. In Scenario 1, one potential issue related to the adoption of the blockers can be that users may not even know the existence of ProtectionPlus since as a non-profit organization, NoTracking.org does not have as many resources as the other companies in the industry to promote their product. Thus, users who are generally not tech savvy potentially will not even know about the blocker because they do not know to search for and download it. In fact, we already observed such issues from our Wizard of Oz studies since very few of our participants have heard of any web tracking blocking tools. Similar issues also hold in Scenario 2 and Scenario 3: users who have low digital literacy may experience issues in finding the feature, understanding

the interfaces, and in how to use the tool to protect themselves. Even if installed as a feature of the browser in Scenario 2, users may not take advantage of it, depending on how the Limestone browser implements it (e.g., Limestone may decide to turn off the blocker by default or make its settings page difficult to find). Users with low digital literacy can be put in risky positions considering the imbalanced power dynamics in the web tracking ecosystems where other stakeholders already have significantly more power than users do.

4.6.6 Implications and Questions for Blocking Technologies Adoption

From the above reflections, we synthesize a set of implications and questions for privacy researchers who are creating blockers to consider. These considerations operate at a level above individual interactions with an interface, focusing on how blockers might interact with other aspects of the world over a longer period of time.

- Consider additional direct and indirect stakeholders beyond users (similar to Value Sensitive Design [91]). For instance, what is the role of browsers, developers, regulators, and policymakers, etc?
- Whose privacy is at stake, and what threats is the blocker trying to protect them from (e.g., government surveillance, behavioral advertising tracking, bad web actors, etc.)?
- Should privacy be treated as the only goal? What other values beyond privacy might be at stake? How to understand and achieve the balance between privacy and other goals?
- What legal environment(s) will the blockers be working in? Does that have implications for what the blocker needs to disclose to users, or how the blocker itself is able to work?

- Who is responsible for offering and maintaining the blocker over time? Deployment options such as having the blocker be embedded in a browser or operating system versus a third-party download, or being maintained by a for-profit versus non-profit company might suggest different types of motives for creating the blockers and may affect user trust in different ways.

Moreover, prior literature on information tracking, in general, arises the discussion around the potential consequences that web tracking and blocking technologies could have more broadly beyond the web tracking industry. For example, a recent study has shown that developers perceived the web tracking embedded in the advertising network as the only viable way to monetize their apps [154]. However, most app developers also acknowledge that they mostly choose to use the advertising networks' default settings to collect users' data [154]. As the tracking online and through smartphone share the same advertising networks, the changes we argue for in the tracking and blocking mechanisms can impact how the ads network operates and collect users' data. The ads network will further influence how developers collect users' information in smartphone app tracking and create a significant impact on their lives. More broadly and in the long run, the impact may also affect tracking through the Internet of Things (IoT) as well, given the recent growth in the IoT adoption, as well as many forthcoming new technologies.

Thus, our speculate exploration essentially provides a lens to understand the adoption of blocking technologies in pervasive tracking more generally and its implications on many facets related to the tracking.

4.7 Discussion

4.7.1 Reflecting Our Speculative Exploration Process

As privacy researchers, we realize the limitations of the existing methodological toolbox for privacy-enhancing technologies (PETS) research. Traditionally, PETS related research pipeline includes an empirical study to understand people’s privacy concerns, system prototyping to address people’s concerns, and user studies to validate the system (e.g., [210, 126, 233]). This approach limits the privacy concerns and risks to what the users have suggested in the study rather than what could have happened in the long run when the actual systems are deployed in different real-world settings.

This chapter presented a case study of a potential new privacy research paradigm. We started our exploration by investigating the web tracking blocker design space through literature, current practices, and a Wizard of Oz study. Grounded in the dimensions we identified, we designed different variations of blockers and their interfaces, then created three speculative scenarios where these variations were used in practice. This sequence of explorations not only gave precedence to what users value (e.g., privacy, autonomy) but also surfaced other potential issues and obstacles when implementing web tracking blockers.

Creating a set of speculative scenarios and interfaces helped us build on and move beyond immediate questions about usability and user expectations of the web blocking system. By thinking about “implications for adoption” of a web blocking technology in different socio-technical settings, we began to consider how the blockers might interact with multiple types of stakeholders in various contexts, as well as how they might relate to different technical and legal infrastructures, and how the blockers might be maintained and used over longer periods of time. These reflections suggested social, technical, and legal implications for researchers who are working on blocking

tools to consider, which were not obvious from the results of the traditional user studies alone. As the responsibility for addressing privacy spans social, technical, and legal domains, expanding researchers' considerations to include these domains helped us work towards a fuller and richer set of ways to try to protect privacy. This case study suggested that PETS research would benefit from considering implications for adoption through speculative design-inspired methods.

4.7.2 Speculative Design of PETS

We recognize that the speculative design methodology has rarely been used within the privacy and security community. To the best of our knowledge, our study as well as Briggs and Thomas' study on future identity technologies [36] are some of the few uses of speculative design approaches in privacy and security research. In contrast, lots of prior work in the design research community has touched on privacy and security issues related to technology development (e.g., [222, 170, 141, 139]).

Based on our own experiences, we believe that speculative design is a valuable addition to privacy/security researchers' methodological arsenal, and can help put privacy/security research in conversation with design researchers similarly grappling with issues in the domains of privacy and security through different methods. The mainstream research paradigm in privacy and security usually involves researchers building functional prototypes and then running user studies to evaluate the prototypes. The focus is often on the technical feasibility and user experiences. In contrast, speculative design purposefully ignores the technical feasibility in exploring the futuristic and technically challenging ideas. It allows people to critique the status quo, imagine and experience an alternative future(s) where they are free from the current market and technology limitations, and raise questions for future technology development. It also helps to surface other potential issues and consequence that are not obvious through traditional user studies. Creating alternative futures

also allows people to pursue the values they espouse—or alternatively, to explore values in possible futures that we might want to avoid. This is particularly valuable in the context of online tracking or the Internet more broadly because stakeholders’ relationships are power-laden, where usually the industry has more power over ordinary users. This power imbalance shapes current online tracking practices. Future speculative design work might explore what privacy looks like when power is distributed differently among users and technology companies.

We advocate that on the one hand, researchers interested in the privacy and security aspects of the technology in the design community can proactively work with researchers in the privacy and security area to better facilitate and deepen the related aspects in their design work. On the other hand, the privacy and security community should broaden their methods and consider adopting speculative design methods in exploring future privacy-enhancing technologies. In particular, these approaches can encourage researchers to think outside the box by not limiting themselves to the industry or market demands or current practices and to raise empirical research questions for future research. Furthermore, it broadens the scope of the investigation, helping researchers think of further downstream effects related to the ways in which the technologies interact with other aspects of society beyond the technologies’ users. It can also help researchers to think more critically about contemporary technologies and shed lights on who is truly benefiting from these technologies. For example, this method can be applied to explore PETS for future automation technologies (e.g., privacy mechanisms for self-driving cars), smart toys (e.g., empowering children and parents to better control their privacy against ubiquitous data collection), social media (e.g., recognizing imposters immediately to prevent loss), and big data (e.g., profiling users through automatic algorithms).

4.8 Conclusion

Online tracking is prevalent and many tracking blockers have been developed to help people manage online tracking. We used a speculative design approach, complemented by an Wizard of Oz study, to understanding the implications of adopting web tracking blockers. Through three speculative scenarios and several interfaces, our exploration surfaces many potential issues that may come with the adoption of web tracking blockers, and further sheds light on the technical, legal, social, economic, and user-related implications for future web tracking blocking tools. Our research presents a case study of a new paradigm for future research in the privacy and security community. Our approach also helps to discover deeper privacy issues hidden in the user studies and illuminate future empirical research questions to ask. We advocate a deeper collaboration between the design community and the privacy and security community for mutual benefits.

Part III
Case 2: Drones

Chapter 5

Drone Privacy: Bystanders' Perspectives

As an emerging technology, drones represent a simple multi-stakeholder environment. On the one hand, drone controllers fly the drones and take pictures or record videos with it; on the other hand, drone bystanders (i.e., those who are not the controllers, such as people who are walking on the street) can potentially be caught by the camera on a drone. To examine people's privacy perceptions of drones, we started with an interview study to investigate the privacy perceptions of drone bystanders [228]. In this section, I summary our study design, data analysis, as well as main findings. For full results, please refer to Appendix A.

5.1 Study Design

Each study session started with a drone demonstration. When the participants came to our lab for the study, we first showed them a drone (DJI Phantom 2 Vision+) statically and explained in detail how to control a drone to fly and take photos. If the weather permitted, we would demonstrate the drone motion by operating it in front of the participants. We then answered any questions the participants might have related to the drone and its capabilities.

We started the interview with the participants immediately after the demonstration. The interview protocol is structured as follows. We began by asking the participants some general questions related to their perceptions of drones, such as “have you ever heard of drones? How do you feel

about drones? Do you see any benefits of drones?” We also asked our participants to compare a drone with two other tracking/recording technologies, smartphones with cameras, and closed-circuit television (CCTV). We then presented five scenarios to the participants, including a drone is being used in recording a promotion event in a shopping mall, delivering goods, recording a friend’s party, reporting a parade, and searching criminals. We asked our participants to explain their perceptions of drones in each scenario. Finally, we asked our participants to describe what kinds of notifications and controls they would like to have and how the drones should be regulated.

5.2 Data Analysis

We conducted interviews with 16 participants (eight male, eight female, 18 - 62 years old, average 29 years old). All interviews were recorded then transcribed. We conducted a thematic analysis [32]. The other co-author and I first immersed ourselves in the data by reading through the transcripts several times. We then coded the transcripts, discussed the codes, and merged our codebook to form a final codebook with 132 unique codes. The codes were then grouped into nine themes, including drone features, drone usage, attitudes towards drones, cultural differences, private vs. public space, privacy concerns, safety concerns, and drone control. The details of the process and the sample codes are documented in the published paper [228] in Appendix A.

5.3 Results Summary

Our participants suggested several perceived benefits of drones, such as the ability to fly and to take high-definition photos and videos in inaccessible or even dangerous environments. They also mentioned many potential application domains, such as aerial photography, package delivery, and emergency responses in the high-traffic area. At the same time, our participants also discussed

their safety, security, and privacy concerns of drones. The safety concerns were mainly about the possibility of hitting by a drone or drones interfering with each other. The security concerns mainly centered around drone trespassing on some sensitive or even forbidden areas (e.g., military facilities) due to its mobility. As of the privacy concerns, Table 5.1 summarized the key findings related to our participants' privacy concerns. The full results and examples of quotes can be found in Appendix A.

High-level Findings	Sub-themes	Key Factors/Findings
General privacy perceptions	Public vs. private spaces	Ownership, sensitivity of the place, and nature of activity are three factors to determine whether a place is private or public.
	Peeking and stalking	Drones may peek through the window, or follow and record an individual's activities.
	Recording and sharing	Drones may take and share photos and videos without bystanders' consent.
Context-based perceptions	Shopping mall monitoring	Most participants perceive a shopping mall as a public space, thus have less concerns. But drone usage should be limited.
	Recording a friend's party	It is not clear whether a friend's house is private or public, thus participants have mixed attitudes.
	Delivering goods	Generally acceptable
	Reporting a parade	Generally acceptable because it is a public space
	Searching for criminal	Generally acceptable, but some participants felt uncomfortable if the search is near their house
Comparison with other tracking technologies	Drone and camera phones	The major difference is the distance of recording, and whether the bystanders can access the owners or controllers
	Drones and CCTV	The flexibility, visibility, intended purposes, and the trust in the controllers influence people's privacy preferences
Expected notification and control	Tracking drone controllers	Drone controllers should be required to register
	Tracking and controlling drones	Each drone should have a unique ID
	Regulations	Regulation should limit the physical dimension of the drone and the area that drones can fly in
Key Discussion Points	Duality of drones	Privacy concerns are not only about the drones but also about the drone controllers and the perceived relationship with the controllers.

Table 5.1: Summary of the privacy perceptions of drone bystanders

Chapter 6

Drone Privacy: Controllers' Perspectives

Prior research has discovered various privacy concerns that bystanders have about drones. However, little is known about drone controllers' privacy perceptions and practices of drones. Understanding controllers' perspective is important because it will inform whether controllers' current practices protect or infringe on bystanders' privacy and what mechanisms could be designed to better address the potential privacy issues of drones. In this chapter, we report results from interviews of 12 drone controllers in the US. Our interviewees treated safety as their top priority but considered privacy issues of drones exaggerated. Our results also highlight many significant differences in how controllers and bystanders think about drone privacy, for instance, how they determine public vs. private spaces and whether notice and consent of bystanders are needed.

6.1 Introduction

Drones are lightweight unmanned aircraft controlled by operators or onboard computers. Drones can enable numerous innovative applications but have also raised significant privacy concerns due to their maneuverability and capabilities of taking photos/videos and sensing the environment. For instance, in our prior work, we interviewed *drone bystanders* (i.e., people who had no experience operating drones but may be surrounded by flying drones) and found that bystanders had various privacy concerns about drones such as stalking, photo/video recording and sharing [228].

However, it is unclear how *drone controllers* (i.e., people who have directly operated drones) think and do about privacy in their practices. Answering this question is important because it will inform (1) whether controllers' current practices may protect or violate bystanders' privacy, and (2) what mechanisms could help controllers better address these privacy concerns.

As a follow-up study of drone bystanders [228], we conducted interviews with 12 drone controllers in the US about their perceptions and practices of drones. We focused on drones for civilian rather than military purposes. Our controller interviewees were primarily concerned about safety and felt that the privacy risks of drones are exaggerated. Most of them also believed that they have the rights to fly drones and take photos/videos in public spaces without the need to get others' permission. While they adopted a legal definition of public space that is primarily based on ownership, our prior study of bystanders found that some bystanders followed a "social" definition in which the nature of a space is characterized by the social relationship within it, for instance, shopping with a close friend in a mall makes the space private [228].

This chapter makes two main contributions. First, it provides a rich account of the perceptions and practices of drone controllers. Second, comparing our results with the prior literature on bystanders [228] uncovers significant privacy "mismatches" between drone controllers and bystanders. We discuss future directions to help bridge these mismatches.

6.2 Related Work

Privacy issues of drones have been discussed in the literature. For instance, the Electronic Privacy Information Center (EPIC), a civil liberty organization highlights aerial surveillance as a critical privacy issue of drones [80]. Legal experts have also voiced ethical and privacy concerns regarding the use of drones (e.g., [43, 72]). Dunlap argues that when drones are used for surveillance, they

can violate Americans' constitutional rights, particularly citing the Fourth Amendment, which protects people from unreasonable searches and seizures [72]. As for regulations in the U.S., the Federal Aviation Administration (FAA) requires drone controllers to register their drones with the agency [1]. The FAA's new rules on small drones focus on safety (e.g., prohibit night operations of drones) [84].

People's privacy concerns of tracking/recording technologies, such as wearable cameras (e.g., [106, 66, 108, 107]), CCTV (e.g., [216]), and RFID (e.g., [14]), have been studied extensively in the literature. However, privacy issues of drones are understudied. Clothier et al. conducted a survey in Australia and found that the respondents' overall attitudes towards drones were fairly neutral but less than one fifth of the respondents reported concerns about drone surveillance or spying [52]. Our prior study of drone bystanders in the US uncovered different privacy concerns about drones in general and under specific drone usage scenarios [228]. While these two studies focused on ordinary citizens or bystanders (i.e., who did not have experience operating drones), the literature says little about drone controllers. Our present study fills the gap by focusing on drone controllers' privacy perceptions and practices.

6.3 Methodology

From January to March 2016, we conducted semi-structured interviews with 12 drone controllers in Syracuse, New York (US). We recruited our interviewees by posting study fliers in places such as university campus and parks. The interviews were conducted in a drone hobbyist club, our lab, and public places such as libraries. Each interview took about 1 hour with a payment of \$10. 10 of our informants were male, and two were female. Their ages ranged from 20 to 62 years old with an average of 28. They also presented diverse occupations, including college students, a professional

photographer, a tax officer, an office administrator and a retired worker.

To ensure the validity of comparison between this study and our previous study of bystanders, we adopted the same interview protocol and data analysis approach [228]. We re-framed many interview questions to focus on the perspective of drone controllers. The interview consists of general questions about drone controllers' perceptions of drones (e.g., "Do you see any benefits or drawbacks of drones?"); purposes and practices of using drones (e.g., "Where do you fly your drone(s)?"); expected notice and control (e.g., "Do you feel people should get others' (e.g., bystanders) permissions before flying a drone or taking pictures/videos?"), and controllers' attitudes towards drone usage under different scenarios. We used the same five drone scenarios based on real-world drone usage from our previous study of bystanders [228]. These scenarios are (1) recording a promotion event in a shopping mall by a store owner; (2) delivering packages by Amazon; (3) recording a friend's party; (4) reporting a parade by a news agency; and (5) searching suspects in a residential area by the local police. We asked our interviewees if they would operate the drone as described in each scenario and why.

We audio recorded the interviews upon informants' permissions. The interviews were then transcribed and analyzed qualitatively. Similar to our previous study [228], we conducted a thematic analysis. 12 Two co-authors (coders) used ATLAS.ti, a popular qualitative analysis software, to manually and independently generate initial codes that capture meanings of the same subset of our interview data at a fine-grained level (usually at the sentence level). Then, the two coders convened, discussed, and converged their codes into a code book of 135 unique codes ranging from drone usage (e.g., photography) to privacy concerns (e.g., identifying people) to drone community (e.g., irresponsible controllers). Next, the coders used the agreed-upon code book to code the interview data. The inter-coder reliability was 0.85. We grouped 135 codes into ten themes: drones in

general, drone usage, safety concerns, privacy concerns, permission, private/public spaces, scenario questions, application design, drone community, and regulations.

6.4 Findings

We present major themes from our interviews and use pseudonyms for our interviewees.

General perceptions of drones. All of our interviewees have flown a drone themselves and most of them own a drone. Overall, they were passionate about this emerging technology. However, some of them preferred not to use the word “drone.” For instance, Mike (28, male, drone hobbyist) avoided the term “drone” because it comes with certain connotations from which he wanted to dissociate. He said, *“usually I associate the word drone with the military drone that performs air strikes and things like that.”* So, instead, he used the term “quad” as he further explained, *“I still try to insist on calling them quads because I won’t fall into what the media has done in calling them and dubbing them drones.”* By using a different and arguably more neutral term “quad,” Mark deliberately separated himself from the sensitive military use of drones and the media’s newsworthy and often controversial accounts of drones (e.g., a drone crashed on the White House lawn [114]).

General use of drone. When asked why they fly drones, many interviewees mentioned that radio-controlled drones are just fun to fly. Other interviewees talked about drones allow them to pursue their personal interests such as photography and DIY (do-it-yourself) projects. Our interviewees noted that safety is their highest priority in drone operations, including the safety of both drones and people. Our interviewees reported relying on their common sense (e.g., avoid flying near a crowd) and caution when operating drones. They were also thoughtful about where to or not to fly their drones. They reported usually flying in public parks and deliberately avoiding places such as surrounding areas of airports and schools with children.

Taking and/or sharing photos/videos. Many interviewees also used their drones to take pictures/videos, mostly for landscape, as Mike explained, “*they’re actually mostly landscape, that’s sort of what I’m interested in as a sort of photographer.*” They kept the photos for personal use, or share with their friends in social media. For example, Tim (23, male, electronic engineer student) described, “*I share the photos/videos on Instagram [@***] and Facebook.*” When asked whether the photos captured bystanders, Tim continued “*I would imagine that I have taken photos/videos with bystanders in them, but nothing extremely close where someone watching the video could recognize anyone. I have shared these photos/videos online, or have given them to friends who want to use them.*” These drone practices may explain their privacy perceptions.

6.4.1 Privacy Perceptions

While our interviewees valued privacy, they also felt that the privacy concerns about drones are exaggerated and even misguided by the media’s sensational reports of drones. They framed their opinions mainly along the lines of public vs. private spaces and claimed that they have the (Constitutional) rights to fly drones in public spaces.

Public vs. private spaces. First, our interviewees had fairly consistent views of public and private spaces based on ownership. Ross (26, male, software engineer), for example, described, “*I define [public space] as like places that are generally open to the public like anybody can walk by and it’s open and free to access.*” Tim put it more bluntly: “*the way I think about it is that public space is public space, you can do whatever you want in public space.*” For private space, Mark had a typical definition, “*the space that’s owned by a particular person or particular business entity then that’s private.*”

However, our interviewees differed in their interpretations of a specific case. Some defined

the airspace over people's personal property as public space, as Dan (23, male, media student) argued, *"my definition is aligned with the legal definition of airspace ownership actually. So as I currently understand that you might own your house and the land, however, you do not own the airspace above you."* While this is legally true in the US, others were more sensitive socially. For example, Jake argued, *"generally I look at the fenced off area as private property and even though you don't really own the airspace above your property it's still not a very nice thing to do."* Jake's comment points to the nuanced social expectations of privacy and his sensitivity in respecting such expectations.

No privacy expectations in public spaces. Most interviewees generally felt that people should have expectations of privacy in private spaces but not in public spaces. Dan was quite vocal about this, saying *"I think the American public needs to really understand that when they're in public they shouldn't expect any privacy, that's just pure and simple."* As a result, our interviewees usually did not ask for people's permission before flying their drones in public space. Alex, for instance, had a clear view: *"So if you are in a public property, you should not get permission from anyone because you are flying in an area that is open to everyone, so you are allowed to use it like it's open to everyone."*

Some of them even stressed that they have the legal rights to take photos in public spaces without people's permissions. For instances, Alex reasoned, *"if you think about First Amendment, we are allowed to photography [sic], as a photographer, anyone in anywhere as long as you are in public property."* Sasha (22, female, photographer) held the same belief, explaining *"again I'm protected by my First Amendment to take, as a free press, to take a picture so I do not. You know, if I'm having a nice day I will tell you, but I know the law, therefore if I don't I'm okay with my own conscience for not asking."* Among other things, the First Amendment to the US Constitution prohibits abridging

the freedom of speech and infringing on the freedom of the press.

Since we do not have the legal expertise, we consulted American lawyers regarding the validity of the above interviewees' claims. Both lawyers disagreed with interviewees' claims regarding their constitutional rights to fly drones in public spaces partly because airspace is not a public space. The US government has "complete and exclusive national sovereignty in the air space" over this country.

In contrast, interviewees mostly agreed that people should have privacy in private space. However, one interviewee, Dan held the view that *"If you don't want your indoor activity to be reviewed because there is a drone outside of your window, I'm sorry you just have to put a curtain down."*

Ask for permissions. While our interviewees claimed that they do not need to ask for permissions to fly drones in public spaces, in practice they do sometimes. Jake told us one story, *"the last time I asked people was over the bridge in [a town], there were two people out there fishing and I said hey I want to bring my quad up and film it, do you mind if I get you guys."* Jake got the permission to film but he explained why he asked, *"I don't want to annoy people and I mean you really can't get into trouble per se...but they can scream at you, they can yell at you, they can be very violent to you."* In this case, Jake asked for permission to show politeness and to avoid unnecessary confrontation. But they were also practical about asking permission, as Jake illustrated, *"When there's a large crowd, it's not worth taking the time to ask everyone individually, just avoid."*

Drones vs. DSLR. Some interviewees felt the accusation of using civilian drones for spying is misplaced. They compared drones with other photographic equipments, such as DSLR (Digital Single-Lens Reflex) cameras. For instance, Alex explained, *"Drones don't have telephoto lenses which can zoom in as well as shoot from the ground, which I think it's sort of misguided privacy concern people have because you can do more damage in privacy, invading privacy more with DSLR and telephoto lens."* Mike stressed the intent of usage, *"So if the intent is the same, you know*

the equipments might not be that different by the way of doing it.”

Scenario-based perceptions. Besides questions about controllers’ general perceptions of drones, we also provided specific scenarios to further investigate their acceptance of drone usage in each scenario. Our interviewees’ responses were relatively consistent across scenarios. They indicated that they would fly the drone in the scenarios. Their decisions were mainly based on whether they have the permission to fly. For instance, for the mall scenario, they said they would fly if they have a Section 333 exempt, which allows individuals to fly drones for commercial purposes in the U.S.

Drone regulations. While our interviewees reported being reasonable and considerate in their own drone usage, they almost unanimously suggested the need for some form of drone regulation. Our interviewees supported the FAA drone registration requirement. Furthermore, some interviewees suggested having a drone license. Others disagreed, for instance, one interviewee argued that drones are not deadly weapons, thus he did not need a license. In addition, some interviewees also advocated for drone controller training so that controllers can know more about how to fly and how to be safe and skillful.

Construct a positive community identity. Several interviewees talked about the overall image of the drone community and disdained irresponsible/reckless drone usage. Jake shared his past experience, *“I drove by the prison and I actually saw a drone like hovering over the prison. That’s definitely something very stupid . Very stupid and that could increase laws for us.”* Jake’s concern highlights the potential externalities (e.g., stricter laws) as a result of some drone controllers’ irresponsible behavior. John agreed, *“they don’t have license, everybody can fly. Then some people would do something stupid that can damage the whole community.”* With drone registration, license and training, their hope was that the drone controller community as a whole will behave responsibly

which can improve the public perceptions of drones and the community.

6.5 Discussion

Our prior study focused on drone bystanders [228], while this study focused on drone controllers. Both groups are important stakeholders of the drone ecosystem and their perspectives are useful in understanding the privacy implications of drones. We highlight many notable differences between the two groups.

Controllers vs. bystanders. First, while bystanders are fine with calling this emerging technology “drones,” some controllers deliberately use “quads” instead of “drones” to avoid the controversial or negative connotations of drones, which are often associated with military drones for spying.

Second, controllers are generally positive or even enthusiastic about this emerging consumer technology of drones. Their highest priority is ensuring safety of drones and people. In comparison, bystanders have mixed feelings about drones. They see potential benefits and applications of drones, but they are also concerned about safety, security and privacy issues that drones can pose.

Third, bystanders have several privacy concerns about drones, such as peeking and stalking as well as taking and sharing pictures/videos. They are also concerned that they may not see the flying drones and their controllers, which limit their abilities to communicate their privacy preferences (e.g., drones not taking pictures/videos that capture them) to the controllers. In contrast, most controllers feel that the privacy issues of drones are exaggerated because they value others’ privacy and rely on their common sense to operate drones appropriately. They also debunk the perceptions that drone cameras can easily and clearly capture people’s faces from the air.

Fourth, when determining their acceptance of specific drone usage, both bystanders and controllers consider whether the drone is operating in a public or private space. However, their definitions of public/private space differ. Controllers mainly use the ownership of a place to differentiate public vs. private space. They believe that private space is legally owned by a private entity (e.g., people's houses), whereas public space is owned by the public (e.g., parks). This type of understanding is more aligned with the legal definitions of public/private spaces. Bystanders' definitions of private space and public space rest on three factors: ownership, sensitivity of the place, and nature of activity in the place. In particular, some bystanders characterize spaces based on the activities and social relationships within the space. For instance, shopping with a close friend in a mall would make it a private space because of the close personal relationship between friends. This highlights the "social" definition of space.

Fifth, when considering the specific drone usage scenarios, bystanders consider three criteria: (1) whether the drone is operating in a public/private space; (2) what is the intended purpose of the drone usage; and (3) notification and consent of the drone usage. In comparison, controllers mainly focus on whether they have the permission to fly the drone in the scenario (e.g., need a permit for commercial use of drones).

Sixth, several bystanders expect to be notified and asked for their permission before a surrounding drone taking pictures/videos even in a public park. While controllers may sometimes do that to be polite, many of them believe that they have the constitutional rights (citing the First Amendment) to fly drones and take pictures and videos in public spaces without getting others' permission. The National Telecommunications and Information Administration (NTIA) recently released a document of voluntary best practices for commercial and non-commercial use of drones, for instance, "If you can, tell other people you'll be taking pictures or video of them before you

do” [161]. Controllers’ belief that they are entitled to freely operating drones in public space may dissuade them from following these best practices.

These mismatches between controllers and bystanders are perhaps not surprising given that they have different roles and interests in the context of drones. But, these mismatches can lead to tensions especially when bystanders’ privacy concerns about drones are not adequately addressed by controllers.

Mitigate bystanders’ privacy concerns. One future direction is to enable bystanders and controllers to communicate directly so that bystanders can express their privacy concerns or preferences and controllers can explain their drone usage (e.g., purpose). Direct communication can help bridge some of the mismatches between the two groups. Since bystanders may not see drone controllers, enabling electronic communication channels (e.g., via a website or a mobile app) would be useful. In addition, NTIA’s best practices of drones is a good step towards educating controllers about potential privacy risk of drone usage and practical strategies to mitigate these risks. However, these best practices are voluntary so controllers may not adopt them. Many of our interviewees expressed their hopes to create a positive image of the overall drone community. These best practices and other privacy mechanisms can be emphasized as improving the image of the drone community, which can help incentivize adoption.

Study limitations. First, our study has a relatively small sample size. While we only interviewed 12 drone controllers, we did not find any significantly new findings from our last three interviews. Second, we only interviewed controllers in the US, therefore it is unclear whether our findings would be similar for controllers in other countries. Third, our interview data is self reported and might be subject to the social desirability bias. In particular, our controller interviewees may withhold sharing their drone practices that can be considered as privacy invasive, for instance, taking pictures

that capture a bystander's face and making the pictures public online. Lastly, self-reported data can divert from actual behavior [194].

6.6 Conclusion

We interviewed drone controllers in the US to understand their privacy perceptions and practices of drones. The results suggest that they treat safety as their highest priority but consider privacy issues of drones overstated. Comparing with our prior study of drone bystanders, we highlight important mismatches between controllers and bystanders on how they view drone privacy. Future work should explore how to bridge these mismatches and mitigate bystanders' privacy concerns.

Chapter 7

Privacy Mechanisms for Drones: Perceptions of Drone Controllers and Bystanders

Drones pose privacy concerns such as surveillance and stalking. Many technology-based or policy-based mechanisms have been proposed to mitigate these concerns. However, it is unclear how drone controllers and bystanders perceive these mechanisms and whether people intend to adopt them. In this chapter, we report results from two rounds of online survey with 169 drone controllers and 717 bystanders in the U.S. We identified respondents' perceived pros and cons of eight privacy mechanisms. We found that *owner registration* and *automatic face blurring* individually received most support from both controllers and bystanders. Our respondents also suggested using varied combinations of mechanisms under different drone usage scenarios, highlighting their context-dependent preferences. We outline a set of important questions for future privacy designs and public policies of drones.

7.1 Introduction

Drones are unmanned aircraft that can be controlled remotely by human controllers or operated autonomously by onboard computers. In recent years, drones have entered the mainstream consumer market. This type of drones often carry cameras and possibly other sensors such as GPS,

accelerometers as well as altitude, temperature and infrared sensors. Drones enable innovative applications but also raise privacy issues. For example, the Electronic Privacy Information Center highlights surveillance as a key privacy issue of drones [80].

In the U.S., the National Telecommunications and Information Administration (NTIA) released a document of voluntary best practices for commercial and non-commercial use of drones, for instance, having a privacy policy that explains an organization's use of drones [161]. Many technical privacy mechanisms for drones have also been proposed. For instance, LightCense uses LED lights on a drone as its ID so that people could identify the drone and its information via a mobile app [136]. However, most of these technical or policy-based mechanisms are voluntary and thus it is unclear whether people will adopt them and even if adopted, whether they would be effective.

In this chapter, we focus on *how drone controllers and bystanders perceive these technology-based or policy-based privacy mechanisms for drones*. We define *drone controllers* as people who have operated drones and *bystanders* as people who have not operated drones but could be surrounded by flying drones. This research question is timely and important because if people perceive these mechanisms as requiring too much effort, being impractical or ineffective, they are unlikely to adopt these mechanisms. As a result, people's privacy concerns about drones may remain largely unaddressed, potentially hindering the acceptance and adoption of drones and limiting their benefits to society. Privacy mechanisms that are supported by both drone controllers and bystanders have great potential to be adopted and useful in practice.

To answer the research question, we developed detailed descriptions of a diverse set of representative privacy mechanisms for drones and conducted two rounds of online survey to investigate how drone controllers and bystanders perceive these mechanisms. In this research, we focus on drones that are used for civilian purposes, excluding military usage. We found that when considering

individual mechanisms, owner registration and automatic face blurring received most support from both controllers and bystanders. However, under specific drone usage scenarios, our respondents also suggested using multiple mechanisms together as they may contribute to different aspects of privacy. But, the choices of mechanisms varied across different scenarios.

This chapter makes three main contributions. First, it sheds lights on how drone controllers and bystanders think about different types of privacy mechanisms for drones. Second, it not only discusses ways to improve these specific mechanisms but also outlines important questions for future privacy designs for drones. Third, it makes a public policy contribution. While most of the studied privacy mechanisms are currently voluntary, they could become mandatory in the future. The findings on people's attitudes towards and perceived effectiveness of these privacy mechanisms can inform the policy development, for instance, mandating certain mechanisms.

7.2 Related Work

7.2.1 Perceptions of Tracking and Recording Technologies

Since drones often carry cameras, they are a type of tracking and recording technologies. Prior studies have identified people's privacy concerns (e.g., leaking personal information) about various tracking and recording technologies, such as Radio-Frequency Identification (RFID) tags [14], credit cards and store video cameras [163].

Prior research has also studied people's perceptions of wearable devices (e.g., glasses). For instance, Denning et al. find that people expect giving their permissions before Augmented Reality (AR) glasses can record them [66]. These wearable devices can also enable "lifelogging" where photos or videos can be automatically taken by these devices (e.g., SenseCam [105]) in a person's everyday life. Hoyle et al. find that people have various privacy concerns about lifelogging, for

instance, sensitive information such as lifeloggers' locations or credit card numbers as well as bystanders' faces or behaviors appearing in the "lifelog" [108]. In addition, robots equipped with cameras can also be considered as a type of tracking and recording technology. For instance, Butler et al. find that people desire mechanisms to help protect their privacy in the presence of remotely tele-operated in-home robots [42]. Our work adds to this literature of tracking and recording technologies but focuses on drones.

7.2.2 Privacy Issues of Drones

Similar to other tracking and recording technologies, legal scholars have argued that drones can infringe on citizens' privacy. For instance, Dunlap posits that when drones are used for surveillance they can violate the Fourth Amendment of the US Constitution, which protects citizens from unreasonable searches and seizures [72]. Wright et al. raise heightened concerns about drones due to the fact that drones could be cheaper to obtain than before and could be so tiny, albeit equipped with high-definition cameras (e.g., "dragonfly drones") [225]. As a result, drones could take detailed pictures of people and it would be difficult for people to notice the drones and to realize they are being recorded by the drones [225].

There are few empirical studies of drone privacy. A survey of Australians' perceptions of drones find that their respondents did not consider drones to be overly beneficial or risky, but some respondents (less than one fifth) did raise a general privacy concern about drone surveillance or spying [52]. Our previous interview study of drone bystanders find that they had various privacy concerns about drones and their perceptions of drones varied in different scenarios [228]. Unlike these prior studies, we focus on specific privacy mechanisms for drones in this research.

7.2.3 Privacy Mechanisms for Drones

Several mechanisms have been proposed that either directly or indirectly protect people’s privacy against drones. For instance, traditional “sense and avoid” systems for drones have been re-designed so that minimum personal data will be retained by the drones [15, 95]. In addition, B4UFLY, a mobile app, was designed to help drone controllers “determine whether there are any restrictions or requirements in effect at the location where they want to fly” [82]. Besides, a type of geo-fencing was proposed to allow individual citizens to designate their addresses as drone no-fly zones, which can be incorporated into the software or firmware of drones and/or honored by drone controllers [167]. LightCense uses a blink sequence of LED lights on a drone as its ID [136]. To learn information about a particular drone, people can use a mobile app to scan the blinking light sequence to identify the drone and look up its information [136]. There are also server-side privacy mechanisms for drones. For instance, Yoohwan et al. propose a system that enables encryption, access control, and image/video transformation of drone recorded data [120]. The NTIA recommends a number of voluntary best practices for drone usage, such as informing bystanders before drones taking pictures/videos if possible [161]. However, it is unclear how people perceive these mechanisms and whether they will be adopted. To fill this gap, we surveyed drone controllers and bystanders, asking them to assess a diverse set of privacy mechanisms for drones.

7.3 Methodology

We conducted two rounds of online survey of drone controllers and bystanders. Both surveys focused on respondents’ assessment of specific privacy mechanisms for drones. After conducting survey one, we learned many things that can improve the survey. For instance, many respondents felt the descriptions of privacy mechanisms were not detailed enough and they raised many questions

about the specifics. Therefore, we conducted survey two, which was very similar to survey one but differed in three main aspects: making descriptions of privacy mechanisms more detailed, testing a slightly different set of privacy mechanisms, and including specific drone usage scenarios and demographic questions.

We recruited survey respondents from Amazon Mechanical Turk (MTurk) where workers were based in the US and had at least 95% task acceptance rate. We also recruited respondents from drone user forums such as the DJI forum and the Quadcopter.com forum. We conducted survey one during March 2016 and received a total of 456 valid responses including 385 bystanders and 71 drone controllers. We conducted survey two during August 2016 and received a total of 430 valid responses including 332 bystanders and 98 drone controllers. Each valid response from MTurk was compensated for \$2. We had about 100 controller respondents from drone forums and administered a raffle of four \$50 gift cards. This research was approved by the Syracuse University IRB office.

7.3.1 Survey Flow

For both surveys, we first provided a working definition of drones as “an unmanned aircraft guided by remote control or onboard computers” and a photo of a DJI Phantom 2 as an example drone. We also told the respondents to focus on civilian not military uses of drones. Next, we asked “have you ever flown a drone yourself?” If a respondent answered yes, then he or she would answer the controller branch of the survey; otherwise, answer the bystander branch. We told controller and bystander respondents to answer the remaining questions by representing themselves as a controller or a bystander, respectively.

We then provided each respondent descriptions of a set of privacy mechanisms in a randomized order. For each mechanism, we asked respondents to answer three questions (5-point Likert scale):

“How practical do you think this mechanism will work? Are you willing to use this mechanism if it is implemented? If this mechanism is implemented, how effective do you think it will protect people’s privacy regarding drones?” Respondents were then asked to provide open-ended answers to explain their ratings. The Likert-scale questions were inspired by our prior interview study of bystanders where the interviewees talked about practicality and effectiveness of, as well as effort/willingness to use a privacy mechanism when they proposed ways to address their privacy concerns [228]. We checked the open-ended answers and found them to be consistent with the corresponding Likert-scale ratings, suggesting the Likert-scale questions were understood correctly.

7.3.2 Survey One

We created brief descriptions of six privacy mechanisms based on the literature and industry proposals. These mechanisms varied by their types (e.g., technology vs. policy, proactive vs. reactive). Below are the descriptions (bystander version). We denote each mechanism in a format of Name (Short name).

Deletion request (Delete): Drone controllers can receive requests from me to delete photos or videos that capture my family, properties or myself via a mobile app [228]. **Gesture opt-out (Gesture):** Have gesture recognition technology incorporated in the drone so that I can choose to opt out of being recorded by using certain gestures (e.g., two hands pose as X), and the drone camera can recognize the gesture and the camera will blur my face or figure in the recording (pictures or videos) [46]. **No-fly-zone (Zone):** I enter my addresses (e.g. home) in a no-fly-zone database so that drones controllers will be warned when they fly the drones near these addresses [167]. **Owner registration (Register):** every drone owner must register with the government by providing his or her real name and contact information. Before flying a drone, the owner must mark his/her

Registration Number visibly on the drone. I can see the registration number on a drone and then find out its owner information [83]. **Controller-bystander app (App)**: a mobile app that allows drone owners to provide information about his/her drone such as owner, purpose, drone model and camera/sensor information as well as the current location of the drone. It also lists drones near me and allows me to learn more information about these nearby drones. I can also directly contact drone owners via the app [228]. **LED license (LED)**: a drone will use a visible color blink sequence of its LED lights to serve as its unique “license” and I can use a mobile app to capture the color blink sequence, identify the drone, and look up the information about the drone (e.g., its ownership or purpose) [136].

For controllers, these descriptions were framed from a controller’s standpoint, for example, “people enter their addresses (e.g. home) in a no-fly-zone database so that I will be warned when I fly the drone near these addresses.” Survey one also had many privacy concern questions. Since they are not the focus of this chapter, we do not report them here. Besides, we did not ask about demographics due to the survey length.

7.3.3 Survey Two

Privacy mechanisms. In survey two, we removed two mechanisms from survey one, i.e., deletion request and gesture opt-out, because they were not well supported by both groups of respondents, as well as they have not been implemented and are challenging to implement in practice. We added two new mechanisms: privacy policy (Policy) and automatic face blurring (Blur). After survey one, in June 2016, the NTIA recommends organizational users of drones to have a privacy policy that describe their drone uses and the related data practices [161]. The face blurring mechanism was modeled after a Google Street View privacy feature that has been used to automatically

detect human faces and blur them [93]. For each mechanism, we tried to describe what it does, how it is implemented, and what controllers and bystanders need to do to use it. To make these mechanisms more comparable, we framed them as administrated/suggested by the US Federal Aviation Administration (FAA). Below are the descriptions.

No-fly-zone (Zone) is implemented using a database maintained by the FAA. If a citizen is not comfortable of having drones flying around her house or apartment, she can go to the no-fly-zone website and enter her home address to designate the area within 10ft of her address (including backyard) as a no-fly zone. She needs to submit a document that verifies her residence (e.g., a utility bill). After the no-fly-zone system validates the entered address, the self-designated zone will be stored in the no-fly-zone database.

The drones incorporate the information of this no-fly-zone database either by directly connecting to the database via WiFi or by downloading and updating the database in the drone firmware on a regular basis. These no-fly zones will be highlighted on the map in the drone control interface. In addition, when a drone flies into a no-fly zone indicated by a citizen, the drone operator will get a warning on the drone control interface. Since there are no laws that require drone operators to honor these no-fly-zone requests, the drone operators may or may not choose to honor these requests.

Owner registration (Register): Every drone owner in the US must register with the FAA by providing his or her real name and contact information. Before flying a drone, the owner must mark his or her Registration Number visibly on the drone. In the event that a drone behaves inappropriately, a bystander may report to a law enforcement department. Federal law requires drone operators to show the certificate of registration to any Federal, State, or local law enforcement officer if asked.

Controller-bystander app (App) is designed to improve communication between drone controllers and bystanders. The app works with three assumptions: (1) drones have a GPS module; (2) drones have a Wi-Fi module; and (3) both drone controllers and bystanders have installed and created an account in this app on their mobile devices. The app is operated by the FAA. By default, GPS and Wi-Fi will be turned on while a drone is flying. The drone will record its location information as well as its recording status (e.g., whether the drone is taking photos or videos). This information will first be transmitted from the drone to the controller's app on his or her mobile device through Wi-Fi, and then sent back to a central database on a regular basis.

A drone controller creates an account in the app with information about his or her drone (e.g., drone model, usual flight areas and times) as well as optional contact information. An app user can choose a pseudonymous user name in the app. Registered users of the app can send each other private messages via the app. In addition, the controller can choose to share photos, videos, or live video feed taken by the drone in the app so that other registered app users can see.

When a bystander creates an account and then logs into this app on his or her phone, the app will check with the central database on a regular basis. All the updated information, including drones nearby, will show up in the app interface. For example, if there is a drone nearby, the drone will show up on a radar map with the distance and direction from the bystander's current location. If the bystander would like to message the drone controller, the bystander just needs to tap on the drone in the radar map. The bystander will see all public information about the controller and the drone and can send a private message to the controller through the app.

LED license (LED): A drone has an array of color LED lights (e.g., blue, green, red) that can be seen by more than 300ft without using any special equipment. These LEDs blink in a particular sequence to help people visually identify the drone. In other words, the blink sequence of LEDs

serve as the drone’s “license.” This system is operated by the FAA. A drone controller can sign up to use this system by registering an account via the system’s website and can optionally provide information about himself or herself as well as information about the drone. When a bystander spots a drone nearby, he or she can use the companion LED license mobile app to capture the LED blink sequence (with its camera), identify the drone, and look up the information about the drone (e.g., its ownership or purpose) provided by its owner/controller.

Privacy policy (Policy): The FAA recommends any organization that uses drones to have a drone privacy policy on their website. The privacy policy should include information about how they use drones, such as what kinds of drones they use; where, when and why they fly the drones; what kinds of data the drones will capture (e.g., pictures or videos) and for what purposes; how long the recorded data will be retained; how the recorded data will be processed and/or shared to others; and if citizens have questions about their drone use, how to contact them. This drone privacy policy can either be a standalone privacy policy or part of an organization-wide privacy policy. Ordinary citizens can visit the organization’s website to find and review its drone privacy policy.

Automatic face blurring (Blur): Drones have a built-in feature that can enable automatic identification and blurring of human faces in the pictures/videos taken by the drone camera. By default, this feature is turned on. The FAA recommends drone controllers to use this feature unless there is a legitimate reason not to do so.

We aimed to model these mechanisms realistically. Some mechanisms have been implemented for drones (owner registration, no-fly-zone, and LED license) or used in other domains (privacy policies for websites, and face blurring for Google Street View). The controller-bystander app has been proposed but not implemented [228]. All mechanisms are voluntary except for owner registration, which is required by the FAA. Some mechanism descriptions (e.g., controller-bystander app)

were much longer than others (e.g., owner registration), but that reflects their relative complexities from users' perspective.

Scenarios. Next, we provided respondents three concrete drone usage scenarios, adopted from our prior work [228]. Below are the descriptions.

Neighborhood safety scenario: Your neighborhood recently had several public safety incidents (e.g., burglaries). The local police department hires a few drone controllers to fly multiple drones with cameras in the neighborhood for public safety purposes. As a result, the neighborhood will be continuously monitored. The drones will be streaming a live video feed to the police department but will not record any photos or videos.

Public park scenario: A drone controller is flying his drone in a public park and taking photos and videos for fun. You and your family, together with several other families with kids are playing in the park. You and your family members may be captured in the pictures and videos taken by the drone.

Real estate photography scenario: A real estate agency company hires a drone controller to shoot photos and videos of a house for sale. When the controller flies the drone to take photos and videos of the house, these recordings might capture your houses and/or your backyard.

These scenarios varied by the type of controllers (e.g., companies vs. individuals), the purpose of drone usage (e.g., personal enjoyment vs. public safety), the number of drones used (e.g., single vs. multiple), the duration of drone usage (one-time vs. continuous), and the nature of recording (e.g., streaming without recording vs. recording). We randomized the order of scenarios. For each scenario, we asked respondents which privacy mechanism(s) they want to use and why. We finished with demographic questions such as age and gender.

7.3.4 Data Analysis

We computed descriptive statistics of the quantitative data (e.g., ratings of privacy mechanisms). We also coded the open-ended answers using a thematic analysis, “a method for identifying, analysing, and reporting patterns (themes) within data” [34]. First, we carefully read through the open-ended answers. Second, we independently open coded a subset of open-ended answers. Third, we discussed and created a code book containing codes that cover the respondent’s overall sentiment of the mechanism (e.g., positive), specific pros (e.g., easy, practical, effortless, similar to existing mechanisms) and cons of the mechanisms (e.g., inaccurate, subject to hack, requiring too much effort, useless, impractical, increasing government surveillance), implementation details of the mechanism (e.g., scope of effective operation, communication channel, mobile app), and suggestions to improve the mechanism (e.g., legal requirement, automatic enforcement, restricting access to the controller data). We then used the code book to code the rest of the open-ended data. se themes.

7.4 Results

We now report drone controller and bystander respondents’ quantitative ratings of and qualitative feedback on different privacy mechanisms in both surveys.

7.4.1 Results of Survey One

Figure 7.1 shows the percentages of controller and bystander respondents who were either “positive” or “very positive” that a privacy mechanism is effective, practical, and that they are willingness to use it. For instance, 51% of bystanders thought owner registration is practical, whereas 42% of bystanders thought so for LED license. Therefore, we say that owner registration received more

support than LED license, from bystanders, based on the practicality measure. In general, owner registration and no-fly-zone received more support than the other four mechanisms tested in this survey, from both bystanders and controllers, across all three measures.

Since we removed deletion request and gesture opt-out from survey two, we will focus on people's qualitative feedback on these two mechanisms here. We will discuss the feedback on the other four mechanisms using the data from survey two because it had more detailed mechanism descriptions and a wider range of feedback than survey one. Whenever possible, we report the percentages of bystanders and controllers expressing a main opinion of a mechanism.

Deletion request (Delete). Many bystanders (17%) felt this mechanism can be useful if their requests are honored. Some bystanders also raised two main issues: (1) there is too much work for bystanders (22%), and (2) controllers may ignore/reject the requests (15%). One bystander summarized both points, saying *“This requires too much effort, and there doesn't seem to be any consequences if the drone owner chooses to do nothing.”* Another bystander highlighted his concern about malicious controllers: *“A drone that is trying to spy on me or, otherwise, has ill intentions is not going to cooperate anyway.”* For controllers, some of them (5.9%) felt this mechanism is unnecessary partly because they only publish photos that they deem safe to post. Besides, some controllers (6%) were concerned that bystanders can abuse this mechanism by sending an overwhelming number of requests.

Gesture opt-out (Gesture). Many controllers (29%) and bystanders (10%) thought this can be a good solution if people know it. The burden is on the bystanders to learn the gesture. However, some bystanders (15%) argued that it is the controllers' responsibility to protect bystanders' privacy. One bystander explained, *“I feel like I shouldn't have to make gestures to protect my own privacy and that I would have to constantly be watching out for drones for this to be effective.”* On the

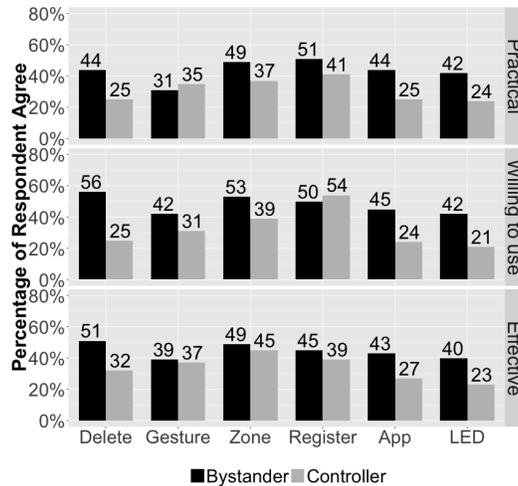


Figure 7.1: Survey one results.

Percentages of bystander and controller respondents who were either “positive” or “very positive” that a privacy mechanism is practical, effective, and that they are willing to use it. The six mechanisms include (left to right): Delete request (Delete), Gesture opt-out (Gesture), No-fly-zone (Zone), Owner registration (Register), Controller-bystander app (App), and LED license (LED).

other hand, some controllers (6%) felt there is really no need for opt-out because drone cameras are usually not good enough to capture people’s faces in the air. One controller explained, *“There is a real lack of knowledge about the cameras on drones. Unless it is a large octo-copter being used by a professional operator with a high priced DSLR camera, then the images/videos you get would be grainy, and if taken from more than about 15ft up unable to identify faces.”* This quote also suggests that an information asymmetry about drones’ capabilities exists between controllers and bystanders.

7.4.2 Results of Survey Two

In survey two, 42% of bystanders were male and 58% were female, whereas 66% of controllers were male and 34% were female. In terms of age, controllers (20% 18-25, 52% 26-35) were slightly younger than bystanders (14% 18-25, 42% 26-35). 70% of controllers had less than one year of experience in drone usage and 30% had more than one year of experience. Most bystanders were not familiar with drones.

Mechanisms	Pros	Cons
1. Deletion request (Delete)	+ Helpful if requests are respected (both)	- Too much work for bystanders (both) - Controllers can ignore or reject requests (both) - Too many requests (controller)
2. Gesture opt-out (Gesture)	+ Good solution if people know about it (both)	- Too much work for bystanders (both) - Have to learn the gesture (both) - No need for opt-out (both)
3. No-fly-zone (Zone)	+ Simple and requires little effort (both) + Add control over controllers (bystander) + Similar to do-not-call list (both)	- No law enforcement (both) - Practical issues due to proximity of homes (both) - Large amount of data (controller)
4. Owner registration (Register)	+ Practical in tracking down controllers (both) + Similar mechanisms in other domains (both) + Discourage irresponsible use (bystander) + This mechanism is already in use (controller)	- Not directly protect privacy (both) - Privacy issue for controllers (controller)
5. Controller-bystander app (App)	+ Enhance controller-bystander communication (both) + Improve controller accountability (both)	- Too much work for bystanders (both) - Privacy issues for controllers (both) - Responses not guaranteed (bystander)
6. LED license (LED)	+ Help identify controllers (both)	- LED patterns can be changed or hacked (both) - Phone camera cannot recognize the pattern (both) - Not directly protect privacy (both) - Too many possible patterns (controller)
7. Privacy policy (Policy)	+ Give bystanders peace of mind (controller) + Provide information about drone use (controller) + Hold organizations more accountable (bystander)	- People rarely read privacy policies (both) - Not directly protect privacy (bystander) - Policy not followed (bystander)
8. Automatic face blurring (Blurring)	+ Effective in hiding people's identity (both) + Make people feel more secure (bystander)	- Conflict with controllers' purpose of use (both) - Slow or inaccurate facial recognition (controller)

Recall that survey two also tested six privacy mechanisms, excluding deletion request and gesture opt-out from survey one, but including the other four mechanisms from survey one as well as two new mechanisms: privacy policy and automatic face blurring. Similar to Figure 7.1 of survey one, Figure 7.2 shows the rating results of the six mechanisms in survey two. In general, owner registration and automatic face blurring received more support than the other four mechanisms tested in this survey, from both bystanders and controllers, across all three measures. Since the controllers and bystanders differed in their age and gender distributions, we controlled for these demographic differences by taking subsets of the original data set and checking the subset results. For instance, we extracted the data of all female respondents of age 26-35, and compared the controllers and bystanders within this subset. The subset results were in line with the results in Figure 7.2.

Next, we present respondents' qualitative feedback on each mechanism. Table 7.1 summarizes the perceived pros and cons of each mechanism in survey one and two. We provide examples of these opinions below.

No-fly-zone (Zone). Both controllers (30%) and bystanders (20%) appreciated this mechanism is simple and requires little effort. One bystander highlighted, *"I think the concept of a no-fly database is simple enough, and practical enough because little is required to get your property included in it."* Many respondents from both groups also associated it with the do-not-call list that they were already familiar with. One controller said, *"I like this system and I think it's a unique idea. This would give bystanders the option of 'opting out' of having drones around their space in much the same way as the 'no call list' works for telemarketers."* In addition, some bystanders thought it will add a layer of control and responsibility over controllers. For instance, one bystander said, *"I think this is effective because it puts the responsibility mostly on the drone operator and*

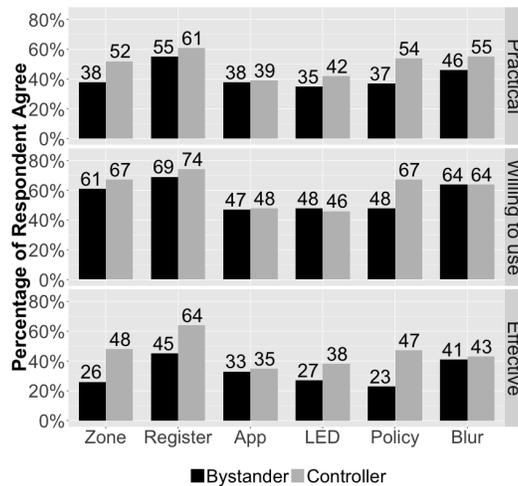


Figure 7.2: Survey two results.

Percentages of bystander and controller respondents who were either “positive” or “very positive” that a privacy mechanism is practical, effective, and that they are willing to use it. The six mechanisms include (left to right): No-fly-zone (Zone), Owner registration (Register), Controller-bystander app (App), LED license (LED), Privacy policy (Policy), and Automatic face blurring (Blur).

allows bystanders to opt in or out.”

However, both controllers (28%) and bystanders (55%) raised concerns about the lack of enforcement because of its voluntary nature. Many respondents from both groups suggested mandating this mechanism by legislation. For instance, a controller said *“I don’t think the ‘no fly zones’ will be respected. There would have to be a law requiring the zones to be respected or it probably won’t work.”* This speaks to the concern that some controllers might choose to ignore this mechanism. Besides, both groups (controller: 14%, bystander: 13%) also raised practical issues due to proximity of addresses. One controller questioned, *“If my neighbor didn’t want a drone flying near their house would that keep me from flying my drone ten feet away above my yard?”* Some controllers (5%) also raised a practical concern about maintaining the large amount of data this mechanism may generate, as one controller noted, *“That would be a massive geographic database, with all the design, operation, and maintenance problems such a thing has.”*

In addition to laws, some bystanders suggested making drones respect these no-fly-zone signals automatically. For instance, a bystander proposed, *“Like, the drone operator gets a warning that they are within so many feet of a no-fly zone, and warnings up until they reach it, then the drone be deactivated if they ignore the warnings and enter the zone.”* While completely automatic deactivation of drones might be unsafe, configuring the drones not to enter a no-fly zone is doable similar to how some drones are configured to stay away from sensitive places such as airports via geo-fencing [69].

Owner registration (Register). Many bystanders (43%) praised that this mechanism can help make controllers more accountable for their drone practices. Some even suggested that people need to take lessons and get a license before they can operate drones. For example, one bystander suggested, *“This will help to hold flyers accountable for their actions while flying a drone and could be extended to require lessons and certification in the actual flight of the drone just like a driver’s license.”* This mechanism was also positively received by the controllers (42%). In fact, many of them self-reported having done the registration, which is required in the US by the FAA.

However, many bystanders (39%) and controllers (28%) felt this mechanism does little to directly protect people’s privacy. One bystander expressed, *“It seems like a good basic requirement, but would not necessarily protect people much.”* Another controller believed it is more for safety than privacy, saying *“owner registration is a good idea but it will not have any effect on ‘privacy.’ It will be more useful in identifying the owner in case of an accident with the drone.”* In addition, some controllers (5%) were concerned about who can access their registration information and explicitly mentioned that only the government can access that information. Furthermore, some controllers worried that this mechanism can increase the government’s ability to track their activities. One controller summarized the pros and cons, saying *“I think it’s a good and a bad thing. Good*

in that if someone is using their drone for illegal activity it would be easy to identify their drone information if they are reported. It's a bad thing because it's another way for the government to monitor people's activities."

Controller-bystander app (App). Both controllers (19%) and bystanders (21%) commended that this app can enable or enhance the communication between bystanders and controllers. For example, one controller said, *"Controller-bystander app is a very effective way of using Drone. It provides a direct way of communication between drone controllers and bystanders."* Some respondents (controller: 3%, bystander: 10%) also felt it can increase the accountability of controllers. For example, one bystander expressed, *"I think the app will provide better protection to bystanders and make the controller more accountable."* Allowing bystanders to see nearby drones and information about their usage would hold the associated controllers responsible for their behaviors.

However, some controllers (12%) and bystanders (2%) raised a potential privacy violation of controllers since their drone practices are tracked. One controller complained, *"I feel that this is a huge invasion of privacy for the drone owner him/herself. It seems that it will record all activities and where the drone is and where it has been and if it was used for pictures/video. This is worse than someone accidentally having their face recorded."* This highlights the challenge of making drone usage transparent while protecting controllers' privacy.

In addition, many bystanders (27%) complained that this mechanism demands too much effort. One bystander commented, *"This requires a lot of work for the bystander. Some people will not know about this app and the fact that they can use it."* Even if they are aware of the app, they still need to install, learn how to use, and actually use the app. Another bystander felt this voluntary mechanism would fail to detain malicious controllers, saying *"This seems like an honor-system thing and I don't think that would solve much with people who are using drones inappropriately."*

They've already proven they won't follow an honor system." This highlights the concern that some controllers may intend to bypass this mechanism. To improve this mechanism, many bystanders proposed that controllers should be required to use it by regulations. For instance, one bystander suggested, *"I think maybe it would have to be mandatory to install and use this app to fly a drone or the operator could face federal charges. Maybe a live feed of what the drone is recording could be useful to bystanders."*

LED license (LED). Many controllers (22%) and bystanders (16%) felt this mechanism can help identify drones and their controllers, as one bystanders noted, *"I think it would help in identifying the drones owner."* However, both groups (controller: 38%, bystander: 50%) also raised practical issues, such as the LED lights can be obscured or altered by the controllers. For instance, one bystander said, *"there are some less honest people out there would be obscure the lights to prevent detection."* This underlies the concern that some controllers may intend to circumvent this mechanism. Another bystander further suggested mandating this mechanism, *"That seems kind of silly, because people who are using drones maliciously will simply not sign up to register their drone. It needs to be made mandatory somehow upon purchase of a drone/built into all new drones."* In addition, a few controllers (8%) and bystanders (11%) suspected that cameras on phones are not good enough to capture the blinking sequence correctly. For instance, one controller said, *"it would be hard for a camera to pick up blinks with a phone camera."*

Some bystanders were also concerned about the effort needed including learning about, finding and downloading and then using the app. One bystander summarized, *"It's not practical to the every-day bystander. It's too much work for the average person to go through and they shouldn't have to go through such lengths to ensure their right to privacy."* In addition, some controllers complained that this mechanism can violate their privacy because people can see their information

via the app. One controller said, *“I wouldn’t want just any bystander with an app to have the ability to look my info up.”*

Privacy policy (Policy). Many controllers (16%) noted a privacy policy can provide bystanders information about drone practices. One controller said, *“I think it is a decent policy. It would be easy to implement and would be good for bystanders who want to know what you’re doing with the drone.”* Besides, some controllers (12%) thought it can provide bystanders peace of mind, as one controller explained, *“I think it gives people more peace of mind about drones knowing they can request information on why they’re being used.”* However, some respondents from both groups (controller: 4%, bystander: 21%) felt it does not directly protect privacy, as one bystander put it, *“it doesn’t protect people of prevent anything.”* Another issue was that people rarely read privacy policies (controller: 14%, bystander: 16%). One controller said, *“I think this is a necessary feature, although I’m not sure how effective it will be. Most people do not pay attention to privacy policies in general.”* This suggests that they felt this mechanism is needed but not sufficient by itself.

Bystanders generally appreciated this mechanism. Some bystanders (8%) also felt it will help hold controllers accountable. One bystander said, *“This could help with accountability and discourage inappropriate behavior.”* However, many bystanders (32%) also questioned whether organizations will follow their policies. One bystander was pessimistic, saying *“It’s highly debatable how many organizations actually even follow their own privacy policies. This would do ZERO, literally ZERO to help curb privacy violations and privacy concerns.”* This highlights the need for enforcement. In the US, the Federal Trade Commission can prosecute companies that do not follow their own privacy policies as deceptive practices.

Automatic face blurring (Blur). Many controllers (22%) and bystanders (18%) valued this mechanism’s potential in hiding people’s identities. One controller commented, *“Auto blur would*

absolutely protect privacy.” One bystander said, “Seems practical enough because it’s turned on by default. I would feel more safe should this feature be implemented.”

However, both groups also had reservations about this mechanism. Some controllers (27%) and bystanders (48%) were concerned that this mechanism can be useless because controllers can easily turn it off. For instance, one bystander said, *“If you can disable the setting, it is worthless. People all like to spy and see things so they won’t care about privacy if they can disable the setting.”* Another issue was that bystanders do not have an easy way to know whether this feature is on or off. Even if it is on, some controllers and bystanders suspected that it can be reversed. One controller commented, *“I’m sure any half way decent hacker can un-blur this picture.”* This comment highlights the concern that some controllers may have the ability to circumvent the mechanism. Some controllers complained that this mechanism is on by default. For instance, one controller said *“It sounds stupid. And what if I’m trying to identify someone? I don’t want anything blurred.”* A few controllers (8%) also questioned the capability of this mechanism. For example, one controller said, *“I just don’t think the facial recognition software can work fast enough to block out all faces as soon as they appear.”* While this mechanism might not be able to blur faces during the recording, it has been shown to work well on recorded images/videos [93].

Drone Usage Scenarios

We next asked respondents to select the mechanism(s) they want to use under three concrete scenarios and explain why. They all chose their preferred individual mechanisms and many also suggested using multiple mechanisms together. Their choices of mechanisms varied across scenarios, showing their context-dependent preferences.

Neighborhood safety scenario. In this scenario, the largest percentages of bystanders chose the following three mechanisms: privacy policy (48%), automatic face blurring (36%), and no-fly-zone (34%). Controllers instead chose: drone owner registration (49%), privacy policy (41%), and automatic face blurring (39%). Many respondents desired both privacy policy and face blurring. They felt that a privacy policy provides information and serves as a notice and face blurring protects their identities. For instance, one bystander explained, *“considering the drone privacy policy, I would like to know how and to what extent the police will be using this footage. Since they will be on constant patrol, I would like to have all faces blurred to protect anonymity and privacy.”*

Public park scenario. In this scenario, bystanders preferred face blurring (82%), controller-bystander app (31%), and drone owner registration (31%). Controllers preferred face blurring (71%), drone owner registration (39%), and privacy policy (29%). Many bystanders and controllers considered face blurring the most effective mechanism partly because its protection for children. One bystander explained, *“The face blurring thing is the best option to protect their children and the families at the park since there is nothing else that can be done about it.”* Some bystanders liked the combination of owner registration and controller-bystander app. For instance, one bystander explained, *“It would be helpful to know the drone is registered with the FAA and the controller-bystander app would be perfect in this case. It would make the bystander feel safer and may even help to make friends.”* This comment also suggests that the app can help people socialize. Another bystander added, *“I believe in asking for something. ‘Please do not record myself or my family, thank you.’ would send a polite and clear message.”*

Real estate scenario. In this scenario, the three most chosen mechanisms by bystanders were: no-fly-zone (64%), face blurring (61%), and privacy policy (38%). For controllers, it was the same set of mechanisms but in a different order: face blurring (51%), privacy policy (49%), and

no-fly-zone (45%). Bystanders felt no-fly-zone can at least signal their intent to opt out. 45% of controllers indicated they would respect no-fly zones. One controller said, *“It would show which houses to avoid, as in, shoot from a different angle if a neighbor is on the list.”* Both groups also valued the face blurring mechanism as it can protect people’s identities. Since this scenario was related to organizational uses of drones, several bystanders and controllers thought that a privacy policy would be helpful. One controller also suggested combining multiple mechanisms for better privacy protection, *“Given the purpose and who is controlling it, I think the privacy policy would be effective, but added protection of face blurring and, if I was so inclined, respecting my no-fly zone would be beneficial.”*

7.5 Discussion

We discuss respondents’ perceptions and relative preferences of different mechanisms as well as important questions for how to address privacy challenges of drones.

7.5.1 Perceptions and Preferences of Privacy Mechanisms

While the privacy mechanisms that we explored are not exhaustive, they cover a wide range of designs ranging from technical mechanisms (e.g., LED license and face blurring) to policy mechanisms (e.g., own registration and privacy policy). These privacy mechanisms can be roughly categorized into two groups based on our respondents’ ratings of and feedback on each mechanism.

While no mechanism was perceived as a silver bullet, owner registration and face blurring gained relatively more support from both bystanders and controllers than other mechanisms. This matters because this result suggests these two mechanisms are more likely to be adopted and to help mitigate bystanders’ privacy concerns. In other words, they have great potential to succeed in practice. Owner registration was already in use and was perceived by both controllers and

bystanders as useful but insufficient by itself. Face blurring was perceived by both groups as useful and something that requires little effort. It has not been applied for drones but should be considered by drone manufacturers as a useful privacy feature.

Privacy policy and no-fly-zone also received some support, albeit more controllers perceived them to be practical and effective than bystanders. This suggests that while controllers may adopt these two mechanisms, bystanders may consider them ineffective. The remaining four mechanisms received even less support, but this does not mean they are completely useless. For instance, in the public park scenario, the second most selected mechanism by bystanders was the controller-bystander app because it allows them to directly communicate with controllers about their privacy concerns about the drone. Our prior research shows that bystanders are concerned that drone controllers can be invisible or cannot be reached for communication [228]. The FAA has promulgated new drone safety rules, such as prohibiting flight over people and night operations, and requiring drones to be in visual line of sight of the drone controllers [84]. These new rules do not require drones nor drone controllers to be visible to bystanders. Therefore, bystanders' concern about invisible controllers remains largely unaddressed. The controller-bystander app can allow bystanders to contact controllers, but our controller respondents did not value this mechanism as much partly because its usage might infringe on their own privacy.

7.5.2 Important Questions for Addressing Drone Privacy Issues

We now discuss important questions and suggestions for designing future privacy mechanisms and policies for drones.

Improving individual privacy mechanisms

How to improve individual privacy mechanisms? We suggest to consider three aspects: effort, practicality, and effectiveness.

Effort. One important question is how much effort a mechanism demands from a bystander or controller. If people think a mechanism requires too much effort, then they are unlikely to use it because privacy is often not their main or immediate goal. Deletion request, gesture opt-out, and controller-bystander app were not rated higher partly because they were considered as requiring too much effort from bystanders. In addition, many bystanders believed that it should be the controllers' responsibilities to protect the bystanders' privacy. However, this can be a risky belief because controllers may care more about protecting their own privacy rather than the bystanders' privacy. One reason that face blurring was highly rated is because it requires minimum effort from controllers and bystanders.

Practicality. Another question is how practical a mechanism is in reality. Many common privacy strategies are challenging to implement in the context of drones. For instance, it is hard to implement user consent when a drone is operating in a public space (e.g., a park) where there are many people present. Do we require the drone controller to get consent from each person before flying the drone or using the drone to take pictures/videos? What if bystanders have conflicting preferences? Another example is providing privacy notices. When drones are flying in the air, it would be difficult for people to see or read any privacy notice on the drones. How to help bystanders identify drones' privacy policies or notices, and understand what privacy mechanisms have been applied is important for future privacy designs and policies for drones.

Effectiveness. The third question is how effective a mechanism is in practice. This is particularly important in the context of drones because most of the existing privacy mechanisms for drones are voluntary. Both controllers and bystanders believed some controllers have the *ability* and/or *intention* to circumvent these mechanism. For instance, un-blurring face blurred images highlights not only the potential technical weakness of the face blurring mechanism but also controllers' *ability* to reverse it. In contrast, malicious controllers who spy people would intentionally ignore no-fly-zone requests, speaking more about controllers' *intention* to avoid the mechanism.

Our respondents suggested using laws and/or technical means to enforce these voluntary mechanisms. For instance, some controllers suggested “hard coding” no-fly-zone information into drones that automatically prevent them from flying into a no-fly-zone. This is known as geo-fencing, which currently works for sensitive locations such as airports and does not include people's homes. Other respondents suggested making laws to mandate and enforce mechanisms such as no-fly-zone, privacy policies, and face blurring.

Combining multiple mechanisms

Our scenario-based results suggest that respondents from both groups had desires of using a combination of mechanisms. For instance, privacy policy and owner registration were often considered helpful but not sufficient because they do not directly protect people's privacy as many respondents put it. Therefore, our respondents suggested combining multiple mechanisms such as privacy policy, owner registration, and face blurring since they can improve different aspects of privacy. For instance, privacy policy can provide notice about drone usage, owner registration can help hold controllers accountable, and face blurring can hide bystanders' identities. Our respondents' choices of mechanisms also varied across different scenarios, suggesting that they had context-based

preferences of privacy mechanisms. Future research can explore packages of mechanisms based on the changing scenario or context.

Bridging the bystander-controller mismatch

Our results also indicate that our bystander and controller respondents often had different perceptions of the same privacy mechanisms (e.g., the effectiveness of privacy policy). In the case of deletion requests, some controllers were concerned that bystanders may abuse this mechanism by sending them an overwhelming number of requests. These differences between controllers and bystanders are perhaps not surprising because of their roles. Their behaviors can be thought as the in-group (controllers) versus out-group (bystander) behaviors in an inter-group process (drone operations) [37]. In drone operations, controllers directly operate drones and presumably focus on utilizing and enjoying drones, whereas bystanders do not directly participate in drone operations and thus prioritize their welfare such as safety and privacy against drones.

One way to bridge the bystander-controller mismatch is to improve the trust between them. Prior research has also shown that lack of trust is an antecedent to privacy concerns [192]. When controllers are organizations, we can learn from the e-commerce literature, which has shown that companies can build consumer trust and thus reduce consumer privacy concerns by using a number of measures such as adopting fair information practices (e.g., notice and consent) [61], presenting privacy policies [76], and displaying privacy notices or seals [209]. We studied some of these ideas, for instance, privacy policy and gesture opt-out (a form of user consent).

Prior literature has also proposed different ways of providing notice and consent to users in ubiquitous computing environments in order to improve users' privacy awareness [128, 123, 183]. Future work can explore these ideas (e.g., broadcasting a user's privacy preferences [123] in a

physical location) for drones. Displaying privacy notices or seals directly on a drone might be hard for people to see or read, but they could be shown on the information page of the drone once people have identified a drone by LED license or controller-bystander app, for instance. When the controllers are individual users, we can learn from ways to increase interpersonal trust such as providing transparency in decision-making (e.g., why use drones to take pictures) and holding people accountable [7]. Many respondents commended that the controller-bystander app and owner registration help hold controllers accountable.

Protecting the privacy of both bystanders and controllers

While bystanders valued their privacy, controllers were also concerned about protecting their own privacy. For instance, when considering owner registration and controller-bystander app, many controllers did not want bystanders (in theory, almost anyone can be a bystander) to know their information. Some controllers also expressed concerns that these mechanisms could increase the government's abilities to track them. Therefore, another important privacy design question for drones is - how to balance the privacy of bystanders and controllers. For instance, one idea to help protect controllers' privacy against bystanders is that bystanders can only report problematic drones to the government using the controller's registered ID but cannot access other controller information. Alternatively, bystanders can only view a controller's information when they are physically close to the flying drone.

Lastly, privacy has been a key research theme in the HCI community. Our research highlights that the design of human-drone interaction should not only consider controller-drone interaction but also the indirect involvements of bystanders, as their privacy can be intentionally or inadvertently violated by drone operations. Identifying privacy mechanisms that are supported by both controllers

and bystanders is thus important to inform the development of public policies and future designs of drone technologies.

7.5.3 Study Limitations

First, we cannot completely guarantee that all controller respondents were actually drone controllers. However, we double checked with the open-ended question on what brand/model of drones they have and they had reasonable answers.

Second, our sample cannot generalize to all drone controllers and bystanders. We recruited respondents from Amazon Mechanical Turk and multiple drone forums. We also focused on the US. Thus, our results may not apply to other countries.

Third, the privacy mechanisms studied in our research are by no means exhaustive, but we chose a diverse set of technology-based and policy-based mechanisms. While we attempted to provide detailed and realistic descriptions of these mechanisms, some descriptions are hypothetical because the described mechanisms have not been fully implemented in practice and we had to imagine their implementations. Besides, the drone usage scenarios are hypothetical, but they were modeled largely after real-world uses of drones.

Lastly, our study focused on people's perceptions of privacy mechanisms rather than their actual adoption behavior. We only collected self-reported data, which can divert from actual behavior, as shown in the privacy paradox literature (e.g., [194]). However, we note that people's perceptions or behavioral intentions (e.g., willingness to use a mechanism) is important to study because they can influence people's real behavior.

7.6 Conclusion

As drones continue to be adopted and used by governments, organizations, and ordinary consumers, how to protect people's privacy against drones is a critical and timely question. We conducted two surveys to investigate how drone controllers and bystanders perceive a diverse set of privacy mechanisms for drones. Our respondents raised various pros and cons of each mechanism. While owner registration and face blurring received most support individually by both groups, many respondents also suggested using a combination of mechanisms, which varied across different drone usage scenarios. We highlight a number of important questions for future privacy designs and policies of drones.

7.7 Acknowledgements

We thank Xingzhi Guo and anonymous reviewers for their helpful feedback. This research is in part supported by an internal grant from Syracuse University.

Part IV
Case 3: Smart Home

Chapter 8

Smart Home Privacy: Users' Perspectives

Home is a person's castle, a private and protected space. Internet-connected devices such as locks, cameras, and speakers might make a home "smarter" but also raise privacy issues because these devices may constantly and inconspicuously collect, infer or even share information about people in the home. To explore user-centered privacy designs for smart homes, we conducted a co-design study in which we worked closely with diverse groups of participants in creating new designs. This study helps fill the gap in the literature between studying users' privacy concerns and designing privacy tools only by experts. Our participants' privacy designs often relied on simple strategies, such as data localization, disconnection from the Internet, and a private mode. From these designs, we identified six key design factors: data transparency and control, security, safety, usability and user experience, system intelligence, and system modality. We discuss how these factors can guide design for smart home privacy.

8.1 Introduction

A smart home consists of different sensors, systems, and devices, which can be remotely controlled, accessed and monitored [41, 115]. The massive amount of data collected by Internet of Things (IoT) devices in a smart home allows entities to infer sensitive information without actually collecting them [56, 152]. Even seemingly innocuous data, such as home temperature and air conditioner

status, could be used to determine whether a house is empty or not [137, 197]. In addition, people have expressed privacy concerns about smart homes, such as continuous data collection, sharing, and even misuse [235, 40, 223]. Privacy has thus been identified as a road blocker in the wide adoption of smart homes [110, 125].

To mitigate these concerns, different privacy mechanisms have been proposed, e.g., introducing noise to shape the smart home network traffic to prevent data inference [17]. However, little is known about what kinds of smart home privacy controls people desire. This is an important question to answer because privacy designs that address these desires are likely to be adopted.

To answer this question, we adopted a co-design approach to empower end users and engage them directly in the design process. Co-design [187] is a collaborative design approach in which stakeholders—such as researchers, designers, and users or potential users who are considered as “experts of their experiences” [206]—share their perspectives and cooperate creatively to generate new designs [195]. Kraemer and Ivan advocated that privacy issues in the smart home context should be approached by considering different stakeholders [124]. In our work, we collaborated



Figure 8.1: A photo was taken during one study session.

closely with many groups of participants with diverse backgrounds in designing privacy mechanisms through a series of co-design sessions.

Our main contribution is that we identified six key design factors from our participants' designs of privacy mechanisms for smart homes. These factors include data transparency and control, security, safety, usability and user experience, system intelligence, and system modality. They reflected our participants' expectations in privacy mechanisms for smart homes and can be used as a good starting point to think about the design space of smart home privacy mechanisms.

8.2 Related Work

8.2.1 Smart Home Privacy Concerns and Risks

Prior literature identified a number of privacy and security risks of smart homes. Arbo et al. pointed out the possibility of identity theft and device reconfiguration, suggesting the need for effective malware management [19]. With an experiment, Apthorpe et al. demonstrated how to infer sensitive user interaction with smart home devices through network traffic analyses with reasonable accuracy [16, 17]. A risk analysis of a smart home automation system by Jacobsson et al. pointed out that human-related risks (e.g. poor password selection) and software component risks (e.g., unauthorized modification of functions in the app) were the riskiest ones [109].

End users' privacy concerns have also been examined. By understanding people's mental model of how smart homes work, Zimmerman et al. uncovered participants' privacy concerns about hacker attacks and data abuse [240]. Brush et al. identified four barriers that defer the broad adoption of smart homes, such as "difficulty achieving security" in smart door locks and cameras [40]. Zeng et al. identified a number of concerns people have, such as continuous video recording, data collection and mining, network attacks on local networks, and account hacking [235]. However, people tend

to outweigh cost and interoperability over privacy and security [235]. Worthy et al. found that the fewer trust participants had towards the entities who used their information, the greater control over information collection participants desired [223]. Zheng et al.'s study, on the contrary, found that their participants assumed their privacy is well protected because they trust their smart home device manufacturers [237].

Other studies focused on specific smart home devices. Malkin et al.'s survey about smart TVs revealed their respondents' uncertainty of data collection and usage as well as the common nonacceptance of data being re-purposed or shared with third parties [143]. McReynolds et al.'s study on smart toys unveiled parents' concerns about the toys' recording and data sharing abilities and children's concerns about being heard by their parents [149]. Lau et al.'s study about smart speakers found that users' rationales behind a lack of privacy concerns could lead them to serious privacy risks [129].

8.2.2 Smart Home Privacy Mechanisms

Researchers have proposed various solutions to mitigate privacy concerns and risks in smart homes. For instance, Apthorpe et al.'s solution decreased the inference of sensitive user activities by introducing a minimum amount of noise data to shape the smart home network traffic [17]. Datta et al. developed a Python library for IoT developers to easily implement privacy-preserving traffic shaping [64]. By injecting synthetic network packets, Yoshigoe et al.'s solution obscured the real network traffic and reduced potential privacy vulnerabilities [234]. Wang et al. built a live video analytic tool for denaturing video streams by blurring faces according to user-defined rules [208].

To reduce improper access to users' data, Moncrieff et al. developed a tool to dynamically manage access privileges based on a number of contextual factors in smart home surveillance



Figure 8.2: The flow of our co-design study, including its various components.

(e.g., occupants’ location and content of ongoing conversation) [156]. Arbo et al. proposed a framework to ensure data security for smart home devices by providing dynamically generated policies and interfaces in which end users could use to set up their privacy zones [19]. Chakravorty et al. designed a system to collect and store users’ data, then only allow users to access their data upon successful re-identification [47].

To increase transparency and user control, Das et al. proposed an infrastructure for IoT devices and sensors to send personalized privacy notice and choice based on individual users’ preferences [63]. McReynolds et al.’s study on smart toys suggested that toys should effectively communicate with both parents and children that toys could record [149].

More broadly, to ensure an overall safe environment of smart homes, Lin et al. suggested that auto-configuration support should be developed for smart home network so that whenever a new device is plugged into the network, the supporting system could auto-configure itself and find the most secure settings for the new devices, such as security protocols and essential firmware updates [137].

The commonality of the above mechanisms is that they were proposed or developed solely by experts or researchers. Our work focuses on end users’ needs and perspectives, helping fill the gap in the smart home literature between studying users’ privacy concerns and designing privacy tools only by experts.

8.3 Method

To explore how people desire to protect their privacy in the context of smart homes, we conducted a set of co-design sessions with a total of 25 participants. Figure 8.2 shows our study flow including participant recruitment and two co-design sessions. Each session took about 1.5 hours and each participant was paid \$15 for each session they participated in. Our study was approved by our university IRB.

8.3.1 Participants

Recruitment. We recruited our participants primarily through word-of-mouth, Craigslist, local community centers, libraries, and senior citizen centers. We framed our study as “a design study for smart home technology” and did not mention the word “privacy” to avoid any potential bias. We designed a pre-screening survey to get participants’ demographic information and their experiences with smart home devices.

Participant data. The ages of our 25 participants ranged from 22 to 76 (mean: 41). 13 participants were cisgender female and the other 12 were cisgender male. They had various occupations, such as university staff, librarians, students, software engineers, retired workers, a security guard, a researcher, and a plumber. They were categorized into three types of participants based on their levels of experiences with smart homes: 12 participants owned smart home devices (*users*), 7 participants were interested in buying smart home devices (*interested users*), and 6 participants did not use or plan to buy smart home devices (*non-users*).

Study groups. Participants were divided into five groups (Group A, B, C, D, and E) primarily based on their schedules and levels of experiences with smart home devices. Each group had four to six participants (A:6, B:4, C:5, D:4, E:6). Each session had at least four participants except that

Session 2 of Group D only had two participants due to schedule conflicts. Group E consisted of participants from a senior citizen center. Due to their mobility needs, we chose to conduct the study in their center with the participants from that center only. All other sessions were conducted in our lab.

All participants were invited to both sessions, however not everyone could attend both sessions due to practical constraints (e.g., conflicting schedules). To mitigate this issue, similar to [148], we started each Session 2 with a 15-minute recap of the discussion from the corresponding Session 1 (e.g., pros and cons and privacy concerns of smart home devices) to bring all participants to the same page. In the end, ten participants completed both sessions. Eleven participants only did Session 1 and four participants only did Session 2.

8.3.2 Session 1

The goal of the first session was to understand participants' privacy and security concerns of smart homes and to conduct the initial brainstorming and design. Each session started with a round-table introduction. We then asked each participant to talk about their experiences with smart home technologies and general perceptions. We then provided a working definition of a smart home, "a home that has different sensors, systems, and devices, which can be remotely controlled, accessed and monitored" based on the literature [41, 115]. We showed and explained pictures of a few smart home devices (e.g., voice assistants, smart thermostats, security cameras, and smart toys) to illustrate this smart home definition and potential uses of these devices [151].

Next, we asked our first group-based discussion question, "what are the pros and cons of these devices in your opinion?" This question was meant to frame the discussion in a neutral/balanced manner. In our pilot study, we found that our participants were overly excited about smart home,

tended to fixate on the pros, and hardly considered the cons. To encourage participants to think more about the risks, we added a follow-up question in the actual study that asked, “have you experienced or heard of any negative incidents of smart home technologies, and what can potentially go wrong with smart home technologies?” This question did not prime our participants to only consider the negative aspects because we asked the pros first and they mentioned many pros. The careful consideration of both benefits and risks helped them to consider the trade-off in later co-design activities.

The next two activities were scenario-based because smart home devices can be used in various scenarios or for different purposes and privacy is highly contextual [164, 165]. As such, we hoped to provide opportunities for our participants to explore nuances of smart homes and their contextualized privacy implications. The first activity was a role play. We presented three scenarios adapted from the literature: (1) an Amazon Echo records a conversation between a couple and sends it to other people [201], (2) a security camera monitors a senior citizen’s well-being at home in case of emergencies [238], and (3) a smart toy records and processes a child’s conversation with it in order to respond, but also allows the parents to hear the conversation via a mobile app [149]. These scenarios were chosen to represent different devices, social relationships and power dynamics in the home (e.g., a couple and a co-worker, an older adult and an adult child, and young children and parents). In each scenario, we designed two to four roles for our participants to choose from. Once each participant picked a role, they discussed the potential privacy issues from the standpoint of the role.

The next activity was to co-create smart home usage scenarios. We encouraged our participants to work in groups of two or three. Each group chose a specific smart home device and co-created a usage scenario of that device with one or two researchers. We used six questions to guide the

scenario creation process, i.e., what the device is, where the device is used, who uses the device, when to use the device, why uses the device, and how the device is used. Participants then presented the scenario and discuss any potential privacy implications in that scenario.

Through the above two activities, our participants discussed a wide range of usage scenarios and privacy issues of smart homes. We then moved on to the co-design activity. Specifically, we asked our participants to brainstorm their desired ways to mitigate these privacy issues and draw their design ideas. We provided a number of creation tools (e.g., colored papers, post-it notes, color pens). A student designer was also on site to provide help with sketching if needed. We deliberately asked our participants to work individually, think outside of the box, and consider different kinds of potential solutions. We also explained that the solutions could be futuristic and speculative without considering the status quo. Each participant then presented their ideas to the group.

8.3.3 Session 2

The goal of Session 2 was to continue the co-creation of privacy mechanisms, moving from ideation to creation of prototypes. We started the session by recapping the discussion from Session 1. To help our participants understand different forms of prototypes, we show various examples (e.g., paper prototypes of smartphone screens and a website). Similar to Session 1, we provided different creative tools and had a student designer on site to help them draw.

Each participant had about an hour to work on their prototypes. We encouraged them to discuss their ideas with other participants and the researchers, and then to create the design individually. Then each participant presented and discussed their prototypes with the group.

8.3.4 Data Analysis

Transcriptions and notes. All sessions were audio-recorded upon participants' permissions. The

recordings were transcribed and then analyzed using a thematic analysis [32] by three co-authors. First, we immersed ourselves in the data by reading through the transcripts multiple times. Then we coded one transcription together at the sentence level to develop an initial codebook. Second, we independently coded the same transcription of another session using the codebook. We added new codes to the codebook in that process. Once finished, we compared and discussed our coding, and converged on an updated codebook. The inter-coder reliability was 0.91 (Cohen's Kappa), which is considered good [86]. Next, we coded the rest of the data using the updated codebook, which contains more than 100 unique codes, such as "self-driving car risks," "voice assistant authentication," "home context," "block data collection," "sharing decision," and "intrusion detection." Once finished coding, we grouped all codes into several themes, such as "data transparency and control," "security," "safety," "usability and user experience," "system intelligence," and "system modality." We examined and ensured the codes were assigned to the correct theme.

Image data. We collected participant drawings of their design ideas. Following Poole et al.'s methodology [174], three co-authors coded all elements in every design, including all components involved (e.g., stakeholders, devices, users), information flow, context, as well as other visual elements (e.g., icons, symbols, colors). Over 80 codes emerged from the analysis, and all the codes were grouped into the aforementioned themes in the analysis of audio transcriptions and notes. The inter-coder reliability was 0.84 (Cohen's Kappa).

8.4 Results

Our work contributes to new understandings of how people conceptualize privacy control mechanisms for smart homes. Our participants started with creating their own smart home device usage scenarios. They created a wide variety of usage scenarios covering different devices and

purposes (e.g., smart security cameras for home safety, smart doorbells and locks for remotely locking/unlocking doors, a smart fridge to automate food refill and alert food expiration, and a smart robot to support indoor navigation for people with visual impairments). This activity allowed our participants to explore and discuss possible ways of using smart home devices and potential privacy implications. These scenarios also served as a basis for the subsequent co-design of smart home privacy controls. Next, we will turn to our participants' smart home privacy designs, focusing on the major factors considered in their designs and identified via our thematic analysis. Table 9.2 shows an overview of these factors. It is worth noting that these factors are not mutually exclusive. One design might consider multiple factors. We present these factors below.

8.4.1 Data Transparency and Control

A major privacy concern shared by our participants was smart home devices collecting data about them. They created various designs to increase the transparency of data collections and user control over their data. Seventeen participants considered this factor (P2-6, 9-10, 12, 15-21, 24-25).

Transparency and user awareness. Seven participants' designs (P2-3, 6, 12, 15, 18-19) were centered around improving transparency of data collection and usage of smart home devices. For example, P15 designed a transparency feature for a self-driving car. She considered the car part of the smart home because the car is often parked/charged at home and she can control the car (e.g., start the engine) remotely using voice assistants (e.g., Google Home). However, she was concerned that the car manufacturer might collect her car usage data (e.g., when she used the car, where she had been to) and then use that data to predict her future activities. To address these concerns, she designed the car with two modes, an invisible mode and a visible mode. When she wishes not to be tracked, she can turn on the *invisible* mode (e.g., by plugging in a dedicated USB drive to the car)

Table 8.1: The six factors identified from our participants’ designs of smart home privacy controls

Factors	Examples
Data Transparency & Control	<ul style="list-style-type: none"> - Transparency and user awareness - Data localization - Disconnection from the Internet - Other user controls of data
Security	<ul style="list-style-type: none"> - Authentication of multiple users - Access control - Network intrusion detection
Safety	<ul style="list-style-type: none"> - Notification of physical break-in
Usability & UX	<ul style="list-style-type: none"> - Considerations of user characteristics - Considerations of user effort
System Intelligence	<ul style="list-style-type: none"> - Context detection - Personalization
System Modality	<ul style="list-style-type: none"> - Hardware devices - Apps, modes, policies

to hide her activities. In contrast, under the *visible* mode (default mode), her driving data can be tracked but she can use an app interface to see what data about her has been collected.

P15 explained, “*so the visible basically tells you what you have done with this car, like a transparency tool... [the tool] can also make sense of how the manufacturer uses the data. Like they can infer whether I’m a night person or not to increase my insurance payment.*” (P15) Her design of the visible mode provides more transparency about the car’s data collection and usage practices. However, she suggested this feature should be provided by third-party companies because she felt the car manufacturers might not tell the truth.

It is also worth noting that, our participants considered the purposes of data collection in their smart home scenarios when designing the privacy controls. For instance, even in the invisible mode of P15’s design, she would share when/where she uses the car with a trusted third party (e.g., for better navigation purpose) but would not with the car manufacturer. In P3’s design, the security

camera monitors his home for safety (i.e., purpose), but the associated app does not show the actual images or videos and only offers text descriptions thereof to mitigate privacy risks.

Data localization. A common element across seven participants' designs (P2, 4-5, 15-16, 20, 24) was data localization, the idea that smart home devices store and process the collected data locally as opposed to sending the data to a remote server. For example, P16 designed a smart door lock with a fingerprint reader to improve the safety of her home. Since the fingerprint reader collects her biometric information, she designed an additional privacy feature to protect her fingerprint data by only storing it locally in the lock. She explained, "*the fingerprint will be stored onsite only. There is no need to have it connected to anything. If you didn't have your original key and you didn't get in with your fingerprint for some reason, the only thing that the company can do is complete pledge it and you would start it as a brand-new device because they would not have access to get into that.*" (P16) According to P16's design, her fingerprint data will reside in the lock, but the smart lock is still network-connected because the company could remotely reset the lock if the user lost her key or the fingerprint reader stopped working. However, we note that remotely resetting locks can pose security risks.

Disconnection from the Internet. This privacy mechanism means disconnecting smart home devices from the Internet, essentially working in an offline manner. Five participants' designs (P2-4, 24-25) included this idea. For instance, P2 was concerned about home security cameras collecting personal data and storing these data in cloud storage, which may not be secure. To address these concerns, he proposed a design of a physical lock that could be plugged into a security camera to protect his personal data. He explained, "*I think what people need is something like a lock that can be plugged into the security camera to lock our data like gender or activities. Now they [security*

cameras] are using cloud services like the iCloud to store my personal data, but I don't know whether they are secure or not because they are stored at some other place, so if I have my own device without the Internet, that is safer. It's like a physical control and my things are stored only in my place." (P2) Several interesting ideas are behind this design. First, the lock is intelligent in selectively filtering out certain types of data (e.g., gender). Second, the lock can disconnect the security camera from the Internet/network. A defining characteristic of smart home devices or IoTs more broadly is their Internet connectedness, which often supports data transfer, remote control, and other system intelligence (e.g., predictions). However, here we see P2's desire to directly control (enable/disable) the Internet connectedness. By disconnecting the security camera from the Internet, P2 felt that he has more (at least perceived) control over the data. Third, the lock (as a standalone hardware device) is physical, which affords more tangible control.

Other user controls of data. Besides data localization and disconnection from the Internet, nine participants (P5, 9-10, 12, 15, 18, 21, 24-25) desired explicit controls of their data, from preventing data collection to deleting collected data. The aforementioned P15's example of the invisible mode of a self-driving car was designed to prevent car manufacturers from collecting data about her car activities. In comparison, P5 designed a conceptual model of smart home privacy mechanisms and emphasized that a key aspect of his model was users' ability to delete data. He explained, "*the user should have a hardware option to delete data. So they don't have to necessarily go to the software to delete it.*" (P5) He believed that users should be able to delete the data collected about them and the deletion feature should be implemented as a hardware option (e.g., a physical button on the smart home device), which would be easier to use than a software control.

8.4.2 Security

Another underlying factor of participants' designs was related to security, including aspects such as authentication of multiple users, access control of user data, and network intrusion detection. Twenty participants (all participants except P15, 18-19, 24-25) considered this factor in their designs.

Authentication of multiple users. Eleven participants (P1, 4-5, 10-13, 17, 21-23) spoke to the social relationships and power dynamics in homes where there could be multiple users sharing one device. They emphasized the importance of enabling proper authentication in order to protect each family member's privacy. For example, P13 was concerned that other members in the household might be able to access her credit card information and order food from the smart fridge. To address this concern, she incorporated voice recognition in her design as an authentication mechanism for the smart fridge. She explained, *“even if someone hacks your details about the credit card to make payment, but it will still need your voice to recognize and authenticate that transaction. So that can't happen unless you do it yourself. Even if someone has credit card details, the transaction won't go through.”* (P13) While P13's voice could uniquely identify/authenticate her, she did not speak to the possibility where someone else might record and replay her voice to impersonate her (i.e., replay attacks).

Access control. In addition to authentication, authorization or access control of who can access what data was another security feature that 16 participants (P1-11, 14, 16-17, 20, 23) considered in their designs. For instance, security cameras (e.g., Nest Cam [162]) often allow anyone who logs into the compatible mobile app to see the same video content. However, P3 wanted to give users different access rights, as shown in the left screen in Figure 8.3. He designed two modes. In the *online* mode, the app shows the video feed from the camera. In the *offline* mode, the app provides a

textual description of the video feed (e.g., a person is walking), as shown in the middle screen in Figure 8.3. The access control (the right screen in Figure 8.3) determines who gets to use which mode. He elaborated, “you can decide who should be in which mode from this access management page, so some will see the text description, some will see the live video.” (P3) This feature is similar to sharing location data at varying granularity (e.g., actual address vs. city) with different entities.

In addition, some participants also designed location-based access controls. For example, P1 explained that in his design of a home automation system, the app to control his smart home devices should have a *local* mode and a *remote* mode. He should have full access to all the functions and data only when he is physically at home, which triggers the local mode. In comparison, when he is away from home, the system will enter the remote mode and should only give him partial access to the devices and none of his data should be transmitted through the Internet. We are not aware of any existing home automation products that support this feature.

Network intrusion detection. Another security feature included in three participants’ designs (P6, 22-23) was the ability to detect external intrusions into the smart home network. For example,



Figure 8.3: The online and offline modes of security cameras (P3)

P6 had a technical background, and he was particularly concerned that hackers might hack into his home network and steal his information, that his neighbors could connect to his home network and invade his privacy, or his devices could send his personal information to third parties other than the device manufacturers. To address these concerns, he designed a smart router with a built-in firewall and an app that worked with the smart router, which could be used to notify users whenever an outside intrusion was detected. He elaborated, *“the app will be able to track data sending from each device and its destination. If there is anyone who is trying to penetrate your network or trying to use your data, collect your data, this app should send you identification. You can probably reboot the device from the app to stop, kind of like a filter.”* (P6) This smart firewall would be able to track each smart home device’s data flow and notify users of any third party attempting to collect data.

8.4.3 Safety

Since our participants designed for the home environment, safety was also a concern. Twelve participants (P1, 7-8, 11, 15, 17-18, 21-25) considered this factor. For example, our participants expected that security cameras or voice assistants should be able to notify either the homeowner or the police department if someone broke into their house. If the users were in an emergency and needed help, they should be able to call for help quickly from these devices. These features are often already supported by existing products. Some participants also tried to ensure the physical safety of the home while preserving people’s privacy. For instance, P18 was concerned about the unknowingly recording of passersby by doorbell cameras. She then created a long list of key points that should be written into policies. She explained, *“I suggested that it could be first a law by the government that owners have to somehow make other people somehow aware whether that’s a sign that says you’re being recorded.”* (P18) While the doorbell cameras can arguably help improve

people's (safety) awareness of their home door area, P18 was concerned about the privacy of other people (e.g., passersby). By calling for legislation that requires a clear notice of such recording practice, P18 attempted to strike a good balance between homeowner safety and passersby privacy.

8.4.4 Usability and User Experiences

Many participants explicitly considered the usability of their privacy designs, ensuring that users have good user experiences with the designs. Twelve participants (P3, 11-15, 17-19, 21-23) considered this factor in their design. There were two broad categories of usability considerations: user characteristics and user effort.

Considerations of user characteristics. Seven participants (P11, 13, 18-19, 21-23) took people's characteristics (e.g., abilities) into account when designing their privacy mechanisms, hoping to make their designs more inclusive to a wide range of users. For instance, P11 designed an in-home robot, which could help people with various tasks in their homes. He then designed hardware access control interfaces to manage who could access the robot remotely. In the group discussion, P13 asked the following questions about P11's design, "*if the person is blind, or is it like an elderly person who cannot walk?*" (P13) These questions prompted P11 to reconsider his design. P11 then reduced the number of buttons in the interface so that it might be easier for a variety of users (e.g., children, older adults). Similarly, P18 originally designed an authentication mechanism for Amazon Echo using a physical fingerprint reader. Later he added a voice recognition mechanism for authentication because he realized that a physical fingerprint reader may be impossible or hard to use for people who lost their fingers or who have mobility impairments.

Considerations of user effort. Another usability consideration was the amount of effort required from users to utilize the privacy designs. The majority of design ideas were based on

automation (e.g., users receiving automatic alerts about their information being used). However, eleven participants' designs (P3, 11-15, 17, 19, 21-23) intentionally required explicit user effort. For example, P19 designed an improved privacy policy (summaries of most important points) for a smart thermostat and he believed that companies should be required to show the policy and users should be required to read these policies to understand the data collection. However, we note that people tend not to read privacy policies.

8.4.5 System Intelligence

Twelve participants (P1, 4-9, 13, 15, 21-23) considered system intelligence in their design. Among participants' privacy designs, we noticed two types of system intelligence: context detection and personalization.

Context detection. Since homes can have various social relationships, contexts, and thus privacy implications, six participants' designs (P1, 9, 13, 21-23) included a component of automatic context detection. For instance, P9 designed smart toys for her children but was concerned that her sensitive data might be accidentally recorded by these toys. For instance, she might be calling the bank with her credit card information while her children play with the toys, which may record and leak her private information outside the house. To address this concern, she embedded a context detection feature as part of her design. She explained, "*like when we want to have a private discussion, they [smart toys] are not allowed to [record].*" (P9) She expected the toys to be able to automatically detect when she is having a private conversation and the toys will pause their recording. Similarly, P23 designed a security camera that can automatically detect that she is not at home and start recording in the home.

Personalization. Since the home might have multiple people with different needs, twelve

privacy designs (P1, 4-9, 13, 15, 21-23) were personalized. For instance, P21, a senior citizen who lived with a portable oxygen concentrator, expected that her daughter can access her security camera to check on her well-being. However, P21 desired personalized preferences in terms of when her daughter can access the camera feed and when she cannot. P22 echoed her support, *“if I don’t want them to see certain things, I can deny them. Don’t have them on camera when I do this, this and this. I got company, I’m eating ice cream, don’t bother, that’d be good. Program it so that we don’t have to worry about it. Like an alarm, you can set an alarm based on what you are doing. That’s what a true friend would do. Let Alexa be your true friend. Tweak it up!”* P22 expressed her desire of setting her personalized preferences of access control via a voice assistant, which then can automatically enforce these preferences.

8.4.6 System Modality

We observed four forms or modalities of how participants’ privacy designs were embodied: hardware devices, apps, system modes, and policies. These modalities are not mutually exclusive. Some privacy designs had two or more modalities.

Hardware devices. Ten participants’ designs (P2, 6-9, 11-12, 14-16) were proposed as hardware devices, such as P2’s design of a physical lock for security cameras, P6’s design of a smart router, and P15’s design of a USB device for self-driving cars. In some cases, our participants intentionally designed their privacy mechanisms as a hardware solution. For example, P15 explained, *“it is just a USB, you plug it in, it will record the data, you can plug it into a computer to read...it’s small, I can take it with me and plug it in whenever I want to hide my activities...I don’t know whether it is possible to connect just using Bluetooth. So with the Internet, it can upload the data by itself, but with the Bluetooth, it can only transmit data from the device to your phone, then your phone can*

analyze the data itself, so the USB is safer.” P15’s choice of a USB was due to its portability and its perceived security (no connection to the Internet).

Apps. Another common modality was a mobile app, often features of the mobile app associated with the smart home device. Twelve participants’ designs (P1, 3-5, 9, 12, 14-15, 17, 21-23) took the form of an app. We have presented examples, such as P3’s privacy design of the security camera app, and P15’s privacy design of the app for self-driving cars.

Modes. Four privacy designs (P3-4, 12, 15) were envisioned as system modes in hardware devices or mobile apps, such as P3’s design of online and offline modes for security camera and P15’s design of visible and invisible modes for self-driving cars. These modes were often binary, privacy mode vs. regular mode. They mapped to some participants’ coarse categorizations of privacy implications (e.g., I need privacy in this case). Some of them explicitly mentioned the incognito mode (of the Google Chrome browser), which likely inspired their designs.

Policy. In addition to technological solutions, six participants’ privacy designs (P10, 18-20, 24-25) were in the form of laws and/or policies, for instance, P18’s example of legislation, which would require smart doorbell cameras to clearly notify passersby that the cameras can record them.

8.5 Discussion

8.5.1 Smart Home Privacy

The home context is complicated for privacy. First, smart home privacy covers not only information privacy (e.g., data collection and sharing) but also physical privacy (e.g., the privacy of the physical space of homes). Our participants paid attention to both types of privacy in their privacy designs (e.g., data transparency, safety).

Second, the complex social relationships and power dynamics in a home, such as parents

and children, brothers and sisters, husband and wife, owners and guests, patients and remote doctors [92], can significantly affect whose privacy is at risk or how privacy can be enacted. Many privacy designs in our study supported multiple user accounts which have been explored for shared home computers [79], but also included multi-user authentication and access control.

Third, different social relationships may suggest varying privacy norms [164]. For example, having visitors in a home changes the social context of the home and its privacy norms. Homeowners might choose not to say things in front of their visitors. Similarly, if the smart home devices record or process the conversations in the home, visitors may feel their privacy is violated. An open question for designing smart home privacy mechanisms is, *whose privacy should be protected and who should make the decision?* While most of the participants' designs were for people who live in the home, we saw some cases where the privacy of other people (e.g., passersby) was considered.

8.5.2 Design Implications

Next, we will discuss how the list of design factors we identified from our participants' privacy designs can be used to guide the design of smart home privacy mechanisms.

Data transparency and control are relevant whenever smart home devices collect data and/or can infer data about people in the home and around the home (e.g., passersby). While notice and choice are well-respected privacy principles, how to best provide and implement them is still an open question for smart homes. In terms of notice, our participants desired more transparency about what data individual smart home devices, as well as the smart home system as a whole, can collect, infer, share and use about them. Therefore, privacy designs should be considered at both the individual device level and the whole system level (e.g., P6's design of a smart router that monitors the entire smart home system).

In terms of user control, many participants desired data localization, the ability to have the smart home devices store and process the collected data in the devices locally. While client-side data storage/processing has been proposed as a privacy-enhancing technique (e.g., in targeted advertising [31] and recommender systems [97]), most smart home devices rely on servers and cloud services to store and process the collected data [41]. In fact, some proposed mechanisms in the literature also require cloud storage [70], which conflicts with our participants' desires. We suggest that designers should consider data localization as a possibility.

In addition, many participants incorporated the idea of disconnection from the Internet in their privacy designs because they felt it will give them a peace of mind because their data cannot leak out via the Internet. We note that this is concerned with the public Internet rather than the private home network (Intranet). This idea challenges a typical assumption that all smart home devices are Internet connected. Do these devices always need to connect to the Internet and should they pause their data collection and sharing if users demand so? We believe that these are important questions that designers should consider. We also note that just because devices can disconnect from the Internet does not mean they cannot collect and send data after they resume Internet connections. Fundamentally, this idea is about giving users the option to say no to data collection and sharing. Disconnection from the Internet is a simple concept that people can understand and perceive better privacy/security.

Security is closely related to privacy. Our participants considered different kinds of security attack scenarios, ranging from other members of the home accessing their data to hackers breaking into the home network (gaining control of their devices and/or stealing their data) to the devices sending their data to external third parties. In response, our participants' designs covered multiple user authentication, authorization (access control of who can see what data), and network intrusion

detection. Our participants were particularly concerned about information related to their health, finance, gender, location, and activities. Most of the current smart home devices lack these security features. We recommend designers to consider these options to address users' security concerns, e.g., if the devices allow direct interactions with users, then user authentication and authorization should be considered.

Safety was a natural concern for the home context. Many participants' designs included safety features (e.g., security cameras identifying suspicious activities). We recommend designers to consider these safety features, but more importantly, we encourage designers to think about whether safety and privacy might be in conflict. For instance, in P18's example of a doorbell camera, homeowner safety and passersby privacy might be at odds. How to reconcile when these two values conflict is another open question for further research.

Usability and user experiences are arguably important for any user-facing design. Our participants desired simple and easy-to-understand privacy mechanisms, for instance, the feature of disconnection from the Internet. In addition, they paid attention to the diversity of users and their varying needs as well as the amount of user effort required to use the designs. Our suggestion here is that designers should consider how to make their design more inclusive to various user groups and how to reduce user effort to use the privacy controls (e.g., designing privacy-friendly default settings).

System intelligence in our study covers automatic context detection and user personalization. While context-aware computing has been extensively studied, some of the designs included intelligent context detection that is currently hard to implement (e.g., security cameras automatically detecting and describing what is happening in a home). Supporting users' personalized privacy preferences has been explored in the IoT space (e.g., [63]) and designers should consider supporting

this feature in their smart home privacy designs.

Lastly, *system modality* presents the form(s) in which these privacy designs are embodied. Our participants covered four modalities: hardware devices, apps, system modes, and policies. Designers should consider this question of modality because it could influence other aspects of their privacy designs, for instance, usability and user experiences. Some of our participants' privacy designs were deliberately envisioned as hardware controls (e.g., a USB device to turn on the invisible mode) due to its perceived ease of use and portability. Many designs were also based on binary modes (visible/invisible, or online/offline). This binary model is easy to understand and use in part because people have experiences with similar models in other domains (the private mode in web browsers).

8.5.3 Policy Implications

Some participants designed privacy policies (e.g., P18's suggested policy on smart doorbell cameras). They believed that the government should play an important role in ensuring device manufacturers behave appropriately, for instance, what they are allowed and not allowed to do. Our participants also discussed the following scenario: if a user encountered some negative incidents (e.g., robbery) due to data collection or sharing by the device manufacturers (e.g., the user's personal information was collected and leaked to the wrong hands, then the user's daily schedule was inferred), will the manufacturers be held accountable? To what extent current privacy laws such as the European Union (EU) General Data Protection Regulation (GDPR) and the California Consumer Privacy Act address these questions remains to be seen.

8.5.4 Reflections on Participants' Privacy Designs

Our co-design study aimed to give voice to people (users and non-users of smart homes), who are often not included in the design of privacy controls. Our participants contributed many novel

ideas, such as stand-alone hardware devices as privacy controls (e.g., a physical lock for security cameras), seamless identification and authentication of multiple users, and automatic context-based personalized privacy controls (e.g., a smart toy selectively pausing its recording based on the detected context and a user's contextual preferences).

However, our participants' privacy designs also have limitations. First, since most of our participants were not technically savvy, so their designs did not cover all the privacy-enhancing techniques found in the literature, for instance, adding noise to the home network traffic to reduce data inferences [17]. Second, their designs did not address all potential privacy risks in smart homes, for instance, the risk of secondary use of data (using the collected data for a different purpose) [143]. Their designs also did not address the case that the manufacturers collected users' data for a reasonable purpose but then shared the data with third parties. Third, some of their designs are currently hard to implement (e.g., security cameras providing real-time textual descriptions of the video feed). However, this was by design because we did not want to limit our participants' creativity. Fourth, many designs could potentially pose privacy or security risks themselves. For example, the smart router monitoring the entire home network could be privacy intrusive itself. Remotely resetting a door lock could also have security risks.

All of these novel design ideas and concrete limitations suggest that when designing privacy mechanisms for smart homes, inputs from both users and privacy experts are needed.

8.5.5 Limitations of Our Research

Our study also has several limitations. First, we asked the participants "what can potentially go wrong with smart home devices," which may prime our participants to focus on the negative aspects. However, we believe that our participants were unlikely to make up issues because (1) participants

were not required to answer this or any question, (2) the privacy concerns were often raised by multiple participants voluntarily, and (3) participants discussed similar concerns in other domains (e.g., web tracking). In addition, as we discussed in the method section, we asked about the pros and cons of smart homes first. Our results showed that our participants covered both the pros and cons in their considerations and their designs often reflected the trade-off between benefits and risks. Second, while our participants had very diverse backgrounds, we did not include anyone younger than 18. While our study did not focus on smart toys, some participants designed for smart toys. Having children as part of the co-design team would have been valuable for the privacy designs for smart toys. Third, all of our participants' designs were low-fidelity paper prototypes rather than interactive high-fidelity prototypes. Therefore, they might have missed potential challenges of their designs or opportunities to improve the designs. Fourth, our study focused on users and might have missed perspectives (factors) from other stakeholders such as device manufacturers. Our student designer who helped our participants with their design was not experienced in hardware designs, so the help was also limited.

8.5.6 Future Directions

The aforementioned limitations point to a few directions for future research. The limitations of our participants' privacy designs suggest that future work should not only continue to explore user-generated designs but also critically evaluate these designs in terms of their feasibility, usability, privacy, and security. These evaluations can shed light on how to effectively combine end users' ideas and expectations with experts' knowledge. These insights can then inform how these designs should be adapted and implemented in practice. Furthermore, future co-designs could consider educating users on privacy/security risks and countermeasures before starting the actual design. This

might lead to additional designs. Lastly, we did not observe any differences in terms of perceptions or design factors among the users, non-users, and interested users. This might be due to the small sizes of different user groups in our study. Future work can further explore potential differences among various types of users.

8.6 Conclusion

Smart home devices are gaining momentum albeit with serious privacy challenges. We conducted a co-design study to understand how people desire to protect their privacy in the smart home context. From participants' designs of smart home privacy mechanisms, we identified six important design factors they considered: data transparency and control, security, safety, usability and user experience, system intelligence, and system modality. We discuss how these factors can guide the design of smart home privacy mechanisms. Future research should try to involve more stakeholders (e.g., device manufacturers) in the privacy design process, and further explore and evaluate user-generated privacy designs.

8.7 Acknowledgement

We thank our participants for their valuable input. We thank Elias Greendorfer, Brianna Dym, Bryan Semaan, anonymous reviewers and shepherd for their thoughtful feedback. This work was supported in part by NSF Grant CNS-1464347.

Chapter 9

Smart Home Privacy: Bystanders' Perspectives

As the Internet of Things (IoT) devices make their ways into people's homes, traditional dwellings are turning into smart homes. While prior empirical studies have examined people's privacy concerns of smart homes and their desired ways of mitigating these concerns, the focus was primarily on the end users or device owners. Our research investigated the privacy perceptions and design ideas of smart home bystanders, i.e., people who are not the owners nor the primary users of smart home devices but can potentially be involved in the device usage, such as other family members or guests. We conducted focus groups and co-design activities with eighteen participants. We identified three impacting factors of bystanders' privacy perceptions (e.g., perceived norms) and a number of design factors to mitigate their privacy concerns (e.g., asking for device control). We highlighted bystanders' needs for privacy and controls, as well as the tension of privacy expectations between the owners/users and the bystanders in smart homes. We discussed how future designs can better support and balance the privacy needs of different stakeholders in smart homes.

9.1 Introduction

Various Internet of Things (IoT) devices have made their way into people's homes, turning traditional dwellings into *smart homes*. These devices infiltrate households and aim to provide efficiency and

usability for homeowners. At the same time, the Internet-connected nature and the amount of data collected by these IoT devices pose great privacy risks to users. A 2015 report by the Federal Trade Commission has shown that fewer than 10,000 households with smart home IoT devices can generate 150 million discrete data points per day [56]. This massive amount of data allows a variety of analyses which are not possible using other data [56].

From the perspective of smart home users, many prior studies have investigated users' privacy perceptions of smart homes and have discovered a number of privacy concerns, such as sensitive data collection [237], data sharing [235], and data misuse and re-purpose [143]. However, little is known about other stakeholders' privacy perceptions in smart homes, such as visitors, tenants, other family members, etc. This is an important aspect to consider in the development of smart homes because these other stakeholders' privacy is often ignored and can even be violated without their knowledge. For example, a recent news article reported that smart home devices that can record people's voice often pick up other people (e.g., spouse, friends, kids) talking in the background [207]. In the real world, such cases often happen in scenarios where guests visit other people's homes and are exposed to other people's smart home devices. This case demonstrates that the privacy risks for other people in smart homes indeed exist, however, their understanding and privacy perceptions are understudied in the prior literature. In addition, people may face other situations that can potentially invade their privacy, e.g., an Airbnb host may have access to security camera data while the tenant may not due to the power imbalance between the owner and the tenant [68].

This chapter focuses on one specific group of stakeholders in smart homes, i.e., bystanders. In this chapter, we use smart home **owners/users** to denote people who directly purchase smart home devices. In other words, owners/users in our study context refer to people who *own* smart home devices. We use smart home **bystanders** to refer to people who do not own or directly use these

devices but are potentially involved in the use of smart home devices, such as other family members who do not purchase the devices, guests, tenants, passersby, etc.

Our study attempts to fill the gap in the literature by specifically investigating smart home bystanders' privacy perceptions. In particular, through a focus group study with eighteen bystanders in six groups, we aimed to understand the concerns that bystanders had towards smart homes under a variety of social contexts. In the last part of the focus group, we adopted a co-design approach [187, 195, 206] and collaborated with bystanders to design privacy mechanisms to mitigate their privacy concerns in smart homes.

This chapter makes three contributions. First, we investigate smart home bystanders' privacy concerns and identify several factors that affected their privacy perceptions, such as trust towards the owners. Second, the design activity results in a number of design factors that bystanders considered when designing privacy mechanisms to protect their privacy. These perceptions and the design factors demonstrate bystanders' needs for privacy and some control mechanisms. We highlight the cooperative design mechanism as a unique aspect of bystanders' privacy designs and advocate for addressing privacy needs for both owners/users and bystanders through potential collaboration. Third, we make a number of concrete design suggestions to better support different stakeholders' privacy needs in smart homes.

9.2 Related Work

In this section, we present the prior literature on privacy issues in smart homes in general and different privacy mechanisms. We then summarize prior research on understanding the bystanders' perspective in introduce the bystanders' perspective in prior research and explain why our results fill a significant gap in the literature.

9.2.1 Smart Home Privacy Risks and Concerns

Given improvements in smart home technologies, researchers have started to look at the potential privacy and security risks associated with changes in technology. These risks include but are not limited to: the possibility of identity theft and device reconfiguration [19], the inference of user activities at home through smart home network traffic analyses [17], as well as risks caused by human factors (e.g., weak passwords) and system flaws (e.g., unauthorized system modifications) [109].

User studies also looked at the privacy issues of smart homes from the perspective of end-users. From the smart home level, people were concerned about the possibility of Internet attacks and data abuse [240]. Zeng et al.'s interview study discovered people's concerns on video recording, data collection, and analysis, as well as network hacking [235]. However, their participants also outweigh cost and interoperability over privacy and security [235]. Worthy et al. found an association between people's trust towards the entities that collected their information and their desired control of such information, and they argued that less trust would lead to a greater level of desired control [223]. Brush et al. further claimed that a "difficulty achieving security" was one road blocker towards large adoption of smart home devices. On the other hand, according to Zheng et al., some people believed that their privacy was well protected by the entities who collected the information, which may result in new privacy risks [237]. In a slightly different direction, Apthorpe et al.'s survey study investigated the privacy norms in smart homes using the theory of Contextual Integrity and found a number of factors that could influence specific norms, such as their purposes of device usage (e.g., in an emergency situation) and device ownership (e.g., how many devices users owned) [18]. To explore what information should be provided to the users in smart homes, Jakobi et al. conducted a long term study and identified users' information demands to understand smart home system

performance and communications [112]. They found that in the initial phase of smart home usage, users preferred to access detailed information of the smart home environment through web-based platforms, whereas in the later stage, users preferred to only know the exceptions where something went wrong [112]. In a recent work, Barbosa et al. discovered that factors such as “consent not given” and “sensitive data collection” could make users less comfortable, and factors such as “user control” and “user awareness” could make users more comfortable regarding the data collection in a smart home [26].

Concerning especially about smart home devices, users have shown their uncertainty of data practices in smart TVs, including data collection, usage, re-purposing, and sharing with third-parties [143]. For smart toys, parents were concerned about data collection and sharing abilities while children were uncertain of whether their conversations with the toys could be heard by their parents [149]. Interestingly, smart speaker users generally expressed no concerns in their perceptions, but the rationale behind their perceptions (e.g., they did not mention any concern because they had strong trust towards device manufacturers) could lead to more serious privacy risks [129]. In fact, when users did not have concerns toward smart home devices, it did not mean that the users do not face any privacy risks. For example, users’ activities can be inferred using their smart home network data and thus pose various privacy risks to the users [16].

9.2.2 Bystanders’ Privacy Concerns

The privacy of bystanders has been studied in a few contexts. For example, Denning et al. found that bystanders assumed augmented reality wearable devices were used for recording [66]. In their study, they reported cases in which bystanders had negative reactions to these devices and expected to be asked for consent before they were captured by them [66]. Bystanders’ privacy was a concern

of people who recorded audio or video for lifelogging [108]. As such, the recorder chose to discard, modify, or not share the audios or videos to respect the privacy of bystanders [107]. Wang et al. studied bystanders' privacy perceptions of drones in a variety of usage scenarios and discovered several privacy concerns held by bystanders, such as peeking and stalking, as well as surveillance in public places [228]. Interestingly, when comparing drone bystanders' privacy perceptions with the drone controllers', Yao et al. identified several mismatches. For example, bystanders were heavily concerned about their faces being recognized in drone footage while controllers believed that drone cameras were satisfactory for such purposes [232]. Motivated by this line of research, in this chapter, we investigate the privacy perceptions of bystanders in the context of smart homes. Besides, we aim to identify, if any, mismatches in the perceptions between users/owners and bystanders to further inform future privacy designs.

9.2.3 Smart Home Privacy Mechanisms

Many technical solutions have been proposed to mitigate smart home privacy risks and concerns. To reduce the potential of inferences, Yoshigoe et al. designed a software-based system to automatically inject synthetic network packets to obscure legitimate network data flow [234]. Apthorpe et al. demonstrated the effectiveness of introducing noise data to shape smart home network traffic [17] to obscure the real traffic and prevent data loss. A more recent work by Datta et al. introduced a Python library so that developers could easily implement traffic shaping for IoT devices [64]. To achieve better access control to users' data, Moncrieff et al. designed a tool to manage access privileges dynamically and automatically according to some contextual factors in home surveillance, such as users' activity and location in the home [156]. Through capturing user-defined privacy zones and generating corresponding policies dynamically, Arbo et al.'s framework aimed to ensure data

security for smart home devices [19]. Chakravorty et al.'s system only granted users access to their data if these users were successfully re-identified by the system [47]. To increase data transparency and user awareness as well as to facilitate user control, Das et al. envisioned an infrastructure to personalize users' privacy notices based on their privacy preferences in IoT devices [63]. Wang et al. built a tool which could blur faces captured by cameras based on users' self-defined rules [208]. McReynolds et al. suggested that smart toys should communicate their recording capabilities to the parents and children [149]. Lastly, Lin et al. suggested a mechanism in which supporting systems should auto-configure new devices added to the smart home network based on the most secure setting to ensure a safe home environment [137].

Arguing that the above privacy mechanisms were proposed or implemented either by researchers or experts without users' input, Yao et al. took a user-centered approach in which they involved smart home users and co-designed privacy-enhancing mechanisms to alleviate users' privacy concerns [229]. Their study suggested several factors and features that users considered in their designs, such as network intrusion detection and data localization [229].

9.2.4 Gap in the Literature

We aim to explore how **bystanders perceive privacy issues in smart homes**. This is an important question for two reasons. First, bystanders' privacy issues are usually omitted. This is because bystanders are not the owners nor users of smart homes, however, they are subject to usage of smart home devices without their knowledge for most of the time. Understanding bystanders' privacy perceptions can broaden our knowledge of smart home privacy issues more holistically. Second, the rapidly growing popularity of smart home devices has created many interesting yet controversial social contexts in which users receive benefits from these devices but may put bystanders' privacy at

risk (e.g., using an Internet-connected security camera in an Airbnb apartment [94] and the adoption of voice assistants in hotel rooms [49]). Understanding the factors that influence bystanders' privacy perceptions of smart homes can provide insights into how to better suit the needs of both users and bystanders collaboratively.

In addition, inspired by Yao et al. [229], we deem to see what privacy mechanisms bystanders desire to mitigate their privacy concerns if exist. The results can inform future privacy designs and illuminate how to support the privacy needs of both bystanders and users. Next, we will describe our methodology in detail.

9.3 Method

To explore bystanders' privacy expectations and how they desire to protect their privacy, we conducted six focus groups with an average of three participants in each group and a total of eighteen participants. We chose to do focus groups instead of one-on-one interviews because we hoped to encourage interaction between participants and spark the discussion by bringing different experiences as bystanders. The average length of the sessions was 1.5 hours. Upon completion, each participant was paid \$15. This study is approved by our university IRB.

9.3.1 Study Settings

Participants recruitment. We recruited our participants primarily through Craigslist, word-of-mouth, and local senior citizen centers. When prospective participants first reached out to us, we asked them to fill a pre-screening survey to obtain their demographic information. We deliberately selected participants from various gender identities, age groups, occupations, and with different levels of smart home experiences. We carefully framed our study as “a focus group study to understand your perceptions of smart homes” without mentioning anything related to “privacy” to

Table 9.1: Summary of participants' demographics

Group NO.	Participants	Gender	Age	Occupation	Experiences	Scenarios
1	P1	M	18-25	Student	Owner	S1, S2
	P2	F	18-25	Student	Owner	S1, S2
	P3	M	18-25	Student	Owner	S1, S2
2	P4	M	36-45	Hospital employee	Owner	S1, S2
	P5	M	26-35	Government employee	Owner	S1, S2
	P6	F	26-35	Student	Experienced	S1, S2
3	P7	F	18-25	Paralegal	Owner	S2, S3
	P8	M	26-35	University staff	Owner	S2, S3
	P9	F	36-45	Postal expeditor	Experienced	S2, S3
	P10	M	36-45	Civil engineer	Owner	S2, S3
4	P11	M	>65	Retired	Non-user	S1, S3
	P12	F	26-35	Unemployed	Experienced	S1, S3
5	P13	F	36-45	Sales	Experienced	S1, S3
	P14	M	56-65	Retired	Non-user	S1, S3
	P15	F	>65	Retired	Non-user	S1, S3
6	P16	M	26-35	Editor	Owner	S2, S3
	P17	F	26-35	Filmmaker	Owner	S2, S3
	P18	F	36-45	Chef	Experienced	S2, S3

prevent potential bias. We summarize the demographics of the participants and their groups in Table 9.1.

Pilot study. Drawing from prior research [66, 228, 107, 233, 229, 187, 195, 206], we developed a list of questions and activities to probe participants to think about the potential benefits and concerns of smart homes from the perspective of bystanders. We ran two pilot study sessions with seven participants, gained a few insights, and then made several changes to our study protocol. First, in the initial protocol, we only asked participants to think from the bystanders' perspective. In the pilot study, we found that our participants tended to think from the owners' perspective. Thus, we added a question and asked the participants to recall the last time they visited other people's places where smart home devices were installed to better situate them as bystanders. Second, the original

protocol asked participants about their general perceptions and concerns of smart home devices. However, in the pilot study, we found that our participants focused more on the negatives of smart homes in the scenarios. Thus, to reduce potential priming, we asked participants to discuss the benefits first in each scenario to ensure they think thoroughly. Third, when we asked our participants to create prototypes to illustrate their ideas, most of them expressed confusions on the definition of “prototypes”. Thus, we added a brief introduction session to show participants a few examples of different types of prototypes (e.g., diagram, low-fidelity paper prototypes, wireframe, etc.) to help them know the expectations from the design activity. The final study protocol is described in detail in the next section.

9.3.2 Study Flow

The goal of the study is to understand bystanders’ privacy expectations in smart homes and their desired privacy controls. We divide each study session into three main parts.

Part 1: general understandings. We started each session with a round-table introduction. We first asked each participant to talk about their understandings, experiences, and general perceptions regarding smart home technologies. Regardless of their prior knowledge, we provided a working definition of smart home [41, 115], “*a home consists of different sensors, systems, and devices, which can be remotely controlled, accessed, and monitored.*” We showed them a few smart home devices and explained their primary functions. These devices included voice assistants, security cameras, smart toys, and a set of smart appliances.

Before introducing the concept of “bystanders”, we first asked participants to discuss the pros and cons of smart homes in general. We then started to shift the perspective towards that of bystanders by asking participants to describe their past experiences and thoughts in other people’s

smart homes. Next, we deliberately introduced our definition of a “bystander” in the smart home context, framing it as “*people who are not the owners nor the primary users of smart home devices but can potentially be involved in the device usage.*” We then asked our participants to think about themselves as the bystanders in the remainder of the study.

Part 2: scenario-based discussions. Similar to [213, 231], we introduced three scenarios in our study to (1) capture participants’ contextual privacy perceptions and (2) better situate our participants and nudge them to think as bystanders. The three scenarios were inspired either by findings in the literature or from the news, including: (1) *the temporary residency scenario (S1)*: you rented an apartment for three days through Airbnb and an Internet-connected security camera was installed in the apartment [94]; (2) *the playdate scenario (S2)*: you took your child to a playdate and there was a smart toy for the kids to play with [87]; and (3) *the cohabitant scenario (S3)*: you live in your own house and your spouse purchased an Amazon Echo [237]. These scenarios were designed to represent a variety of factors, including different application contexts (e.g., temporary resident in an Airbnb apartment, friend’s house, your own house), social relationships (e.g., tenants and owners, guests and owners, husband and wife), and different devices (e.g, Internet-connected security cameras, voice assistants, smart toys). It is worth noting that we did not limit our participants to these devices. All participants were told that they could also discuss other devices they would like to add in each scenario. We also did not explicitly mention or investigate the case of hidden devices (e.g., devices that were purposefully hiding in a place by the owners or not obvious to the bystanders) because we did not want to prime our participants to think about devices that were not obvious to them which could adversely influence their perceptions. Due to time limitations, each group was asked to discuss two of the three scenarios. For each scenario, we asked our participants to discuss

the benefits of smart homes as bystanders, then we moved forward to discuss their concerns.

When finishing the discussion of the scenarios, our participants demonstrated to have grasped the concept and the role of bystanders for the study. These prior activities resulted in participants' understandings of a wide range of potential benefits as well as potential concerns, including their privacy concerns and expectations. We then focused on the privacy concerns and expectations emerged from the discussion and continued the study with a co-design activity.

Part 3: co-design of privacy mechanisms. The goal of this activity was to co-design privacy-enhancing mechanisms with our participants based on their privacy concerns and expectations in smart homes as bystanders. We first situated participants in a friend's house with a number of smart home devices (i.e., voice assistants, security cameras, and smart toys, all adopted from the previous three scenarios) presented, then we asked them to brainstorm their desired features to mitigate their aforementioned privacy concerns and created prototypes to illustrate their ideas. We chose to use a different and more general scenario with the same devices for the design activity rather than using the three scenarios from the previous part for two considerations. First, each group of participants only discussed two of the three scenarios, thus none of the three scenarios were shared by all three groups. Besides, involving all participants in the same scenario made it easier to synthesize our findings. All previous scenarios were designed with different combinations of contextual factors in mind which made it difficult to understand the rationale behind participants' designs. We provided a set of tools (e.g., Post-It notes, color pens, paper board, color papers) for their convenience. We encouraged participants to collaborate and discuss their ideas with others and with researchers, break the existing technological and policy limitations, and potentially design futuristic and speculative solutions.

9.3.3 Data Analysis

Study recordings. All study sessions were audio-recorded after obtaining participants' consent. Then two co-authors transcribed all recordings and conducted a thematic analysis [32]. We read all transcriptions a few times to familiarize ourselves with the data, then coded one transcription (i.e., the transcription of one complete focus group) together at the sentence level. After generating the initial codebook, we independently coded the same subset of data. When new codes emerged from this process, we added them to the codebook. Upon completion, we compared and discussed the coding and merged their codebook. The inter-coder agreement was 0.81 (Cohen's Kappa), which is considered good [86]. Then we coded the rest of the data using the updated codebook. The final codebook contained over 120 unique codes, such as "bystander action", "trust in the owner", and "wish for data local storage". We further grouped all codes into seven themes, including "general perceptions", "perceived norms", "bystanders' awareness", "privacy-seeking behaviors", "cooperative mechanisms", "bystander-centric mechanisms", and "demographics". We deliberately checked all the codes to ensure they were assigned to the appropriate groups.

Image data. We collected participants' prototypes of their designs. Using the same analysis method in Poole et al.'s study [173], two co-authors coded all the elements in the prototypes. These elements covered everything that was covered in the prototypes, including all components (e.g., stakeholders, devices), visual elements (e.g., buttons, colors), information flow (e.g., information type, flow directions), and other parts (e.g., notes). We followed the same coding procedure as described above and resulted in a codebook with over 100 codes. We further grouped all codes into six themes. The inter-coder agreement was 0.85 (Cohen's Kappa), which is considered good [86].

9.4 Results

Next, based on our thematic analysis, we first focus on the themes related to our participants' privacy perceptions, then we present the main themes from our participants' privacy designs.

9.4.1 Participants' General Perceptions

Our participants discussed several benefits of adopting smart home technologies, including enabling home automation, providing remote access, ensuring home safety, etc. They acknowledged several benefits of using smart home devices in all scenarios. We summarize the perceived benefits in this section. The full list of the perceived benefits and risks discussed by our participants is attached in Table 12.1 in the Appendix. For example, in the temporary residence scenario, bystanders mentioned that if the apartment was a shared space, having some smart home devices (e.g., an Internet-connected security camera) could provide them a peace of mind for safety purposes. Our participants' concerns of smart home devices varied among individuals and across different scenarios. In general, bystanders' had more privacy concerns in the temporary residence scenario and the playdate scenario than the cohabitant scenario. Bystanders also expressed more concerns regarding the video and audio data collected by devices with microphones and cameras (e.g., voice assistants, security cameras) but barely any concern with other devices (e.g., smart coffee makers).

Through our analysis, we identified three major aspects to shape bystanders' privacy perceptions of smart home devices: their perceived norms in different contexts, their awareness of smart home devices and device behaviors, and the potential ways to control their privacy. In the following section, we unpack the three aspects and describe how each aspect shapes bystanders' privacy perceptions.

9.4.2 Three Aspects of Bystanders' Perceptions

Perceived Norms.

Perceived norms refer to bystanders' believed values or standards in a given context. For example, our participants felt that as bystanders, they should not directly control the devices without the owners' permission. Such norms are deeply rooted into specific contexts and contain four primary facets: (1) perceived device utility, (2) perceived social relationship, (3) perceived trust, and (4) length of stay. Changes to any one of the facets may cause changes to bystanders' perceived norms, which further influence their privacy perceptions. We present the four facets below.

Perceived device utility. The first facet of the perceived norms is bystanders' perceived device utility. This facet was brought up by eight participants (P1-3, P5-6, P8, P14, P17). Bystanders held different opinions on whether smart home devices were needed in this context. For instance, in the temporary residency scenario, several bystanders believed the legitimacy of having Internet-connected security cameras installed in the apartment for security reasons (e.g., if the shared space was broke in or needed surveillance) as long as the cameras were not in the bedroom or living room. However, some other bystanders were completely against the use of cameras inside an Airbnb apartment since they preferred to have their privacy. In the playdate scenario, P6 shared her opinion about smart toys:

“My concern would just be that the kid grows up used to having invasive devices present so I would prefer the Alexa not in the house and the toy not in the house ... because I would imagine the purpose of the toy is to get children used to having smart devices and other smart amenities. Personally, I would want my child to be more concerned about their privacy.” (P6)

P6's perceived utility of smart toys significantly affected her perceptions. She believed that this

device was designed to immerse the children in an environment full of “smart amenities” so that they would get used to these devices, instead of being an object that the children could simply play with. In the long run, children would be used to the data collection and potential privacy violations associated with such devices. Thus, she would be concerned.

Social relationship. The social relationship in the study mainly refers to bystanders’ relationships with the owners of smart home devices. Such a relationship also helps to shape bystanders’ perceptions. This facet was discussed by seven participants (P1, P3, P5-7, P10, P14). For example, P10 commented on how social relationships impacted his perception:

“Listening to children doesn’t make sense to me. I would like that kid playing with another toy. If it is my friend’s kid and they assure me that the toy is fine, I would be ok. If I’m not close with the other parent, I would try to politely get my kid not to play with it. ” (P10)

P10 believed that smart toys needed to listen to and record children’s conversation to provide the “smart features.” His privacy perception of the smart toys from a bystander’s perspective largely depended on the relationship between himself and the owners. In this case, he would allow his children to play with the smart toy if the owner was a close friend. In the cohabitant scenario, P16 was against the use of any smart home devices due to the potential collection of his data. However, he also acknowledged that he would still use the voice assistant if his wife bought it and wanted to use it, which further confirmed the role of social relationships in influencing bystanders’ perceptions.

Perceived trust. Perceived trust refers to bystanders’ perceived trust level towards different potential stakeholders involved in smart homes, such as the owner, the device manufacturers, as well as the potential mediators (e.g., Airbnb as the company in the temporary residency scenario).

This was discussed by seven participants (P1, P3, P6-7, P10-12). For example, in the temporary residency scenario, bystanders discussed their trust towards the owner of the property. P12 refused to book an apartment even if she was told explicitly that there was a security camera in the home. She explained:

“I like my privacy. You can say there’s ownership and it’s their building, but I guess I don’t trust people, expect the worst of people I guess. I wouldn’t like that.” (P12)

P12’s lack of trust towards the apartment owner was the primary reason for not accepting the usage of a security camera. Besides, bystanders also mentioned how their trust towards the manufacturers and the mediators influenced their privacy perceptions (e.g., they tend to trust the household company names and believe in their privacy policy), further confirming the findings from prior research [129, 235].

Length of stay. Length of stay is a unique facet in smart home norms from bystanders’ perspectives. It refers to how long a bystander stays at one particular location. This was discussed by three participants (P2, P4, P15). Our results suggested that the length of stay also impacted bystanders’ privacy perceptions. Generally, bystanders were more concerned with smart home devices if they were exposed for a longer time, although different participants had different interpretations towards what a “long time” meant. For example, P2 believed that in the temporary residency scenario, “three days” could be considered as a long time, thus using an Internet-connected security camera and other smart home devices was not acceptable. However, P4 had a different opinion on this and explained:

“I think this would be different if this was at a friend’s house or I was renting an apartment and the apartment owner was like we have to have cameras on you at all times because this is a short

amount of time. If this was a long time, like three weeks, or if it was someone I don't trust and there were no laws against it, then that is a red flag.” (P4)

In this response, P4 not only explained his perceived “a long time” being three weeks but also further echoed the aforementioned trust facet. His privacy perceptions as a bystander would vary if these two facets changed. P4’s quote also hinted that regulations would play a role in his perceptions. However, we did not observe recurring theme around users’ legal expectations. Future research may dive deep into this area.

The above examples further suggested that, due to the complex social dynamics in smart homes, the social norms were not always clear, which made privacy management more difficult in smart homes when considering both bystanders and owners. We will further unpack the implication in the Discussion section.

Bystanders’ Awareness of the Smart Homes

The second aspect influencing bystanders’ perceptions is related to their awareness of the surrounding environment. Such awareness further includes their awareness of smart home devices one owners’ property and their knowledge of smart home device behaviors.

Awareness of device existence. This was discussed by thirteen participants (P1-4, P6-8, P10, P11-P13, P16, P17). Many bystanders acknowledged that often times, they did not pay attention to or look for any smart home devices even though those devices were becoming more and more ubiquitous. It is worth noting that although we did not explicitly investigate the issue of hidden devices, the awareness issue still emerged from our study as one main factor that impacted bystanders’ perceptions. Our participants discussed their thoughts about the consequences of not knowing the existence of these devices. As P8 stated, controlling the devices was not as big an

issue since he could negotiate with the owners. However, unawareness of the devices was a vastly different story. One thing worth noting from our study was the fact that people struggled to tell that something was, in fact, Internet-connected and essentially “smart” as these devices made their way into people’s homes in a variety of formats. In the playdate scenario, P7 explained her concerns as a bystander:

“The Google Home and Amazon Alexa are controlled by awake words, they look like devices. That thing [smart toy], it goes back to the awareness factor. If I walked in, I would never know that is a smart toy. I don’t know where it is going, I don’t know what it is recording, I don’t know if someone knows where my children are. That is when it gets concerning. Because those things like the Google Home and Alexa, people can track where you go. That [a smart toy] gets a child involved. That is where I get concerned as a bystander, I want to be aware of the things.” (P7)

P7, who owned a few smart home devices, acknowledged that she did not know that the “dinosaur-shaped toy” was a smart toy. Given her prior knowledge regarding the tracking capability of similar voice assistant products, she would be concerned that her kids were accidentally exposed to another tracking device without her even knowing about it.

Awareness of device behaviors. Relatedly, bystanders also lack awareness of device behaviors and thus are not sure whether they are facing any risks or not. Nine participants (P2, P4, P6-9, P12, P14, P18) discussed this aspect. In the family cohabitant scenario, P6 shared her own story in her parents’ house:

“My dad has a Google Home which, he doesn’t use it to control music although sometimes it will, he might use it to ask a question about the weather or if we are having a conversation and he doesn’t recall something he will ask the Google Home. It is kind of annoying when he isn’t there,

I unplug it because it is kind of weird like if we are talking just amongst ourselves and he says something vaguely like “okay Google” which is the activation thing, it will start listening and it is kind of weird. He works during the day so if I am there on the weekend. He has two one in his bedroom and one in his living room. I don’t spend any time in his bedroom so I only unplug the living room one. I don’t find the device useful and I don’t know anything about it. I don’t know if it is always on or whatever.”

P6 was against the usage of Google Home in her parents’ house. She chose not to expose herself in front of the Google Home by either unplugging the device or not staying in the same room with the device as long as her father was not in the house. This was primarily due to her lack of awareness of the device capabilities and what the device might bring to her.

9.4.3 Privacy-Seeking Behaviors

The third aspect of bystanders’ perceptions of smart homes is their privacy-seeking behaviors. Privacy-seeking behaviors refer to different ways people adopt to mitigate their privacy concerns and protect their privacy. Unlike the owners or users of smart home devices who directly set them up or turn them off the devices if they chose to, bystanders generally do not have access to directly control the devices, or simply do not believe that they should control the devices. In our study, several bystanders (P1, P4, P6, P9, P13, P15-16) mentioned a few ways of how they seek to protect their privacy. One common way in the temporary residency scenario was to cover the security camera if needed. In the family cohabitant scenario, P5 chose to place the voice assistant in a place where he only stayed to make food as a way to protect his privacy against the voice assistant.

In the playdate scenario, many bystanders mentioned that they would directly talk to the owners to either obtain more information about the toys or simply ask the owners to turn off the toys.

However, bystanders also mentioned the potential caveat in doing so, as P1 stated:

“I feel like I am not in the place to be like “hey can you turn it off?”, so I probably won’t, but it still makes me feel uncomfortable.” (P1)

This quote demonstrated that simple and direct privacy-seeking behaviors might create socially awkward situations. In this particular case, an awkward situation was caused by the perceived imbalanced power structure in the owner’s home. As a result, bystanders ended up giving up seeking for privacy controls.

The three aspects discussed above, on the one hand, shape bystanders’ privacy perceptions; on the other hand, provide insights into bystanders’ privacy expectations in a variety of smart home contexts. Building on top of the above results, we present the findings from the follow-up co-design activity during which bystanders carried out various ideas to enact their privacy concerns and meet their privacy expectations.

9.4.4 Privacy Designs

In our study, we include a co-design activity to help us understand bystanders’ desires in privacy controls and mitigation strategies. The activity provides a chance for researchers and bystanders to design together and cope with bystanders’ privacy concerns. It is worth noting that although the designs are for a pre-defined smart home scenario, bystanders discussed the possibility of extending these designs to other scenarios.

We synthesize all designs and extract the design factors from our analysis. We first group all the factors based on design purposes, i.e., the privacy problems to be solved. We identify three purposes which can be mapped to the three aspects in bystanders’ privacy perceptions, i.e., *expressing preferences* and *asking for device control* aiming to clarify bystanders’ perceived norms;

Table 9.2: Summary of bystanders’ privacy design factors, organized based on design purposes and large categories.

Categories	Purposes	Factors
Cooperative mechanisms	Clarify norms	- Express preferences
		- Ask for device control
Bystander-centric mechanisms	Increase awareness	- Detect nearby devices
		- Inform device behaviors
	Provide controls	- Limit data collection
		- Control data processing

detecting nearby devices and *informing device behaviors* aiming to increase bystanders’ awareness; as well as *limiting data collection* and *controlling data processing* aiming to empower bystanders with more control over their privacy.

We then conceptualize all design factors and purposes into two larger categories: cooperation mechanisms and bystander-centric mechanisms. Cooperation mechanisms refer to the designs that require communication between bystanders and owners to collaboratively resolve bystanders’ privacy concerns, while bystander-centric mechanisms refer to the designs that require only bystanders’ effort alone to meet their privacy expectations. The summary of the results can be found in Table 9.2. It is worth noting that these design factors are not mutually exclusive. Many design ideas carried out by bystanders cover a few of these factors.

9.4.5 Cooperative Mechanisms

One primary reason for bystanders’ concerns is the lack of communication between bystanders and owners. This is due to either the lack of communication channels or the potential social awkwardness and confrontation in face-to-face communication. Thus, bystanders’ designs provide technological alternatives to enhance communication. Through effective communications, bystanders hope to establish or clarify contextual informational norms in the owners’ smart home with respect for their

privacy. From this perspective, seven participants considered the cooperative aspect in their designs, focusing on expressing their privacy preferences and asking for some device controls.

Express preferences. Seven bystanders (P1, P5-6, P9-10, P13-14) wished to express their privacy preferences to the owners through their designs. For example, in the playdate scenario, P5 designed an app interface in which he could specify his preferences regarding smart toys and other recording devices in the home. Once the preference was set, the owner would receive notifications in which the owner had the choice to either honor the preference by changing the smart toy settings or ignore the preference.

Ask for device control. Another cooperative element in bystanders' designs centers around asking for some controls of the devices in the owners' home. Six participants (P1-5, P10) included this aspect. Such designs provide a unique perspective in bystanders' privacy expectations since when discussing their perceptions, most bystanders acknowledged that they should not expect controls of others' devices. However, these privacy designs reflected that bystanders expected some controls over owners' devices to protect themselves. For example, P4 designed an app (Figure 9.1) to detect smart home devices in the owner's house. When he connected to the owner's home Wi-Fi, the app would provide a list of connected devices. If he was not comfortable with any of the listed devices and preferred for it to be turned off, he would make a request to the device directly through the app. The owner would be notified as well and would need to approve the request.

There are several insights behind this design. First, the app starts with a transparency feature that can detect all the devices connected in the home network so that bystanders are more aware of their environment. We will further unpack the awareness and transparency aspect in the Discussion section. Second, as a bystander, P4 prefers direct communication with the devices in the owners' home and the ability to turn the device on and off if needed as a way to protect his privacy. However,

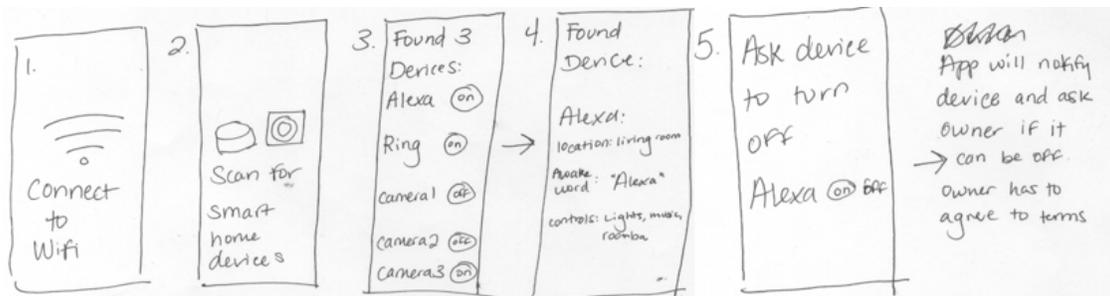


Figure 9.1: The app design by P4.

this idea can be considered as an intrusion for the owners, which further highlights the mismatch and tension between the owners and bystanders. Third, the fact that P4 designed an app to reflect his expectations of controlling the device rather than directly talking to the owner face-to-face suggests the potential of using our everyday technologies to avoid social confrontations in smart homes.

9.4.6 Bystander-Centric Mechanisms

The bystander-centric mechanisms refer to designs that only require effort from bystanders. Several bystanders' designs include many bystander-centric features, with a primary goal of addressing bystanders' concerns around the transparency issues in the owners' home as well as the lack of control of their data. More specifically, these designs focus on detecting nearby smart home devices, informing bystanders of device behaviors, limiting data collection, and controlling personal data process.

Device Awareness.

Eleven of the designs (P1-8, P15-17) had a component to increase awareness of the smart devices in multiple ways. P4's example (Figure 9.1) from the previous section required bystanders to connect to the owners' home Wi-Fi to actively detect the existence of smart home devices inside the house. P14 designed the content of a text message, which included details regarding the types, schedule, and location of the devices. P8 designed an app that could help him realize the existence of a camera

inside the owner's place, and if any devices existed, its operating status:

“You go into each room and the app lets you know if there's a security camera in their home. And going further, maybe it lets you know the location of every camera in the house. But that's dumb because then any burglar can do that. And it lets you know if it's on or off.” (P8)

It is worth noting, however, that P8 also commented on the potential negative consequences of such an app: if he could use it, then any burglar could use it to detect the location of the cameras for a potential break-in. This concern was discussed in three groups, indicating that: (1) if such a design were to be implemented, more advanced security mechanisms should be included; and (2) bystanders' designs need to be critically examined by privacy and security experts as well as practitioners in order avoid security loopholes should any design idea be implemented.

Device behaviors.

Bystanders also design with a desire to know the smart home devices' behaviors to make sure that they are informed. Eight designs (P1-2, P4, P7-8, P11, P17-18) included this feature. For example, one feature in P1's app design focused on improving the transparency of smart home devices, as he explained:

“I talked about an app, so it would be like an app that has access to all IoT devices in your home, so when you talked about the scenario if you walked in your friend's house and there was a camera and you weren't sure. If you walked up with an app, it would tell you all the devices they detected and the hours they are running, if it is on right now, who is using the device right now, the purpose of it so you are aware. This is for everyone around the device to know what it is used for.”
(P1)

P1 was concerned that even if he noticed a smart home device, he would not be aware of the

specific device's details. He would like the app to inform him about different aspects of these devices to increase his awareness, including the device status, purposes, schedule, etc. P7 designed a similar app to monitor device behaviors and inform him of any data collection. Besides, her design also had a piece for providing recommendations regarding how to avoid unwanted data collection. For example, in a case of security camera detection, the app would recommend, *“move to the kitchen if you don't want to be recorded, and don't say secrets.”*

Limit data collection.

Four designs (P3, P6, P7 P14) embedded a piece to limit data collection by smart home devices, particularly bystanders' data. For example, P6 designed an incognito mode for bystanders to alleviate concerns of being tracked:

“I know for a fact that some tech companies track your device's location so by default they track your location. And I think a feature in your smartphone that allowed you to go smart home incognito so if you go to your friends house, and they have a security camera that they don't turn off then I guess you have to get over that privacy concern but the max amount of privacy would be if you have this feature turned on in your phone, your whereabouts won't be in their system. That data doesn't mix with theirs.” (P6)

P6 acknowledged the fact that she was tracked by technology companies, thus she designed an app for her smartphone in which she could set a “smart home incognito mode”. Once activated, this mode could ensure that smart home devices would not track her.

Control data processing. Another feature is to control the processing of bystanders' data, including data sharing, access, storage, and deletion. Four designs (P3, P10, P13, P16) included this component. For example, P10 created a futuristic design which trapped the collected data inside the



Figure 9.2: The signal blocker designed by P10.

house (Figure 9.2):

“This is a signal blocker that stops information from leaving the house, kind of like bubbles around the house basically. Information can still get in. Signal blocker app will go along with it. I could get a notification that says ‘Allow’ or ‘Deny’ information to go out. Potentially it could add other people’s devices to your signal blocker or they could share their system with me, kind of like Google Docs ‘view’, ‘view and edit’ link share.” (P10)

P10s’ design represented an entire smart home system as a signal block shield. In this design, the “bubble” referred to an invisible shield that stopped the information from going out. Bystanders could not control whether to use smart home devices or stop the data collection, but through the app, they could potentially limit the collected data inside their friend’s house. That would make P10 more comfortable.

P13’s design of an app reflected the ideas of data deletion and storage. She explained:

“I would be happy if audio and video were deleted after a couple of days if the person went on and said I want to save this clip for this specific reason. Or if the data is only stored locally on their network or in the mesh network where it is connected between the devices itself or phone or

computer. So it is not actually going outside of your home network and not being sent out anywhere so I know this data is not being shared or used maliciously.” (P13)

P13’s design reflected that her data would be properly handled by either smart home devices or the owners. By deleting the collected data after a certain amount of time and keeping the collected data stored in the owners’ house, this app offered her peace of mind as a bystander since she would know that her data was secure.

It is worth noting that while our participants have different levels of experiences with smart homes, we do not observe notable differences between the designs of those who had more experiences and those with fewer experiences. We suspect this is due to 1) our small number of participants, 2) the qualitative nature of the study, and 3) our scenario-based discussions which reduces the potential influences of participants’ prior experiences with smart home devices. We encourage future work to further dive into this issue using different methodologies (e.g., survey) with a larger sample size.

9.5 Discussion

Few prior research has hinted the needs to study the privacy issues from other stakeholders in smart homes, such as visitors and other family members [129, 229, 235]. Our study attempted to respond to such needs. To the best of our knowledge, our work is the first of its kind that specifically focuses on the understudied bystanders’ privacy in the context of smart homes, aiming to understand their privacy perceptions and desired ways of addressing their privacy concerns. Our results highlight the different aspects of shaping bystanders’ privacy perceptions in smart homes. Besides, our co-design activity results in several design factors that the bystanders desire when designing privacy mechanisms to meet their privacy expectations in smart homes.

This study primarily focuses on smart home bystanders, while prior work primarily focuses on smart home users. Both groups are necessary and important stakeholders in smart homes and their perceptions and desirable design factors provide novel insights for the smart home industry, practitioners, and researchers. In this section, we compare our results to the findings from the literature (primarily the work of Zeng et al. [235], Zheng et al. [237], and Yao et al. [229]) to highlight major differences between the two groups. We choose these prior work because 1) they are among the pioneering work in understanding smart home users' privacy perceptions in the literature, and 2) they are similar to our work in terms of methodology. The goal of the comparison, instead of generating an exhaustive list of similarities and differences between bystanders and users, is to show that owners and bystanders may have different privacy needs. We also take the first attempt to answer the question posed by Yao et al. [229], i.e., "whose privacy should be protected and who should make the decision?"

9.5.1 Comparing Bystanders' and Users' Privacy Perceptions

Our results suggest that bystanders, despite their limited engagement with other people's smart homes, still hold their privacy concerns. We compare our findings of bystanders' privacy perceptions with other results of users' privacy perceptions and summarize three major differences: (1) contextual variations; (2) bystander's data access; and (3) privacy-seeking behaviors.

Contextual variations.

Prior research on smart home users was mostly situated in their own homes. As a result, users' privacy perceptions were tied to their dwellings without too many variations. This means users chose to adopt smart home devices to accomplish certain goals in their home, such as home automation, surveillance, home safety, and remote access [235, 237]. As such, individual user's

privacy perception centered around their specific use cases and was less likely to change once formed. In a sense, there is limited contextual variation from the users' perspective. In comparison, due to the nature of bystanders, they could switch their roles under different social relationships (e.g., family members living in the same home, visitors to another friend's home, or temporary renters of Airbnb homes). Thus, bystanders could face strong contextual variations. Our study showed that bystanders' privacy concerns varied in different scenarios and that their expectations and information needs were also significantly affected by many contextual factors, such as perceived social relationships with the owners and length of stay.

Expectation of bystanders' data access.

Literature has discovered that users' mental models contained several entities that could access their data collected by smart home devices [235, 237]. These entities included device manufacturers, third-party advertisers, government, Internet service providers, etc. No evidence in the literature suggested that smart home users expected bystanders to obtain access to the collected data. However, bystanders from our study expected a certain level of control of either the users' devices or their data collected by these devices. Such expectation was perceived as reasonable by the bystanders as they were captured by the devices, and their privacy could be at risk.

Privacy-seeking behaviors.

Literature has suggested that the majority of users did not seek privacy protection [235]. This was primarily due to users' overwhelming trust towards the device manufacturers. They believed that the manufacturers would provide satisfactory protection to their data and privacy. Even though users realized these privacy issues, many of them were not concerned about these issues and thus did not seek privacy protection [235]. Bystanders in our study, however, were somewhat different. On one

hand, many bystanders claimed that they have taken some actions in their past experiences or in hypothetical scenarios, such as covering the security cameras, talking to the owners, etc. This was because bystanders' trust towards the owners as well as the mediators in some cases could cause privacy concerns, thus they would actively seek their privacy. On the other hand, we also noticed that some bystanders had concerns and would also like to seek protection, however, they did not do so due to the social pressure or awkwardness (e.g., not being in a place to directly talk to the owners), and they felt that there were no other options for privacy protections.

This comparison indicated several mismatches for privacy perceptions between the bystanders and the users/owners in smart homes. These mismatches highlight the fact that bystanders also have privacy needs and desire some controls, as well as the tension between bystanders and users in smart homes. The mismatches also point to potential opportunities for privacy designs to balance the privacy needs of both stakeholders. To answer the first part of the open question posed by Yao et al. [229] (i.e., **“whose privacy should be protected”**), we argue that the privacy of both smart home users and smart home bystanders need to be well protected, since ignoring bystanders' privacy can heighten tensions between bystanders and owners, which would potentially change the social dynamics at the home. In the long run, there can also be various push backs for the adoption and use of these devices across a variety of contexts.

9.5.2 Unpacking Bystanders' Privacy Designs

The co-design activity in our study results in several design factors that bystanders desire to mitigate their privacy concerns. In this section, we would like to take a deeper look at these design factors to illuminate future privacy designs for smart homes.

Cooperative mechanisms.

Rationale. In our study, bystanders' privacy perceptions changed across different scenarios. However, as we noted before, the norms were not always clear in different contexts, especially when considering various social factors embedded in those contexts. Thus, bystanders often desired to cooperate and even negotiate with the owners/users regarding their privacy in a given context. For example, when bystanders did not give consent for their voice to be recorded, they hoped to send a request to the owners and ask for approval to limit the audio recording. As Nissenbaum argues in the theory of Contextual Integrity, it is important to understand the contextual information norm to decide whether one's privacy is breached within a given context [164, 165]. This contextual information norm includes three independent parameters: (1) actors, including the subject, sender, recipient, embody the context; (2) information types, i.e., the types of information that are collected; and (3) transmission principles, i.e., the principle that is either socially acknowledged or required by law [164, 165]. From the bystanders' perspectives, once the *actors* are determined, they would like to explicitly express their privacy preferences to specify *information types*. Considering the *transmission principle* in the smart home, bystanders further sought users' approval for device controls instead of directly controlling the devices at their command. Thus, through communication and potential negotiation between bystanders and users, the cooperative mechanisms were designed to clarify the contextual information norm in smart homes which were not always clear, so that bystanders' privacy expectations could be better achieved.

Comparison among cooperative mechanisms. In the literature, the attempt to cooperatively negotiate privacy management has been made in several contexts to enact the privacy needs of different stakeholders. For example, in the context of photo tagging on Facebook, to protect the

privacy of users who were tagged in photos and, at the same time, still make the photo-sharing and tagging possible for the photo owners, Besmer and Lipford proposed a design which allowed the tagged users to send the photo owner a request and ask that if they were tagged in the photo, the photo should be hidden from certain people [29]. Xu et al. proposed the design principle for privacy-enhancing tools, which stated that users should act as a member of a group and have collective control of their information [227]. In the context of drones, Yao et al. proposed two mechanisms with a cooperative nature: Deletion Request (i.e., drone bystanders could send request to the controllers to delete their footage) and Controller-Bystander App (i.e., drone bystanders could obtain more information of and, if needed, communicate with the drone controllers) [233]. Both these two mechanisms were designed to enact the privacy issues of drone bystanders but at the same time, balance the privacy needs of drone bystanders and the functional needs of drone controllers.

In our research, the cooperative mechanisms were mostly similar to those mechanisms mentioned above. For example, P4's design asking for device control was similar to the Controller-Bystander App in the context of drone privacy designs [233], and P5's design to express his privacy preferences was similar to the photo tagging tools [29]. It is worth noting that, despite being largely embraced by the researchers as well as targeted users [233, 29], most of these cooperative mechanisms have not been implemented in the real world. This is in part due to the feasibility of these mechanisms, which have either technological barriers or policy obstacles, as well as many other potential design issues, e.g., different preferences of drone controllers and bystanders, the amount of user effort required in cooperation, etc. In our study, we also anticipate similar feasibility issues and design questions, such as whether bystanders have different privacy needs than the owners/users and if the bystanders' privacy needs to defeat the purposes of owners' smart home devices. However, the recent large adoption of Amazon Echo in hotel rooms [49] indicates the urgency of effective privacy

mechanisms for bystanders. We advocate that future research should look deep into these issues, examine user-generated ideas more comprehensively in terms of feasibility, usability, and other potential issues (e.g., required user efforts), then propose new cooperative mechanisms to better suit the needs of both smart home users/owners and bystanders.

Bystander-Centric mechanisms.

Designing for awareness. The other set of design factors focused on bystanders only, with a primary focus on increasing awareness and transparency as well as providing some controls to bystanders. On one hand, as notice and choice are well-recognized privacy principles, increasing bystanders' awareness of nearby devices and device behaviors reduce their uncertainty and alleviate their privacy concerns to a certain degree. On the other hand, what level of awareness should be provided to bystanders and how to do so remain open questions and require further investigation, especially when considering if such awareness tools fall into the wrong hands and cause unpleasant safety incidents. For example, P8's design provided awareness to the bystanders by showing them the location of every security camera in the house, but at the same time, if this app was accessed by burglars, then the safety of the house might be compromised. Future research is needed to provide better and more comprehensive solutions.

Designing for control. In terms of controls, we found it intriguing that bystanders designed for having active controls on other people's property, even though the controls were towards bystanders' data rather than the users' devices. This was an indicator to us that bystanders also expected to have agency in other people's homes. However, it is worth noting that, while providing some agency to bystanders could potentially help them to be able to better control their data from being collected and shared, this could conflict with the owners/users' purpose of using the devices or even invade

the users' privacy. For example, P10's design of the signal blocker limited the data being shared with entities outside of the house but also potentially interfere with the owner's data as all data was collected by the same devices.

Comparisons bystanders' and users' privacy designs.

In the end, we also make a comparison between bystanders' design factors in our research and the owners/users' design factors in Yao et al.'s study [229]. We found that although the design features (e.g., information device behaviors, local data storage) under the bystander-centric mechanisms category in our study were similar to the ones in their study, the design features under the cooperative mechanisms category were unique. As discussed before, the cooperative mechanisms aimed to enhance the communication between bystanders and owners/users so that they could negotiate and hopefully fulfill their individual privacy needs. To answer the second part of the question posed by Yao et al. [229] (i.e., **“who should make the decision”**), we argue that such decisions should be made cooperatively between end-users and bystanders with a consideration of the specific context they are in.

9.5.3 Design Implications

Privacy in smart homes is an important topic in the CSCW community. Many studies have looked at the privacy concerns and perceptions of the owners/users regarding smart homes as a whole [235, 237] and individual smart home devices [129, 149], as well as users' desired ways of mitigating their privacy concerns [229]. Nevertheless, all of these studies hinted at potentially different privacy perceptions from the bystanders' perspective, yet none of these studies have explored the differences. Our study provides rich and novel empirical evidence demonstrating that 1) bystanders in smart homes have various privacy concerns which are further influenced by several

contextual factors, and 2) bystanders desire some forms of privacy controls. Our discussion also dives into the mismatches between the perceptions of bystanders and owners/users, highlighting the tension between these two stakeholders and the needs for enact privacy issues cooperatively.

Based on our study results, we make the following concrete design suggestions for future privacy-enhancing mechanisms in smart homes.

Transparency. In various smart home application contexts, the existence of smart home devices, as well as their behaviors, should be more transparent to bystanders. For example, in the temporary residency scenario, one concrete example is that the owners should proactively provide information about smart home devices, such as their location, purposes, whether data is collected and stored, etc. In the case that the owners do not know some of this information themselves, the device manufacturers should provide this information along with the device, e.g., in a poster with a QR code. The owners can potentially place physical signs alongside the smart home devices in the apartment so that the tenants are more aware of such devices, and provide the tenants with options to opt-out from these devices being used. In the case where mediators exist (e.g., Airbnb in the temporary residency scenario), the mediators can also assure by notifying bystanders if the owners delete their recordings.

Expressing preferences. When designing privacy mechanisms in smart homes, it is critical to consider the needs of both owners/users and bystanders collaboratively to design for both groups. For example, smart home device platforms (e.g., Samsung SmartThings) can potentially create apps for both owners/users and bystanders so that the latter can express their preferences and potentially communicate with owners/users.

Different modes. Smart home devices owner/users and the devices themselves should also be proactive in considering and protecting bystanders' privacy in various ways. One concrete solution

is for smart home devices to have different modes. The devices will be fully functional in user mode, but in bystander mode, the devices' functions can be selectively disabled. For example, the voice assistants will stop recording if a different voice is detected, indicating the possibility of friends visiting; the security camera will only record footage of a designated area in an Airbnb apartment (e.g., the hallway), ensuring the safety of the apartment while protecting tenants' privacy.

It is worth noting that the design suggestions above focus primarily on bystanders. However, one open question to ask for future research is, *how will these mechanisms benefit the owners/users and other stakeholders (e.g., mediators) and whether they will accept these mechanisms?* Grudin's prior work on groupware argued that people who used groupware would not benefit the most, and in some cases, technology that was designed to support one group of people might negatively impact another [101]. Thus, ensuring these privacy mechanisms also benefits owners and other stakeholders are crucial for their adoption.

9.5.4 Limitations and Future Work

Our study is the first to explicitly focus on understanding how bystanders perceive their privacy in smart homes and their desired ways of mitigating their concerns. As such, our exploratory approach has a few limitations.

First, our three scenarios in the study were by no means exhaustive in terms of different application contexts of smart homes and the types of bystanders. Other example scenarios and bystanders could include, for example, a UPS delivery man being caught by a smart doorbell every day, or children being caught by security cameras in their neighbors' house. Future research can either investigate a more diverse set of scenarios or come up with different contextual factors so that participants can assemble their scenarios.

Second, given the purpose and the exploratory nature of the study, our focus group and the co-design activity only included a bystanders' perspective. The co-design activity, although insightful, was also limited by the duration of the study. Future research can consider running extended participatory design workshops with owners/users, bystanders and other potential stakeholders (e.g., developers) to surface the tension and gain a more holistic understanding of their perceptions and desired designs.

Third, given the scope of this study, we did not critically evaluate bystanders' designs in terms of their usability, feasibility, and potential consequences. Future research can focus more on critical evaluation of the designs emerged from our study as well as prior studies [229] and come up with new designs to better fulfill the needs of different stakeholders.

9.6 Conclusion

Prior literature in smart home privacy has focused on end-users, leaving other potential stakeholders, such as bystanders, understudied. In this study, we focus on smart home bystanders, i.e., people who are not the owners/users of these devices but can potentially be involved in the use of such devices. We aim to understand bystanders' privacy perceptions in various contexts and their desired ways to mitigate their privacy concerns. Through six focus groups with 18 participants, our study results in a number of contextual factors that can potentially influence bystanders' privacy perceptions of smart homes. In addition, we also identify a set of design factors that the bystanders consider in their desirable privacy mechanisms to address their concerns. These factors can further be categorized into two types, i.e., cooperative mechanisms and bystander-centric mechanisms. Our research highlights bystanders' needs for privacy and some means of controls, the tension between smart homeowners/users and bystanders, as well as how cooperative mechanisms can be used to better

support and balance the needs of both groups. We propose several concrete design suggestions based on our results, i.e., having both owners/user-oriented apps and bystander-oriented apps in smart home platforms.

9.7 Acknowledgement

We thank our participants for their experiences and insights. We are also very grateful to Bryan Semaan, Nata Barbosa, Smirity Kaushik for their assistance as well as the anonymous reviewers for their thoughtful feedback. This work was supported in part by NSF Grant CNS #1464347.

Part V
Discussion

Chapter 10

Summary of Results

To summarize, this dissertation explores the privacy issues and privacy design space from a multi-stakeholder perspective following the steps below:

I chose to start with a single user view in traditional privacy research in the OBA context. Following a conventional privacy research paradigm, I aimed to investigate users' understandings of how OBA works to identify any privacy risks in their beliefs. Based on the results, I proposed and built prototypes of an information-based web tracking blocker that challenged the status quo of tracker-based web tracking blockers. Then adopting a speculative design approach, I explored how a multi-stakeholder perspective changed our understandings of the privacy issues and the designs of web tracking blockers in the OBA context. The results indicated that when considering other stakeholders in the OBA eco-system, users' privacy issues could have a significant impact on other stakeholders, such as the advertisers, the technology designers, and the policymakers. Such an impact might further influence the privacy designs since different stakeholders might choose to implement the privacy tools in a variety of ways, which would diminish the purpose of the original designs. This initial exploration demonstrated that a multi-stakeholder perspective could contribute a different view of looking at privacy issues and privacy designs, which are currently understudied in the privacy literature.

To build better understandings around the multi-stakeholder privacy perspective, I then chose to study an emerging domain that contained a multi-stakeholder setting, the drones. I conducted interview studies with drone controllers and bystanders and discovered the mismatches between their understandings of the privacy issues related to drones. The results demonstrated the mismatches in several aspects (e.g., whether privacy was an issue, whether drone controllers have malicious purposes of flying drones), as well as the tension between drone controllers and bystanders. Building on the results, we proposed several novel privacy designs for drones to better balance the privacy needs of both parties.

Finally, I moved on to smart homes. Smart homes represent a relatively complex social environment that consists of various social relationships, power dynamics, and stakeholders. It is essential to understand the privacy perceptions, needs, and expectations of multiple stakeholders in smart homes. Furthermore, since existing smart home privacy designs are all proposed by either researchers or technology designers, whether they can address the needs of the actual users remains unknown. As such, I chose to conduct co-design focus group studies with smart home users and bystanders to dive into their privacy perceptions, then worked with them and designed the privacy mechanisms they desire. By synthesizing the results from the co-design sessions, I aim to understand how a multi-stakeholder perspective can inform privacy designs in the context of smart homes.

10.1 Synthesis of Results

After examining the multi-stakeholder situations in three domains individually, I synthesize the results across the three domains and make sense of the commonalities among all contexts. Table 10.1 presents the key findings from these studies as well as a high-level synthesis. In this section, I will dive into the commonalities across all three domains and discuss how privacy designs can embrace

the commonalities in various multi-stakeholder environments.

10.1.1 Commonalities among users in three domains

From the users' perspective, they generally overlook the privacy of other stakeholders. Besides, in most cases, users seldom consider how their use of technologies or their privacy decisions may have impacted other stakeholders. In some contexts, users also need to provide sufficient information regarding what data their devices collect and how they will use the data to keep other stakeholders informed.

From other stakeholders' perspectives, one primary observation is that stakeholders other than the primary users also have privacy needs, which should be carefully and thoroughly examined and considered. Besides, stakeholders other than primary users should proactively express their privacy needs and preferences. Finally, users may be impacted by the privacy decisions of different stakeholders, which is also generally ignored.

Table 10.1 points to several questions that need to be considered in various multi-stakeholder contexts.

- Privacy designs should allow different stakeholders to express their preferences. How do we enable various stakeholders to communicate their privacy needs and preferences? Which stakeholders should be taken into consideration?
- Privacy designs should balance the needs of different stakeholders. In many cases, that will require communication and negotiation. How to provide communication channels that allow various stakeholders to negotiate their needs or preferences? What is the cost of communication and negotiation? Is it always worth the effort to communicate and negotiate? How to reduce the effort which is needed to negotiate? Should the communication be

enforced?

- Privacy designs should also educate different stakeholders on how their behaviors and privacy decision can be mutually influential. How to measure the influence? How to educate different stakeholders?

	Stakeholders	Privacy Issues	Factors	Implications
Case OBA	1: Users, social, technical, legal, and economic stakeholders	<ul style="list-style-type: none"> - Users would like to know what information is collected rather than who collects the information - Users have the needs to protect their privacy, but did not consider the other stakeholders in the larger space 	<ul style="list-style-type: none"> - The website they are visiting - What entities are involved in the data collection - What information is collected 	<ul style="list-style-type: none"> - Privacy tools to balance the needs of various stakeholders - Users need to understand how their privacy decisions impact the larger space and other stakeholders
Case Drones	2: Drone bystanders, drone controllers	<ul style="list-style-type: none"> - Drone bystanders have various privacy concerns - Drone controllers do not think that bystanders have privacy risks, but bystanders believe that they have the needs to protect themselves 	<ul style="list-style-type: none"> - Bystanders' unawareness of the existence of drones or their data collection - Bystanders' relationship with the controllers; the data collection purpose - The location of data collection which relates to the expectation of the bystanders 	<ul style="list-style-type: none"> - Privacy tools to enhance the communication between drone controllers and bystanders - Mechanisms that maintain the drones' functionalities while protecting bystanders' privacy - Privacy mechanisms should also protect drone controllers' privacy.
Case Smart homes	3: Smart home users, smart home bystanders	<ul style="list-style-type: none"> - Smart home users have privacy needs - Smart home bystanders have privacy needs - Smart home users do not think bystanders have too many issues, but bystanders have the needs to protect themselves - Most bystanders do not take actions to protect themselves due to many social reasons 	<ul style="list-style-type: none"> - Relationships between the users and the bystander - Location - Types of data - Length of stay - Levels of difficulties to discover privacy issues 	<ul style="list-style-type: none"> - Cooperative mechanisms are promising - Users should consider bystanders' privacy preferences

Table 10.1: Synthesis of results across all three study domains

Chapter 11

Implications

11.1 Design Implications

Building on top of the results and the synthesis above, I will discuss the following design implications for future privacy designs for various socio-technical systems when considering multiple stakeholders.

Cooperative approach is promising. Here, the cooperative approach does not only refer to the cooperative mechanisms in the study of smart home bystanders. Instead, it also relates to designs that consider multiple stakeholders rather than only considering the end-users. The purpose of the cooperative approach is to provide channels and platforms so that different stakeholders can negotiate their privacy needs, expectations, and controls cooperatively. Besides, in a cooperative approach, default privacy settings should be built to reduce the time and effort needed for effective negotiation.

Relationship-based privacy designs. A critical aspect of multi-stakeholder privacy management is to understand the contextual privacy needs of different stakeholders. In this dissertation, one common observation across all three cases is that our participants made their privacy decisions based on their perceived relationship with the other stakeholders (e.g., in the smart home context, bystanders' privacy expectations rely on who are the users/owners). Then, an essential aspect of

managing one's privacy is relationship management. As such, privacy designs should consider how to better support people's needs for relationship management. One concrete example is to set up context profiles and stakeholder profiles in privacy designs. Each profile will have different privacy preferences and can be applied based on the contexts and the stakeholders involved. One can also set up relationship profiles. Each profile will be associated with a specific kind of relationship and will be activated based on the current relationship with other stakeholders. This approach is similar to the role-based access control in computer security [188, 85].

User education. This dissertation shows that when making privacy decisions, most of the participants did not consider how their decision may impact other stakeholders. In the future, privacy designs should educate their users to better understand the relationship among different stakeholders and how they can affect other stakeholders. A concrete example, following the line of research on privacy nudges [210, 12, 8, 236], is that when a drone controller wants to publish the drone footage online through an app, the app can pre-process the recording and notify the user, "your video contains three bystanders whose faces are identifiable. Please make sure you have obtained their consent to publish this video."

Regulations to support the multi-stakeholder effort. The US NTIA has published the multi-stakeholder process for a variety of technologies and use cases [4, 3, 2]. However, regulations are needed to support the adoption of such multi-stakeholder processes better. A concrete example is that the law can enforce any smart home manufacturer to develop a publicly accessible privacy infrastructure that allows people to discover nearby devices and to enter their privacy preferences. When a nearby device is collecting data that violates their preferences, the privacy preferences will be sent to the owner of the device for communication and negotiation.

11.2 Methodological Contribution

In this dissertation, I employed a variety of methodologies to uncover the privacy perceptions, expectations, and designs of different stakeholders in various socio-technical systems. In this section, I highlight how the methodological approach in this dissertation is different from what has been commonly used in usable privacy literature as well as how it contributes to general privacy research.

Throughout the dissertation, I used a combination of qualitative, quantitative, and design methods to directly engage different stakeholders in exploring their privacy perceptions and expected privacy designs. The outcome has been very beneficial.

In the online behavioral advertising chapter, the interview study provided insights into users' folk models of their privacy perceptions, concerns, and preferences. The results pointed to the potential for a privacy-enhancing tool. Then, with the complement of a speculative design activity, I further explored that if a privacy tool was to be implemented, what would be the potential implications on various aspects of our society. These implications further implied how privacy-enhancing tools should be designed today.

In the drone chapter, the interview studies uncovered the privacy perceptions and expectations of both drone controllers and bystanders and provided insights into how privacy-enhancing tools for drones could be designed. The following survey study further tested several design ideas with both stakeholders. The results pointed to two mechanisms that were well received by both stakeholders. The results also provided implications for future research and design.

In the smart home chapter, the co-design method has not been widely adopted in the privacy research community. The smart home studies in this dissertation were among the first ones to use

co-design to illustrate users' privacy expectations and desired designs. By conducting co-design workshops with both smart home users and bystanders, I uncovered many new aspects of the privacy expectations of both populations as well as key design factors. These design factors shed light on the design of future privacy-enhancing technologies for smart homes.

The overall approach in this dissertation not only has its unique values and helps uncover the perspectives of different stakeholders but also differs from other approaches in the privacy literature that also consider multiple stakeholders. For example, multiple stakeholder consideration was also a critical component in a Privacy Impact Assessment (PIA) [224]. However, PIA is a procedure of expert reviews. When a system turns into a functional prototype or has been developed, privacy experts, using standard questions from PIA, evaluate the potential privacy risks caused by the system to help system designers and developers understand the privacy landscape. In this process, no actual users are involved. As a result of using PIA, privacy experts may identify privacy risks that real users would not care about, and omit other aspects which may not be privacy risks but real users do care about. Future research in the privacy space, especially when user input is needed, should consider a combination of various methods. In particular, methods such as speculative design and co-designs are not commonly used in privacy research. However, proper execution of these methods can yield valuable insights and complement other more common research methods such as interviews and surveys.

11.3 Theoretical Contribution

The theory of Contextual Integrity (CI) argues five parameters to define the privacy norms in a specific context: 1) data subject, 2) data sender, 3) data recipient, 4) information type, and 5) transmission principle [164, 165, 166]. Through the lens of the CI theory, prior research has been

focusing the contextual privacy perceptions, factors that influence these perceptions, and designing and implement new privacy mechanisms to protect people's privacy [144, 135, 111, 24, 103, 104, 78, 50].

In this dissertation, building on the CI theory, I explored the privacy issues in three different domains through a multi-stakeholder perspective. The CI theory focuses on the contextual norms of information flow and appropriateness. These norms not only provide a concrete way to understand and define the privacy issues within the studied domains (especially when considering the potential interactions among different stakeholders) but also point out the direction of how privacy mechanisms should be designed in various contexts. As such, the CI theory forms the theoretical foundation of this dissertation.

Moreover, this dissertation has also discovered some incidents that cannot be clearly defined or explained using the CI theory. Drawing from these incidents, I discussed how the results of this work could complement the CI theory. Thus, this dissertation also makes theoretical contributions.

To better guide the following discussion, I ask the primary question from a theoretical perspective, **dose the theory of Contextual Integrity still hold itself when considering multiple stakeholders?**

One thing worth noting is that the majority of prior research using the CI theory has focused on the privacy of a single user. In the scope of this dissertation, I investigated the privacy issues of not only the users but also other stakeholders that might be affected across contexts. Building on top of the prior research, this dissertation also explores whether the CI theory still holds in the multi-stakeholder environment in today's complex socio-technical systems. In the following section, I critically examine the findings of this dissertation. I argue that while in most contexts, one can use the CI theory to define the privacy norms, it is difficult to do so in some other contexts. I will

provide examples to elaborate in detail.

11.3.1 When is CI theory promising?

This dissertation builds on top of the line of prior research on the CI theory. It explores to what extent the CI theory can guide how we approach privacy issues in today's complex socio-technical systems when adding the multi-stakeholder perspective. Many examples in prior studies aligned with the CI theory. Table 2 below summarizes several examples from all three contexts and explains how they connect to the CI theory. Each example in Table 11.1 represents a specific norm in the context. For example, drone bystanders believed that the drone controllers should not take their photos without their consent. In that study context, since the bystanders did not give consent to the controllers, they considered their privacy to be breached. In the smart home bystander example, bystanders believed that Amazon Echo should obtain their consent in addition to the users' content to capture their voice data. In that study context, one of our participants indicated that since he was often not aware of being recorded (as a bystander), he considered his privacy being compromised.

Example	Data Subject	Data Sender	Data Recipient	Information Type	Transmission Principle	Privacy Violation
OBA	First-Party pull model of P16	Amazon ads server	Facebook	Product information	First Party ads server will hold the data confidential	Yes
	Third-Party model of P4	Amazon	Facebook	Personal preferences	Third-party will hold the data confidential	Yes
Drones	Drone bystanders' concern of their photos being taken	Drone	Drone controllers	Visual and audio information	Drone controller will not distribute the photos/videos	Yes
	Drone controllers taking photos	Drones	Drone controllers	Visual and audio information	Drone controllers will not distribute the photos/videos	Yes
Smart Homes	Smart home users' voice data collected by Amazon Echo	Amazon Echo	Amazon	Audio information	Amazon will hold the data confidential	Yes
	Smart home bystander' voice data collected by Amazon Echo	Amazon Echo	Amazon	Audio information	Amazon will hold the data confidential	Yes

Table 11.1: Examples through the lens of the CI theory

11.3.2 When is CI theory challenging?

The CI theory is not predictive. Instead, it uses five parameters to define a contextual norm, which is then used to determine whether something is a privacy violation. However, through the analysis, I found several cases in the study results in which privacy norms were unclear or could not be easily defined based on the existing parameters in the CI theory. There are primarily two reasons. First, in some cases, the additional stakeholders can complicate one or more of the five parameters in the CI theory and makes it challenging to define norms. Second, in some other cases, the perceived norms are affected by parameters other than the five parameters in the CI theory. Below, I will discuss several examples from these two aspects, then make suggestions on how to complement the existing CI theory. Finally, I will call out areas that need further research to elicit the privacy norms in a multi-stakeholder environment better using the CI theory.

Multi-stakeholder perspective complicates the privacy norm identification.

To discover the informational norm in a context, it is necessary to identify the five parameters. For example, when a patient (data subject) finds a new doctor, the prior doctor (data sender) can transfer the patient's complete medical record (information type) to the new doctor (data recipient) upon the patient's authorization (transmission principle). As such, when the five parameters are clear, the privacy norm will then be defined.

However, the multi-stakeholder consideration adds another layer of complexity to these five parameters and makes it more challenging to identify norms. Such complexity exists in many contexts. For example, when an Amazon Echo collects the voice message of the users who have consented Amazon to collect their data, it can also collect voice messages by other people in the background, and these people are likely to have not consented their data being collected. In this

example, the data subject is not clear, since both the users and the people in the background are subject to the data collection. The transmission principle is not clear either, since although the users may have consented the data collection, the people in the background may be unaware of their data being collected. Finally, the information type, in this case, is also challenging to identify, especially in the case in which the users and the bystanders interact and have a conversation with each other. It is not clear what and whose data is collected and who should consent to the data collection. After all, in this example, identifying the privacy norms will be very challenging.

In another example, when a drone is flying above a park and taking photos of the park, it will also be likely to capture many people who are in the park now. In this case, the data subject is not clear since it not only involves the park landscape but also all the people in the park. Identifying individuals whose data is collected will be very challenging. The transmission principle is not clear either since it is difficult to either notify the individuals who are captured by the drone or have them consent to the data collection. Even if all individuals are informed, there are still variances in deciding whether they are in a completely public space or a private area of a public place. As a result, the privacy norms in the drone case remain unclear.

These examples indicate that, as we start to think about the multiple stakeholders other than the end-users, identifying privacy norms becomes challenging and more complicated because of the mix of the data subject, opaque information types, and vague transmission principles. If we think more broadly, a related question may arise: what about situations where there is not an accepted norm yet? What changes or addition can be made to the CI theory so that it can further advance the understandings and identifications of privacy norms through the multi-stakeholder lens? I will discuss this direction in a later section.

Considering the Multi-stakeholder Perspective in the CI theory.

The results from this dissertation suggest many cases in which contextual informational norms are challenging to identify under the CI theory. This is primarily due to the various needs, perceptions, and expectations of different stakeholders in a given context. In the next section, I will present three examples of such cases and discuss how the results from this dissertation can help identify the contextual informational norms.

The first example is related to the privacy issues in drones. In this example, when drone controllers take photos in a park using their drones, they may capture bystanders' images. However, our study suggests that the contextual informational norm is not clear. One of the determining factors is whether the park is private or public places. To some bystanders, although they consider the park as a public space, they still consider the proximity of themselves as a private or semi-private space since people who are within the vicinity can hear their talking or see their phone screen. As such, they expect to be notified and give consent before their data is collected. However, drone controllers tend to define the park as a public space based on the legal definition. Thus they believe that they have the right to fly a drone in the park and take pictures, and bystanders should not expect privacy. Such differences, caused by controllers' and bystanders' conflicting perceptions of the nature of the park, are not captured by the five parameters in the CI theory.

The second example is related to tenants staying at an Airbnb apartment. In this example, one of the critical factors that influence our participants' privacy perception is their length of stay. In this case, if the tenants only stay in the Airbnb apartment for a short period (e.g., one day), they perceive the use of a security camera acceptable as long as it is not in the living room because only limited amount of data can be collected in one day. As such, their privacy is not compromised. However,

when the tenants have to stay in the Airbnb apartment for a long time (e.g., seven days), it is no longer acceptable to have a security camera in the apartment, because over a long period, additional information about them, such as their schedule, preferences, or activities, can potentially be inferred or directly collected. As a result, in this case, the informational norm is not clear either.

The third example is related to some guests visiting their friend's smart home. The home contains a security camera systems with multiple cameras across the house for security purposes. Our study has suggested that guests' privacy expectations were determined mainly by their perceived relationship with the homeowners. If the guests have a close relationship with the homeowners, then the use of a security camera system is acceptable. Otherwise, the guests expect to be notified before they arrive. It is also worth noting that in our study, even in some cases, the guests felt somewhat uncomfortable regarding the data collection by the security camera, they did not think it is a privacy violation because of the perceived close relationship. In this case, the information norm is also significantly impacted by the perceived closeness of the relationship, which is not covered by the CI theory either.

In summary, in the above examples, the existing five parameters in the CI theory may not be straightforward or even sufficient to determine whether privacy has been compromised, because it did not capture the interests and preferences of the affected stakeholders (Example 1 – drone controllers and bystanders), individual differences (Example 2 – tenants and hosts), and contextual and societal needs and values (Example 3 – guests and owners). As such, the decision of whether privacy is compromised in these contexts is more complicated because of these additional factors.

As such, identifying the contextual informational norms in certain contexts, especially those that contain multiple stakeholders, warrants further research. To illustrate one example direction for future research, consider a **contextual alteration** factor. To make it clear, the goal of the contextual

alteration factor is not necessarily adding a new parameter to the CI theory. Instead, it aims to provide an angle for researchers to capture things that might influence the contextual informational norms but can be easily overlooked, such as the interests, needs, and preferences of any affected stakeholders. These factors may alternate people's privacy expectations in the context. Below, I will elaborate on the factor of *contextual alteration* with the above examples and re-examine how it helps to identify the contextual informational norms.

In the first example of drone study, when a drone is flying in a park and taking photos and video of the surrounding environment, the contextual informational norm can be defined by five parameters, i.e., data subject - drone bystanders, data sender - drones, data recipient - drone controllers, information type - visual and audio information, and transmission principle - drone controllers will not distribute the recorded data. The *contextual alteration* factor (Table 11.2), in this case, is whether the park is perceived as a private place or a public. In our study [228], some drone bystanders believed that the area in their proximity should be considered as private places even though they were in a public park. In this case, the contextual informational form has changed. It needs to be re-defined, particularly in terms of the transmission principle, since once the data collection happens at a perceived private space, then the drone bystander will expect to give consent before the controllers can record them.

In the second example of the use of security cameras in an Airbnb apartment, when the host installs a security camera in the apartment, the contextual informational norm is defined by the five parameters, i.e., data subject - tenants, data sender - the security camera, data recipient - hosts, information type - visual and audio information, and transmission principle - hosts notify the tenants about the usage of a security camera and its purpose of ensuring apartment safety. The *contextual alteration* factor, in this case, is the tenants' length of stay. In our study, our participants argued that

	Condition 1	Condition 2
Data Subject	Drone bystanders	
Data Sender	Drones	
Data Recipient	Drone controllers	
Information Type	Visual and audio information	
Transmission Principle	Drone controllers will not distribute the photos/videos	
Contextual alteration	Perceived as private space	Perceived as public space
How does the informational norm change?	Bystanders expect to be notified	Bystanders should not expect privacy

Table 11.2: An example from the drone study

	Condition 1	Condition 2
Data Subject	Tenants	
Data Sender	Security Camera	
Data Recipient	Hosts	
Information Type	Visual and audio information	
Transmission Principle	Hosts notify the tenants about the usage of security camera	
Contextual alteration	A short period of time (e.g., 1 days)	A long period of time (e.g., 7 days)
How does the informational norm change?	Unchanged	Tenants expect to not to be recorded

Table 11.3: An example from the airbnb scenario

if they only stayed in the apartment for a short time (e.g., one day), they would accept the usage of a security camera for security purposes. As a result, the contextual informational norm remains unchanged. However, if they stayed in the apartment for a long time (e.g., seven days), then they would consider the usage of a security camera as a privacy breach. The contextual informational norm has changed, as our participants believed that additional information types would also be collected, such as their schedule and activities.

In the third example of guests visiting their friend’s house, when the homeowner installs a

	Condition 1	Condition 2
Data subject	Guests	
Data sender	A security camera system	
Data recipient	Homeowners	
Information type	Visual and audio information	
Transmission	The homeowner notifies the guests that the data is only used for home safety purpose	
Contextual alteration	Close relationship	Not close relationship
How does the informational norm change?	Unchanged	Guests expect to be notified

Table 11.4: An example from the friends visiting scenario

security camera system, the contextual informational norm is defined by the five parameters, i.e., data subject - guests, data sender - a security camera system, data recipient - homeowner, information type - visual and audio information, and transmission principle - the homeowner notifies the guests that the data is only used for home safety purpose. In this case, one *contextual alteration* factor is the perceived social relationship closeness between the homeowner and the guests. Our participants have mentioned that, if they have a close relationship with the homeowner, they would have no objection towards the data collection by the security camera system. The contextual informational norm remains unchanged. However, if they do not have a close relationship, then they would feel their privacy being compromised. In this case, the contextual information norm had also changed, as our participants expected that prior consent should be obtained before they entered the house or the security camera systems should be turned off during their stay.

In these examples, there are contextual informational norms defined by the existing five parameters in the CI theory. The notion of multiple stakeholder consideration does not necessarily change any of the five parameters at first. However, it may introduce new variables that are currently not captured by these parameters. The *contextual alteration* factor captures these new variables, which

may change the privacy expectations of different stakeholders and, eventually, the contextual privacy norms as well. As a result, privacy violation determination needs to be re-examined. As such, it is worth noting that the purpose of the *contextual alteration* factor is not to defeat the original CI theory. Instead, it forces us to re-examine the contextual informational norms through the additional multi-stakeholder lens and carefully consider the potential changes of contextual informational norms caused by the various needs and interests of different stakeholders.

11.3.3 Research Agenda

Based on the above discussion, I propose a research agenda to 1) investigate how to define contextual informational norms in complex socio-technical systems when considering multiple stakeholders; and 2) explore whether contextual alteration factors also exist in other contexts and if so, how they influence the contextual informational norm.

- Capture information practices through some diagnostic tools. These tools shall be used to identify the information flow among different entities in a computing system. One possible way is through network traffic analysis to understand the starting and ending points of a particular data flow.
- Communicate information practices meaningfully with different stakeholders. Research is needed to investigate how to provide notice, choice, and consent to different stakeholders. As different stakeholders have divergent needs for various types of information in privacy notice and consent, it is very critical to identify what information is useful for each stakeholder. Moreover, as each stakeholder may have their preferred way of accessing notices, one thing that needs special attention is to ensure easy access for all stakeholders.

- Learn about the privacy expectations of different stakeholders. Research is needed to understand how various stakeholders perceive privacy and, thus, have different privacy expectations. More importantly, it is critical to identify potential mismatches or conflicts in their expectations. Research is then needed to carefully and thoroughly consider the conflicting needs and seek solutions to resolve the conflicts. One concrete example inspired by our study results is the cooperative mechanism, in which different stakeholders can negotiate their privacy needs cooperatively.
- Detect discrepancies between the actual information flows and the expected information flows. From an empirical perspective, research is needed to investigate how different stakeholders perceive the information flow in various computing systems. From a technical perspective, tools are needed to capture the real information flow in various computing systems. If discrepancies exist, notifications should be sent to different stakeholders in proper ways.
- Define meaningful transparency for different stakeholders. Research is needed to understand what transparency means to different stakeholders thoroughly. This is very important as many privacy designs, policies, and laws center around providing transparency. Research is also needed to understand the potential consequences of offering transparency, since it is possible that improving transparency for one stakeholder may violate the privacy of others (e.g., the controller-bystander app in the drone study). Thus, balancing the needs for privacy and transparency is very crucial.
- Decide whether the contextual alteration is needed in other contexts. Our prior studies have demonstrated the potential of contextual alteration in the context of drones and smart homes to capture the values and perceptions caused by individual differences. Research is needed to

explore whether such a phenomenon exists in other contexts. If so, research is also needed to discover ways to capture contextual alteration in different contexts and to investigate how contextual alteration impact the informational norms.

11.4 Limitations

This dissertation has several limitations.

First, the dissertation primarily focuses on two types of immediate stakeholders, i.e., the users and bystanders, in three socio-technical systems (i.e., online behavioral advertising, drones, smart homes). Although there is one section in the dissertation that explores the implications on the external stakeholders, those external stakeholders were not directly involved in the actual empirical studies. This is primarily due to the scope of this dissertation. Future work can expand on the selection of stakeholders and involve other potential stakeholders in the empirical studies.

Second, although various design implications have been drawn and key design factors have been discovered, those are not users' actual behaviors, because no actual systems were implemented. Future research can expand on the scope of this work by building either functional prototypes or real systems, then launching the systems through field study to collect users' real-world behaviors in order to evaluate the designs.

Third, this dissertation only focuses on three specific types of technologies. As a result, the results might not generalize in other technological contexts. Future research should be conducted to examine other technology domains to further test the results from this dissertation.

Chapter 12

Conclusion

In this dissertation, I first identified a gap in privacy research literature. Most privacy research has been focusing on the privacy of end-users, and few have considered the privacy of multiple stakeholders and the privacy implications caused by the interactions and potential confrontations among different stakeholders. As a result, how the multi-stakeholder perspective changes our understandings of privacy and the designs of privacy mechanisms remains an understudied area.

Through three cases in different domains (i.e., online behavior advertising, drones, and smart homes), in this dissertation, I examined how different stakeholders in each domain perceive privacy and form their expectations differently, and how such variances inform privacy designs. For example, in the context of online behavioral advertising, we learned that privacy designs should consider the needs of various societal stakeholders to avoid potential resistance of adoption. In the context of drones, we surveyed several privacy mechanisms. We found that both stakeholders better perceived those mechanisms that considered the privacy needs of both drone controllers and drone bystanders. In the context of smart homes, we learned that cooperative mechanisms are promising in addressing the social confrontations and conflicting privacy needs among different stakeholders in smart homes.

Drawing from a few examples in the results, in which identifying contextual information norms can be challenging using the theory of Contextual, I argued that the multi-stakeholder perspective

had introduced new variances to the understandings of privacy. I advocate that when applying the CI theory in an multi-stakeholder environment, it is important consider contextual alteration factors. These factors may be hidden and not obvious but may change the contextual informational norms.

Finally, I discussed design implications and a research agenda moving forward.

Appendix A

Yang Wang*, Huichuan Xia, Yaxing Yao, and Yun Huang

Flying Eyes and Hidden Controllers: A Qualitative Study of People’s Privacy Perceptions of Civilian Drones in The US

Abstract: Drones are unmanned aircraft controlled remotely or operated autonomously. While the extant literature suggests that drones can in principle invade people’s privacy, little is known about how people actually think about drones. Drawing from a series of in-depth interviews conducted in the United States, we provide a novel and rich account of people’s privacy perceptions of drones for civilian uses both in general and under specific usage scenarios. Our informants raised both physical and information privacy issues against government, organization and individual use of drones. Informants’ reasoning about the acceptance of drone use was in part based on whether the drone is operating in a public or private space. However, our informants differed significantly in their definitions of public and private spaces. While our informants’ privacy concerns such as surveillance, data collection and sharing have been raised for other tracking technologies such as camera phones and closed-circuit television (CCTV), our interviews highlight two heightened issues of drones: (1) powerful yet inconspicuous data collection, (2) hidden and inaccessible drone controllers. These two aspects of drones render some of people’s existing privacy practices futile (e.g., notice recording and ask controllers to stop or delete the recording). Some informants demanded notifications of drones near them and expected drone controllers asking for their explicit permissions before recording. We discuss implications for future privacy-enhancing drone designs.

Keywords: Drone; UAV; tracking; perceptions; privacy.

DOI 10.1515/popets-2016-0022

Received 2015-11-30; revised 2016-03-01; accepted 2016-03-02.

1 Introduction

1984. Small flying machines rove around Airstrip One where Winston Smith lives, and peek through the windows [49].

2015. A small flying machine crashed in the White House where the US President Barack Obama lives [33].

*Corresponding Author: **Yang Wang:** SALT Lab, School of Information Studies, Syracuse University, E-mail: ywang@syr.edu

Huichuan Xia: Syracuse University, E-mail: hxia@syr.edu

Yaxing Yao: Syracuse University, E-mail: yyao08@syr.edu

Yun Huang: Syracuse University, E-mail: yhuang@syr.edu

The flying machine that George Orwell imagined in his classic novel 1984 and that crashed in the White House lawn is known as *drones*. The Merriam-Webster dictionary defines a drone as “an unmanned aircraft or ship guided by remote control or onboard computers.” Drones are sometimes known as Unmanned Aerial Vehicles or Unmanned Aircraft Systems.

As Figure 1 illustrates, drones often carry cameras to take pictures or record videos. Originally designed for military purposes, this technology has been increasingly adopted for non-military uses. For instance, drones are used to cover ongoing events for journalism [51], record birthday parties [34], deliver packages to customers (e.g., Amazon Prime Air [4]), and to assist police in patrolling and investigation [32].

In this paper, we focus on lightweight drones with operators for civilian uses including public (governmental, e.g. police), civil (non-governmental, e.g., commercial), and recreational (a.k.a, model aircraft) purposes [20]. This type of drones dominates the consumer market and can have broad and emergent impact on ordinary citizens. While no official drone sales data is available, the Consumer Electronics Association (CEA) estimated that 700,000 drones were sold in 2015 in the US [14]. From now on, we use the term drones to denote this type of drones unless specified otherwise.

Because of drones’ small sizes and capabilities in flying and taking high-definition images and videos, government agencies, policy makers, consumer advocacy groups, and legal scholars have raised serious concerns about drones’ usage. For instance, the US Federal Aviation Administration (FAA) is tasked to devise rules for drone use by 2015. Ann Cavoukian, the Privacy Commissioner of Ontario, Canada, has advocated that designers should adopt a *Privacy by Design (PbD)* approach from the beginning of the drone design process to protect people’s privacy [12]. Legal scholars have raised ethical and privacy concerns regarding the use of drones (e.g., [6]).

However, the extant literature mostly from legal scholarship focuses on conceptual analyses of drones and their im-



Fig. 1. A DJI Phantom 2 Vision+ drone that we used in our study

plications. There is a lack of empirical studies that examine people's perceptions of this emerging technology with one exception being a recent survey study of Australian citizens' perceptions of drones [13]. However, little is known about how people in the US feel about drones, particularly around privacy. Understanding people's privacy perceptions is integral in informing future privacy-friendly drone design and regulation. Our research aims to fill this critical gap.

During June to August 2015, we conducted 16 semi-structured interviews to examine people's perceptions of drones. To help our informants get familiar with this technology, we showed them a real drone (see Figure 1) and illustrated its capabilities in flying and taking pictures and videos before the actual interviews. In each interview, we solicited our informants' general perceptions of drones as well as their perceptions under five specific usage scenarios that we adopted from drones' existing real-world uses. We also asked them to compare drones with two tracking/recording technologies that they are already familiar with, smart phones with cameras and closed-circuit television (CCTV), as frames of reference. Lastly, the informants were asked about what kinds of notifications and controls they would expect from drones operated by others as well as what aspects of drones should be regulated.

Our results suggest that our informants had mixed feelings about drones. On one hand, they saw clear values in drones as they identified many benefits and promising applications of drones. On the other hand, they also raised a multitude of safety, security and privacy issues. Our informants were not only concerned about the drones per se, but also the drone controllers that are often invisible. Drawing from Orlikowski's conceptualization of duality of technology, we highlight the *duality* of drones, suggesting that drone design and regulation should consider both drones and their controllers.

This paper makes three contributions. First, we provide a detailed account of people's privacy perceptions of civilian drones. Second, we highlight the *duality* of drone and its value in unpacking people's privacy perceptions of drones. Third, we discuss implications for privacy-enhancing drone designs.

2 Related Work

To situate our work in the literature, we review three lines of related research: (1) people's privacy perceptions of tracking/recording technologies, (2) challenging issues of drones, and (3) privacy mechanisms for drones.

2.1 Perceptions of tracking technologies

Since drones are usually equipped with cameras, they can be classified as tracking/recording technologies. The only user study of drones that we are aware of is a recent survey study of Australian public's perceptions of drones [13]. Overall, the respondents did not consider drones to be overly beneficial or risky [13]. However, some respondents (less than one fifth) did raise a general privacy concern about drone surveillance or spying [13]. Prior studies have identified people's privacy concerns over other tracking and recording technologies. For instance, based on interview and survey data, Nguyen and Hayes suggest that people are concerned about leaking personal information about themselves with institutional and end-user tracking and recording technologies, such as credit cards, store loyalty cards, and store video cameras [43].

In studying Internet users' perceptions of online tracking and online behavioral advertising (OBA), McDonald and Cranor find that while people accept that free online content needs advertising, they reject the idea that they need to give up their data for that exchange [42]. Ur et al. show that people have a conflicting sets of opinions towards OBA, describing it as smart, useful, scary, and creepy [61].

Felt et al. find that mobile phone users have varied yet strong privacy concerns in using their phones, particularly, the potential tracking and leakage of their text messages, e-mails, and photos stored in the phones. Users ranked highest risks in using their mobile phones as all contact information being deleted and messages or calls being sent out by malware without their awareness [21]. Tsai et al. show that people also have privacy concerns in using location-sharing technologies, but their concerns vary across different scenarios [60].

Moving on to the physical world, results from a survey conducted shortly after 9/11 show that the majority of respondents approved expanded use of camera surveillance (CCTV) in public [64]. Angeles has found that people have varying level of privacy concerns over the use of Radio-Frequency Identification (RFID) tags [5]. In particular, less information-sensitive people will favor the benefits from RFID more, and will be more willing to buy and pay more to RFID tagged products; whereas more information-sensitive people are more concerned about their privacy with RFID [5].

Prior research has also explored people's perceptions of wearable devices (e.g., glasses or cameras). Hong suggests that since most people have little experience with wearable devices (e.g., Google Glass) before, their perceived value and perceived risks of these devices may change over time [28]. In a study of Augmented Reality (AR) glasses, Denning et al. find that people expect giving their permissions before being recorded by AR glasses [16]. When participants compared AR glasses with CCTV or surveillance cameras, they did not in-

dicating any evident difference in their attitudes. They felt that these technologies are always recording in public and that the introduction of AR glasses did not affect their expectations of being exposed to various recording technologies [16].

These wearable devices can also be used for “lifelogging” where photos and/or audio/video recordings are automatically taken by the devices as a person goes about doing his/her daily activities (e.g., SenseCam [26]). Hoyle et al. find that people have many privacy concerns about lifelogging [31]. For instance, they are concerned about sensitive information appearing in the “lifelog,” such as their locations or credit card numbers. They are also concerned about the privacy of bystanders since their faces or behaviors may be captured in the “lifelog” [31]. In a follow-up study, Hoyle et al. also discover that “lifeloggers” are motivated to share their “lifelogged” information for impression management purposes [30].

Last but not least, robots when equipped with cameras also have tracking and recording capabilities. Edward Hall proposes proxemics to refer to people's use of space in mediating their contact with others [25]. For instance, if strangers enter into someone's personal or intimate spaces, then the person would feel uncomfortable [25]. Researchers in the field of human-robot-interaction (HRI) have used this concept in studying the interactions and relationships between humans and robots. Studies have found that a robot's form, speed, and height have different degrees of impact on people's perceptions of the robot (e.g., [10]). In a recent study, Butler et al. find that people desire mechanisms to protect their privacy against remotely tele-operated in-home robots [9].

This literature on tracking/recording technologies suggests that people are likely to have privacy concerns with these technologies, but people might have different levels of concerns. People's privacy concerns might also vary across different scenarios. These findings inform us to examine both general and scenario-based privacy perceptions of drones.

2.2 Challenging issues of drones

Besides the Australian user survey of drones [13], the extant literature on drones has largely focused on privacy and security issues from a legal perspective. The legal scholars posit that drones could potentially violate the Fourth Amendment that protects citizens from unreasonable searches and seizures when drones are used for surveillance. Therefore, the Fourth Amendment rights should regulate and restrict drone usage [18]. They also criticize the FAA for not taking more responsibility and initiative in monitoring drone use. For instance, Barbee comments that the potential use and misuse of drones are both considerable and must not be neglected, yet neither the FAA nor the Congress has paid sufficient attention

or taken any action to address the relevant challenges, particularly privacy issues [6]. Research has also suggested that drone developers are somewhat aware of the laws but tend to ignore ethical issues. They would default to some legal considerations of privacy based on their justification of whether the subjects would predict that they are being photographed or video recorded by a drone [15]. Other scholars have heightened concerns due to the fact that drones could be cheaper to obtain than before and could be so tiny yet still with high-definition cameras (a.k.a., “dragonfly drones”). Therefore, drones could potentially get even more detailed pictures of the people being monitored and it would be even harder for people to notice the drones and be aware of them being watched [65]. These legal analyses are informative but lack empirical data about privacy perceptions from ordinary citizens, particularly in the US context. Our study aims to provide this type of empirical data.

2.3 Privacy mechanisms for drones

An number of technical mechanisms have been proposed to protect civilians' privacy specifically regarding drones. For instance, to help drone controllers operate drones appropriately, the FAA has developed B4UFLY, a mobile app that helps drone controllers “determine whether there are any restrictions or requirements in effect at the location where they want to fly” [19]. Besides, ordinary citizens can sign up their addresses as part of the no-fly zones for drones which may be incorporated into the firmware or software of drones and/or honored by drone controllers [47]. To provide citizens more information about drones, LightCense is a system that uses flash lights as a drone's ID. People can look up information about the drone by decoding the sequence of lights via a mobile app [40].

As an example of a server-side mechanism, Yoohwan et al. propose using a combination of encryption, access control, and image/video transformation. Specifically, the system would encrypt the images or videos taken by drones and then deliver them to a privacy server. To access these photos or videos, the privacy server would require a shared key. Depending on the privacy policies of the drone's surveillance area, the pictures and videos can be transformed from high-definition to blurring or totally blank out [37].

While our present research does not design a specific privacy-enhancing mechanism for drones, our study can offer insights to inform future designs of such mechanisms.

3 Methodology

From June to August 2015, we conducted 16 in-person interviews to explore people's privacy perceptions of drones in Syracuse, New York (US). Each interview took about 1 hour with a study compensation of \$10. This study was approved by our University IRB.

3.1 Participants

To recruit a diversified set of informants, we posted study fliers and randomly invited adults in public places such as university campus, streets, and parks, to participate in our study. Half of our informants were male, and the other half were female. Their ages ranged from 18 to 62 years old with an average of 29. Our informants represented various ethnic backgrounds, such as White Americans, African Americans, Latino Americans, and Asian Americans. The majority of them were university students, but we also had a news reporter, a student counselor, an office administrator, and a retired worker.

3.2 Interview protocol

To help our informants get familiar with drones, we used a DJI Phantom 2 Vision+ drone as a prop in our interviews (see Fig 1). This drone has a HD camera and can provide live video feed via a dedicated mobile app.

Each interview was structured as follows. First, before the interview, we showed our informants the physical drone and explained in details how the drone could be controlled to fly and take pictures/videos. If the weather permitted (e.g., not too windy or rainy), we also flew the drone in front of the informant. We also encouraged our informants to ask any questions about drones before formally starting the interview. This kind of in-situ investigation could give informants a more realistic impression about the technology and would be more natural to probe people's perceptions, particularly when people are not very familiar with the technology [16].

Specifically, our interview protocol consists of three parts: (1) general questions about people's perceptions of drones; (2) context-based questions about people's attitudes towards drones under different scenarios; and (3) questions about specific aspects of drones, such as comparisons between drones and camera phones or CCTV as well as expected notice, control, and regulation of drones. The interview questions are included in the Appendix A. Using a semi-structured interview approach, we also asked follow-up questions to continue any interesting discussion.

3.2.1 General questions about drones

We began our interview with general questions to explore informants' understanding of and attitudes towards drones. For instance, we asked questions such as, "Have you heard of drones? What is the first thing that comes to your mind when you hear about drones? How do you feel about drones? Do you see any benefits or issues of drones?"

These questions were mainly adapted from two prior studies on Online Behavioral Advertising (OBA) [62] and Augmented Reality (AR) Glasses [16], respectively. We chose to build on these two studies for a few reasons. First, both studies conducted interviews with ordinary citizens in the US. Second, at the time of the studies, OBA and AR Glasses were trendy and somewhat controversial technologies which ordinary people might not have much knowledge or experience. Third, drones, OBA and AR Glasses all can be used to support or benefit people's lives as well as to track people and potentially invade people's privacy.

We also asked informants to compare drones with two widely used and known tracking/recording technologies, smart phones with cameras and closed-circuit television (CCTV), as frames of reference. This comparison between drones and more familiar technologies was inspired by a pioneering study of risk perception [22]. The main reason we chose camera phones and CCTV is that since they are widely used, ordinary citizens are likely to be *familiar* with them so they can be used as references. We did not choose RFID, Google Glass, or other wearable cameras (e.g., SenseCam [26]) because people may be unfamiliar with them just as drones.

3.2.2 Scenario-based questions

There are both theoretical support and empirical evidence that privacy is contextual. For example, Helen Nissenbaum eloquently points out that human behaviors, e.g., a transaction that occurs, is always situated in some concrete context, e.g., certain geographical area and specific constituted norms within a political, cultural environment [45]. As we discussed in the related work, prior privacy studies have also shown that people's privacy preferences of technologies can vary significantly under different contexts or scenarios (e.g., RFID [52], location-sharing systems [60]). These studies motivate us to develop different scenarios and understand people's context-based privacy perceptions of drones under these scenarios.

We created and presented five specific and realistic drone usage scenarios. We asked our informants if they would accept the drone usage for each scenario and why. We adopted news reports of real-world drone usages in creating the five scenarios: a drone is being used in (1) recording a promotion

event that you attend in a shopping mall by a store owner [66]; (2) delivering goods to you by Amazon [4]; (3) recording a friend's party that you attend [34]; (4) reporting a parade that you attend by a news agency [51]; and (5) searching lurking criminals around your residential area by the local police [32]. These scenarios covered a diverse set of contexts, including indoor use (mall) vs. outdoor use (parade); private home (party) vs. public area (parade); the drone controlled by individuals (friend), the government (police), or a vendor (Amazon); and the drone use benefiting self (goods delivery), other people (friend), or other entities (mall).

3.2.3 Expected notification and control

To help inform future drone design and regulation, we also asked informants about what kinds of notifications and controls they would expect and what aspects of drones should be regulated. Specifically, we asked questions, such as “do you expect to be notified about the time periods during which drones can/will be operated” and “do you expect to be notified about the types of information that the operating drones might collect?” These questions were inspired by the *Drone Aircraft Privacy and Transparency Act of 2013*, which was proposed but not enacted in the US. We also asked questions about expectations of consent and control, such as “do you expect to be asked for any kind of ‘explicit consent’ to allow drones to fly near you” and “do you expect to see detailed explanations if a drone takes pictures or videos that can capture you?” These questions were adapted from a study on RFID [29].

3.3 Data analysis

We audio recorded all the interviews upon informants' permissions. We also took notes during the interviews. The interviews were then transcribed and analyzed qualitatively. In general, qualitative research or analysis is particularly useful in exploring the why and how questions of a social phenomenon. Qualitative research usually does not claim representative results in the statistical sense but allows researchers to make propositions that can be further investigated by quantitative methods such as surveys or experiments.

In our case, we conducted a *thematic analysis*, which is common for qualitative research [7]. Thematic analysis is “a method for identifying, analysing, and reporting patterns (themes) within data” [8]. First, we immersed ourselves in the data by carefully reading through the interview transcripts, actively looking for and taking notes of meanings and patterns.

Second, two co-authors (coders) used ATLAS.ti, a popular qualitative analysis software, to manually and indepen-

dently generate initial codes that capture meanings of the same subset of our interview data at a fine-grained level (usually at the sentence level). These codes are considered as the most basic elements of the phenomenon under our study. Then, the two coders convened, discussed, and converged their codes into a code book of 132 unique codes ranging from individual drone features (e.g., cameras) to usage of drones (e.g., parcel delivery) to concerns of drones (e.g., stalking). Next, the two coders used the agreed-upon code book to code the interview data. ATLAS.ti allows us to identify and extract all excerpts associated with a given code. For instance, one interview quote for the code “public space” is “*everyone is free to go in and out of that place whenever they want to and they're basically free to do whatever they want unless it's against the law.*”

Third, we explored how different codes can be merged into high-level themes. We grouped 132 codes into nine candidate themes: drone features, drone usage, attitudes towards drones, cultural differences, private vs. public space, privacy concerns, safety concerns, and drone control. For instance, the theme of drone control included the following codes: notification, accessible to everyone, actions to protect people, controller flexibility, delete recordings afterwards, destination of Information, expected control, know controller, sound of drones, time to fly, and regulations. We wrote codes on colored post-it notes and sorted them into groups/themes on a wall, creating an affinity diagram [36].

Fourth, we reviewed the candidate themes by reading the interview excerpts of each theme to see whether they coherently present the underlying theme. We then adjusted the themes and our affinity diagram accordingly. For instance, the code “battery life” was originally part of the theme of drone features. After reviewing the interview quotes associated with the code “battery life” (e.g., “*they should know if it's safe or not to use and I don't know if they use batteries*” and “*And again it could run out of batteries*”), we moved this code to the theme of safety. Figure 2 shows the final affinity diagram.

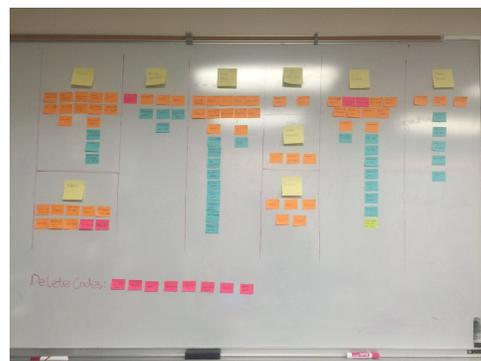


Fig. 2. The final affinity diagram of the themes and their codes.

4 Findings

In this section, we report the major themes emerged from our study. We will use fictitious names for our informants.

Our informants suggested a number of perceived benefits of drones. When asked about potential benefits or applications of drones, our informants focused on a few characteristics of drones, such as their relatively small size, agility, and capabilities to fly and to take high-definition photos and videos in inaccessible or even dangerous environments. They envisioned several drone applications, such as aerial photography, goods delivery, and emergency responses.

Our informants also raised several issues of drone usage related to safety, security, and privacy concerns. Their safety concerns mainly revolved around drones hitting people or interfering with other aircraft. The informants attributed these risks to two sources: the components and/or features of drones (e.g., propellers), and drone controllers' inappropriate or reckless behaviors. Closely related to the safety issues, our informants also brought up concerns about security issues, mainly about drones trespassing on some forbidden or sensitive places such as government or military facilities. When drones encroach on personal spaces which individual informants defined for themselves, a sense of privacy violation arose.

4.1 General privacy concerns

Privacy was a salient and consistent topic across all of the interviews, regardless of the diverse ethnic or occupational backgrounds of our informants. Their discussions about the privacy implications of drones centered around the following themes: (1) the definitions and boundaries of public vs. private spaces, (2) peeking and stalking, and (3) recording and sharing of photos and videos without people's awareness and/or consent. These concerns were related to both their *physical* privacy and *information* privacy. Informants' physical privacy concerns were primarily about the feeling that their private spaces would be intruded by drones. This sense of physical privacy intrusion is similar to when an individual's personal space is invaded by a stranger [24]. In terms of information privacy concerns, our informants were mainly concerned about the collection, use, and sharing of their personal data, such as their locations and pictures/videos that capture them.

4.1.1 Public vs. private spaces

Due to drones' flexibility and mobility, they could intrude into people's private space, compromising people's physical pri-

vacuity. Territoriality (public vs. private spaces) is a key factor that our informants considered in determining their expectations of privacy and the privacy violations of drones. There was a general consensus among our informants that if a drone takes pictures or records video or even just flies in a *private space*, then the drone would be considered as invading the residents' privacy. While intuitive, this view begs an important question - what is considered a private space? Our informants had various definitions of private space and these definitions centered around three factors: ownership of the space, sensitivity of the space, and nature of activity in the space.

Ownership. In general, our informants agreed that their homes (either owned or rent), or personal properties (e.g., a car) are their private spaces. For one group of informants, *ownership* (or temporary ownership such as rental) alone determines private spaces. For instance, Scott (62, retired worker) explained, "*the boundary between public and private space is like a fence of property an individual owns.*" Similarly, Bill (25, computer science major) claimed that "*my private space would be my private properties, my home, my rental house, my car, etc. and public space is owned by the government or public administrations, like school, city square, etc.*" Dan (27, public relation professional) also focused on ownership: "*like a lot of places outside your own property or where you live is kind of a public space because you don't own it, like a park is a public space.*" Emily (18, first year in college) extended her private space beyond her home, explaining "*the surrounding place around my home is still my own private space.*" Mary (28, teacher) went to an extreme in saying that "*I just assume the space outside of my house is public space.*" For these informants, ownership was reasonably easy to identify and so was their private spaces.

Sensitivity. When a drone operates in a public place, our informants generally felt that the drone is less likely to cause privacy violations. However, some informants effectively treated *sensitive* public places with children or other vulnerable populations as private spaces. For instance, Hannah (19, biology major), who used to assist in an elementary school, considered a school as a private space: "*especially in elementary schools a lot of people don't want their children to be recorded or taking photos of them unless you have the consent of their parents.*" This view is compatible with the current US privacy legislation for children, such as the Children's Online Privacy Protection Act (COPPA) which requires the consent from a parent or guardian for collecting personal data about children under 13.

Nature of activity. For some informants, even seemingly non-sensitive public places (e.g., mall) can be their private spaces because of the nature of their activities. For instance, Lily (24, information science major) not only considered the place but also the activity she is engaging in the place at that

time. She elaborated, *"I only regard it a public space when I go to a so-called public area, like a square and in the meantime I am participating in some public event, like a promotion event, or a parade. Otherwise even I am in a public area, I still regard it, particularly my surrounding as my private space."* Lily's viewpoint is related to the notion of proxemics which refers to the personal space (or distance) that people maintain around themselves [25]. Lily desired proxemics against drones even in public places when conducting her personal business.

4.1.2 Peeking and stalking

There were two behaviors of drones that our informants regarded as intrusions to their privacy: peeking and stalking. In terms of peeking, almost all the informants generally loathed being watched or recorded by a drone, peeking through the windows of their private spaces. For instance, Cindy (21, finance major) explained her concern, *"because a drone could fly so high, even if I am living on the top floor of a building, I would still worry that a drone may peek me through the window when I am doing some private things, like taking a bath."* Emily shared the same attitude. She said *"If I'm in my own private home I won't like a drone peeking into my house."* Dan drew an analogy between a drone's peeking and a neighbor's peeking, *"it's the same thing in houses, people don't want the neighbors to peek in. In a society where everything is documented all the time, you know, you don't need another thing adding to that."* By peeking, drones can invade into people's private spaces and lives.

Stalking means that a drone could follow and record an individual's activities. Bill painted a dreadful picture against a backdrop of the current social and cultural landscape in India where he was originally from: *"One concern would be stalking. In India, parents care about their daughter very much, and if they see a drone stalking their daughter, they will be very angry and use every means to find the controller of the drone and punish him, even if the controller is unintentional."* Even perceived or unintentional watching or stalking may lead to revenge and grave consequences. In the US, there were cases where people shot down drones over their backyards [55]. However, Bill's example brought drone controllers to the foreground. They are the ones who will be held accountable for the drones' behaviors.

Mary also denounced stalking or watching and alluded to the problem of the lack of control for drone ownership. She explained, *"there are some very emotionally unstable individuals out there so to have everybody able to own a drone and that I could have some crazy person watching me, yeah that's a problem."* Mary's concern is not unfounded. While the FAA requires licensing of commercial use of drones and registra-

tions of drone controllers, practically anyone can buy a drone for personal use in the US. While the media often focuses on drones' potential use in government surveillance, Mary called attention to misuses of drones by individual controllers.

4.1.3 Recording and sharing

Last but not least, our informants were concerned about drone controllers taking and sharing of photos and videos without people's awareness and/or consent. This concern mainly stemmed from a sense of uncertainty about how drone controllers would collect and use people's information, because a drone can be remotely controlled and can fly out of sight while recording. In other words, both the drone and its controller might be invisible to the people being watched or recorded. Furthermore, people may not be aware of the recording nor have access to or control over the drone recording about them.

Abby (20, environmental science major) explained the invisibility of drones and the lack of awareness and control of the recording, *"People can't tell it's there and will not be aware that they are being watched by such a tiny drone and the footage by drone may be used for whatever purposes without their consent or knowledge."* Abby's concerns pointed to two more fundamental issues: drones' capabilities in capturing pictures or videos of individuals inconspicuously, and bystanders' lack of knowledge of and access to drone controllers.

Dan further highlighted the importance of the drone controller behind the scene, *"drone kind of could be a robot or could be pre-programmed to just like stay in one area all the time so you'd like to know who's behind this and I think that's kind of an important point to raise because people fear what they don't know, so if they don't know who's steering it, that raises some concerns."* Cindy was also concerned about the drone controller and the potential usage of drones' recordings. She elaborated, *"If the person who controls the drone is a person that I don't know, I will be concerned maybe he will use this video to do some other things like put it in the advertisement or yeah, or some illegal things. So I must ensure that the person who controls the drone is someone reliable or the person that I know."* This foregrounding of drone controllers underscores the *duality* of drones. People not only deal with drones but also the people who control the drones. Privacy is deeply relational [59]. Cindy's concern was aggravated or attenuated by the *social relationship* between her and the drone controllers (e.g., strangers or people that she knew).

In addition, the audio part of videos can be another concern. For instance, one informant said *"if you're having a conversation with someone like on the quad and that's still like your own kind of private conversation...well maybe not ever but most of the time I wouldn't want like audio recorded."*

Lily raised a quite different concern not due to the flights of drones per se but due to pictures being taken by drones and later being posted on the Internet. She explained, *“I know someone immigrated from Afghanistan, and they don't like their pictures being posted on the Internet because they are still in touch with their families, and it's a security issue for them. If they are studying in the university and the tribes are here, (post their pictures on the Internet) will be bad for them.”* Lily's example reminds us the privacy and security risks are engendered in their rich social and cultural settings. Exposing people's seemingly public activities such as studying in a US university could potentially put people and their families at risk. The lack of knowledge and control of data collection and sharing by drones paralyzes people's desires and abilities to manage boundaries between themselves and others in achieving the right level of privacy.

4.2 Context-based privacy perceptions

After asking our informants' general perceptions of drones, we then provided five specific drone usage scenarios to further probe informants' context-based privacy perceptions of drones. Our informants' perceptions of drones varied across these scenarios. The differences mainly stemmed from three sources: (1) whether the scenario occurs in a public or private space; (2) what is the intended purpose of the drone usage; and (3) notification and consent of the drone usage.

4.2.1 Scenario 1: shopping mall event monitoring

We described this scenario to our informants, *“Imagine you are in a shopping mall where a promotion event is going on, like on the Black Friday. The store owner decides to use a drone to monitor and record this event, and you happen to be in that event.”*

12 out of 16 informants considered a shopping mall as a public space, and they expressed little concern about being recorded by the drone in this scenario. However, four informants considered shopping mall a private space. For instance, Cindy explained, *“Shopping mall is a relatively private space for me if I go shopping with my intimate friends and I don't want to be audio or video recorded if I am having some private conversation with my friends.”* Again, Cindy's reasoning points to the relational aspect of space, or the social space is characterized by the social relationship therein. The intimate personal relationship makes the space intimate and private.

Several informants expressed that the drone's recordings should be restricted to the promotion event and that the drone should not appear around sensitive areas, such as the dress-

ing room. A few informants would *prefer* to be notified about drone recording, as Emily put it, *“I guess you could like put out some sort of notification to the people...or making people aware that you will be recorded.”*

4.2.2 Scenario 2: recording a friend's party

In this scenario, we told our informants, *“Imagine you are at your friend's party, and your friend decides to use a drone to record the party.”* Five informants perceived their friend's party as a public space. For example, Sue held that, *“it is a public setting so I don't really see that would bother me too much.”* However, the other 11 informants felt it is something in between. For instance, Cindy said *“I think it's kind of between the private and the public. There are many friends I know so it can be taken as private, but there are lots of people, so I also think of it as public.”* This perspective suggests that it is not always a clear cut whether a space is public or private.

Another important factor they considered was the purpose of the recording. Sue commented, *“I could see some concern about that but if you're concerned about your image, you should not consent to go to that party and get drunk in public and anyone can record it.”* Grace could accept this scenario if it is for personal use. She explained, *“I think if it's for his personal use I think it's okay because you've shown consent in being in that space and being with other people.”*

4.2.3 Scenario 3: delivering goods

We told the informants *“Imagine that Amazon decides to use drones to deliver goods that you have bought in its online store to your house.”* All informants felt that a drone delivering goods would fly close to their home, which was unanimously regarded as their private space. However, most of them felt this drone usage is cool. Dan shared his excitement, *“that's like an efficiency thing, that's making life more convenient and better, and it serves a good purpose and I guess yeah, I think that would be pretty cool, that's a cool way of using technology.”* The informants did mention the potential safety risks such as drone crashes. In addition, some informants felt it is not necessary for the drone to carry a camera. For instance, Abby noted that *“I think the drone should have the whole map in its GPS system, you know, you don't have to use the camera.”*

Among the five scenarios, our informants expressed the least demand for notification and consent in this scenario perhaps due to its “useful” nature. However, Grace still wanted notifications: *“there should be some sort of disclosure on how it's going to be sent because normally when you purchase something on Amazon you choose your shipping method.”*

4.2.4 Scenario 4: reporting a parade

For this scenario, we told the informants *“Imagine you are in a parade. Some news agency decides to use a drone to record the parade.”* In general, our informants were least concerned about this scenario. They all agreed that this scenario occurs in a public space. In addition, all informants except two felt that the drone's recording in this case is acceptable because the purpose is for journalism and if they have already decided to participate this parade, they would rather the parade be recorded and publicized.

However, Mary and Cindy had some reservation. Mary was concerned about her vanity. She explained, *“That's just like vanity...just like wish that I would have done my hair that day so that my appearing on news would be great.”* Mary cared about the presentation of herself, a form of impression management practice that Goffman observes [23].

Cindy was more concerned about whether the parade is controversial or not. She elaborated, *“Say I am in a feminist parade. If my face is recorded by the drone and it is put online or in the newspaper, or in other media. The people from the other side, the anti-feminist side, if they meet me later in street, they may revenge me.”* This example alludes to the lack of prior notice of and control over drone recordings as well as the potential ramifications. If she had prior notice, she might have reconsidered her participation in the parade.

Our informants mainly requested for prior notice of using their images in news, especially when the images reveal their identities. For instance, Sue (21, biology major) suggested, *“I guess you would need people's permission before you use their face and you know you post their image on website or social network whatever.”*

4.2.5 Scenario 5: searching for criminals

We told informants *“Imagine that there are some suspects or criminals lurking in your residential area. The police department decides to use drones to search for these people in your residential area.”* Most informants felt this is acceptable, and some of them even applauded this practice. For instance, Joe (59, newspaper reporter) explained, *“That could be better because a drone could get there immediately even before police car gets to the scene. So I think it will be very useful. If waiting for the police, by the time the police comes, maybe the subjects have already fled away.”*

However, three informants considered the surrounding area of their residence as their private space and felt uncomfortable having a drone patrolling around their house, invading their physical privacy. For example, Emily said, *“I would be uncomfortable but I understand why they have it, like if it was*

in my neighborhood for example, and there was a drone flying around I would be thinking like what's going on.”

Given the recent Snowden revelation of the extensive government surveillance, it is perhaps not surprising to see that most of the informants requested explicit notification and consent for this practice. Bill, articulated his expectation, *“I would like to have explicit information from the police department. Because the home is my private space, having a drone flying above me and recording is like having a policeman watching me around my living place all the time...I don't like to be watched or surveilled by a drone, especially without my permission and prior knowledge.”* In Bill's view, the patrolling drone is like a pair of flying eyes watching his life at home.

Many issues that surfaced from these scenarios include physical privacy, purpose of data collection and usage, notice and control. These issues have long been recognized when examining privacy implications in technologies. However, this begs the question: are drones any different from other familiar tracking technologies that have raised similar privacy issues?

4.3 Comparing drones with other familiar tracking technologies

We asked our informants to compare drones with two familiar technologies that have tracking/recording capabilities: smart phone with a camera, and closed-circuit television (CCTV). Our informants pointed out both similarities and differences.

4.3.1 Comparing drones with camera phones

From the informants' perspective, the main similarity is that both drones and camera phones can take pictures and record videos. The major differences, however, lie in the distance of recording, and the visibility or accessibility of the owner or controller. For instance, Dan explained eloquently, *“It would be a lot further away with a drone and they're hidden away but still get a really good shot. So I feel like there's that kind of not knowing that this is happening as opposed to a cell phone where you get a much larger chance of seeing that person taking that photo of you.”* Dan's reasoning again speaks to the flying (drone) cameras at a distance that can take pictures or videos of individuals inconspicuously (i.e., without their awareness).

Emily focused on the controller of the device. She said, *“When you see some people taking picture of you using their phone, you can go to them and ask them to delete the pictures or videos. But when a drone is taking picture or recording video of you, you cannot control it, and it can easily fly away, or the pictures and videos have already been trans-*

ferred to the controller's mobile phone and you may not even know where the controller is." She emphasized the hidden controllers that are behind the scene, inaccessible to the people being recorded. As such, the invisibility and inaccessibility of drone controllers make people's usual privacy practices (e.g., ask camera controllers not taking or deleting the photos) futile.

4.3.2 Comparing drones with CCTV

While both drones and CCTV can take videos, the informants suggested a few notable differences, including their flexibility, visibility, intended purposes, and trust in the controllers. For instance, Mike focused on the flexibility. He said, "*Drones can fly anywhere so it can dynamically record everything. As far as I know, CCTV can only record as long as you're in that area.*"

Emily explained the difference in purpose, "*Well there's obviously a reason why they have the (CCTV) cameras in there, it's like for security and safety and like I know why they're there and I think it's kind of like the norm. People are already like kind of used to having security cameras.*" Emily highlighted that CCTV is a familiar technology now and people have established norms or expectations of it, whereas drones are too new to have agreed-upon norms.

In terms of visibility, Grace articulated the difference, "*A security camera is put in a place where it's very visible, usually places will post some sort of sign that says there's a security camera. So there's that disclosure and you understand if you step into that space you are going to be recorded. But I think a drone has the ability to enter someone's space rather than the person going into a space and then not having that disclosure that it's being recorded.*" What Grace illustrated is a metaphor that CCTV passively waits for people to be recorded whereas drones actively enter into people's space to record them. In a way, this nature of drones shifts the initiative from the people to drones, weakening the people's control of the situation.

Our informants also pointed out differences in their trust of the controllers. "*I wouldn't mind what you would be doing indoors, that's for security, but outside then it's beyond your control because inside you know who is having that, who is handling that...definitely we know who is handling the drone. Maybe the security person or something like that, right. So you trust them, you would not fear unless you are the one who's going to be the trouble inside.*" This informant highlighted the importance of trust or the lack of it for drone controllers.

Because drones and their controllers can be more difficult to recognize and approach, we next discuss what people would expect in terms of controlling and regulating drones.

4.4 Expected notification and control

Our informants proposed a few controls of drones and their controllers. They also suggested regulating drones in terms of their size as well as their flying altitude, area, speed, and time.

4.4.1 Tracking drone controllers

Because of the potential for safety issues and malicious use, several informants suggested the need for drone controllers to register so that they can be tracked and held accountable. In addition, training and certification were recommended for operating drones. Interestingly, our informants used driver's license or gun license as an analogy to drone controller license.

For example, Cindy treated a drone license like a driver's license: "*just like drivers need driving lessons, I think the drone controller also needs a license because if you didn't control it very well and it can make some damage to the environment and may also intrude other people's privacy.*"

Those who compared a drone's license with a gun license felt both technologies can be used for bad purposes. For instance, Hannah proposed a drone license to keep track of the controller, "*You can't just buy drones whenever and wherever, you have to have like a license or it would be registered under your name with that serial number so people can identify whose was that, so people can't use it and like if they used it for something bad it would be under your name.*"

4.4.2 Tracking and controlling drones

Our informants also proposed ideas to track drones such as using a unique ID, and detecting and monitoring drones via a mobile app. Mary was one of the informants that proposed a unique ID for each drone. She explained, "*For example, let's say some people use a drone to do bad things, and you can track the drone by the serial number. It would be an evidence that you have used your drone to do bad things, such as you used your drone's camera to see the forbidden spot.*"

Grace hoped for something more convenient. She imagined, "*I would hope to see in an app or something to show what the drone could see or record from flying above my home.*" What she requested is a technology that could essentially discover nearby drones and monitor what the drone camera can see, but more fundamentally, she asked for more awareness of drones and their operations.

While our informants provided suggestions for tracking drones, they felt that they cannot stop drones from flying or taking pictures. For instance, Lily said "*I can't stop them [drones] from taking pictures. I will just stay away.*" Abby also

expressed her inability to stop drones but suggested that drone manufacturers may do something to prevent misbehaviours of drone controllers. She speculated, “*Let’s say they [drone manufacturers] have a backup chip to capture the images from the phone, so if a controller is doing any illegal activity that would be stored in this purposeful part of this chip, which a controller has no access to but the manufacturer can get into that. That would be the only way to curb the bad activity.*”

Given the powerful nature of drones, regulation is warranted. While the FAA has been finalizing their drone regulation, what do people expect from the regulatory front?

4.4.3 Regulations on drone features and operations

When asked what aspects of drones should be regulated, our informants’ responses focused on two aspects: drones’ physical properties such as size, speed, and color; and drones’ operational properties such as flying height, area, and time.

Some informants held that the size of a drone should not be too small or too big. Abby explained, “*If a drone’s size is too small then that would be weird because you cannot see it and that’s definitely used for spying. I don’t want a drone like a Boeing though, that would be pretty bad if it flies low and it would be scary. It should not be bigger than this table [one yard long, and 16 inches wide].*” Besides, Hannah felt the need to regulate drone speed. She said, “*I would want to regulate the speed of drone, about how fast it goes...may be not too fast.*” Several informants mentioned about noise control. Mike explained, “*Personally you may not want to hear any noise. I may be sleeping so if somebody is outside flying a drone, I feel like so it’s just getting annoying.*”

Finally, a few informants proposed to color specialized drones, e.g. drones used by the police, so that they can be more identifiable by the public. Hannah explained, “*Using color that is specifically for police, just dye your drone would help people. So when I see that drone flying I know that this is for the police and I’d be kind of okay with that.*” The colors signify purposes or ownership which could produce trust: “*I fly somebody’s drone then how can you be sure that that drone does not carry a gun. How can you trust those things...that’s why I say if you use different colors you trust them that it’s some government or some legalization purpose so you trust*”

In terms of drone operations, our informants proposed to regulate its flying height, area, and time. For instance, a few informants emphasized that a drone should not be allowed to fly at night for privacy and noise considerations. For example, Sue said, “*It would be weird if the drone is flying at night...particularly if it is flying around my private space.*” Drone regulations should consider these dimensions.

5 Discussion

While all of our informants had heard of drones before, it is still a relatively new technology to them. Overall, they had mixed feelings about drones. On one hand, they saw clear values in drones as they identified many benefits and promising applications of drones, such as aerial photography, goods delivery, and emergency responses. On the other hand, their positive perceptions were overshadowed by a multitude of safety, security and privacy issues that they raised. Drones can be something our informants love or loathe. They used a wide range of adjectives to describe drones: from cute, cool, fun, useful and beneficial to weird, risky, suspicious, invasive, disturbing, chaotic, and dangerous. None of our informants completely ruled out drones as they always saw some drone usage under certain conditions as beneficial, but some participants expressed the sense of inability to have control over or stop drone usage.

5.1 Unpacking privacy concerns

While the news stories and government regulations tend to emphasize the safety and security concerns, the privacy issues have received less attention. Our findings suggest that the informants had various privacy concerns regarding drones. There are several characteristics of these concerns.

First, we highlight the *duality* of drones. Our informants’ perceptions of or concerns about drones were not only about drones per se, but also about the perceived relationships with the drone *controller*. Wanda Orlikowski posits the duality of technology, a recursive relationship that exists between technology and human action. On the one hand, technology mediates human action, however, it is also changed by human action [48]. Similarly, we suggest the duality of drones. Drones mediate drone controllers’ interactions with citizens, and they are also changed by drone controllers’ actions. One aspect of this duality represents the physical form and properties of the device (drone) as designed and manufactured by people and/or organizations, whereas the other aspect of this duality emphasizes the social construction of drones by the adopters and controllers through the different meanings they attach to the technology. In other words, drones manifest and extend the controllers’ agency and intention. As such, our informants often negotiated their privacy with the imagined and often hidden drone controller, mediated by the drone. It is also worth noting that drones can be used or controlled by different kinds of entities such as individuals, organizations, and governments.

Related to the social construction of drone, people’s privacy perceptions are deeply relational [50, 59]. In the friend’s

party scenario, the social relationships between the friend and the guests parochialize the private party/space and eased some informants' concerns. In the police scenario, the perceived relationships between citizens and the government (where government being a "Big Brother" or a safety guard) affected informants' perceptions.

Second, drones could violate both people's *physical privacy* [3] and *information privacy* [56]. Physical privacy often refers to the concepts of solitude and bodily privacy [3]. Information privacy relates to the collection, use, and sharing of one's personal data [56]. In their seminal paper on privacy, Warren and Brandeis advocate for "the right to be let alone" [63]. The fact that drones can fly close to people or enter into their spaces can be viewed as intruding people's solitude. Jerry Kang discusses privacy in physical space as "the extent to which an individual's territorial solitude is shielded from invasion by unwanted objects or signals" [35]. Drones can be the unwanted objects.

One of the factors that our informants considered when judging the acceptance of a drone usage is whether the drone is operating in a public space or a private space. In general, our informants detested drones flying close to their homes because the drone cameras could peek through the windows to see or record them doing private things, such as bathing as one informant exemplified. This would invade people's bodily privacy. The pictures and video taken by drones about people would obviously affect their information privacy especially when people do not know that they were recorded and how the recordings will be used.

We also note that our informants differed in their delineations of public and private space. Daniel Solove points out that the boundary between individuals' private and public spaces was permeable and pivotal in their privacy concerns [57]. Research has also shown that technologies have been blurring the boundary [27]. A few informants believed that they own their private space within a larger public space, such as parks or shopping malls. Lofland notes that technologies have transformed urban space into a "privatism," because phones, vehicles etc. have "made the withdrawal from participation in the public realm a genuine option" [41]. This is also related to the personal space that people want to maintain [25]. One informant talked about not wanting to be watched by drones if she goes to shopping with an intimate friend. She rejected the drone invading the intimate space between her and her friend while shopping in public places.

This idea of personal space also relates to Lofland's conceptualizations of urban spaces. Lofland differentiates three types of urban spaces: public, private, and parochial spaces [41]. These spaces are characterized by the social relationships therein: strangers (public space), close friends and family members (private space), and people who share com-

monalities, such as neighbours even if they do not know each other (parochial space) [41]. Accordingly, when a person is doing something personal (e.g., shopping with a close friend) in a public place (e.g., a mall), for that person, however, it is still a private space because of the social relationship (close friendship) that defines the space.

The importance of these different kinds of spaces is also related to the social norms within them. In Erving Goffman's *The Presentation of Self in Everyday Life*, he describes how we present ourselves according to the norms in the different spaces [23]. We argue that the presence of a drone can alter how people perceive the norms in which they are embedded. For instance, a drone can bring a sense of "publicness" into a private space. As a result, the drone creates tensions between expected norms of public and private spaces.

Privacy is also highly contextual. Our informants' perceptions of drones varied across different scenarios. They construed and negotiated their private and public spaces differently across the five scenarios. Helen Nissenbaum's theory of Contextual Integrity underlines the contextual nature of privacy [46]. She identifies two types of contextual norms for privacy: "norms of appropriateness, i.e., what information would be appropriate to be revealed in a context; and norms of flow or distribution, i.e., the flow of personal information in certain context needs to be reasonably justified. If either of these norms has been violated, then users' privacy is considered to be infringed" [46]. The informants paid particular attention to the purposes of the drone uses. For instance, in the friend's party scenario, some informants would accept drone recording only if it is for personal use.

5.2 What makes drones interesting?

Drone is certainly not the first tracking/recording technology that raises privacy concerns. What makes drones particularly interesting or unique compared with other technologies, such as camera phones and CCTV? Our informants identified a combination of factors. First, drones are powerfully mobile. Drones with cameras can be viewed as *flying eyes* that could flexibly and even un-noticeably fly into public and private spaces, watch, record and even share what people are doing. However, it is not the drone that is flying, but rather the controller is flying the drone. Even when a drone is flying autonomously, it is executing a plan programmed by the controller. As such, the flying eyes not only represent the drone camera but also the eyes of the drone controllers.

This leads to the second factor - the *duality* of drone. Citizens are not interacting with the drones in that they are essentially interacting with the (hidden) controller. Moreover, both the drones and their controllers can be invisible and/or inacces-

sible to the people being watched. Compared with other tracking/recording technologies, the invisibility and inaccessibility of drone/controller is exacerbated. People may not be able to detect drones from afar or approach the drone controller to find out what the drone is doing. This lack of awareness and approachability paralyze people's abilities to negotiate and enact their privacy. Acquisti et al. point out that people have considerable uncertainty about their privacy [1]. Such uncertainty is in part due to information asymmetry where technologies have made personal data collection and usage invisible [1]. Information asymmetry is not new but drones aggravate it.

Ryan Calo argues that drones may actually help to waken and restore individuals' privacy awareness because previous privacy violations are hard to visualize, thus giving consent or notification to individuals may not generate a concrete sense [11]. Our informants worried about the invisibility and inaccessibility of drones/controllers and as a result, the difficulties in getting notification and providing consent. There is currently no standard or a reliable way to enable notice and consent for drones. This makes our informants' suggestions on the design and regulation of drones particularly valuable.

5.3 Implications for design

High-level principles. Based on our findings, we first propose the following high-level privacy principles for drones:

- making both the drones and their controllers more discoverable, approachable, and accountable;
- enabling communication between drone controllers and ordinary citizens/bystanders;
- making drone designs sensitive to local social and cultural norms.

First, given the duality of drones, designs should make drones and controllers more discoverable, approachable, and accountable. Information about drones and controllers should also be made available. Adopting the notion of "accountabilities of presence," we suggest that the presence of the drone and its controller signifies a participation to a social relationship with the citizens [59]. As a result, the citizens can hold the controller accountable and ease their concerns.

Second, since the invisibility and inaccessibility of drones and drone controllers paralyze some of people's existing privacy practices (e.g., ask the camera controllers not to take or delete photos), we advocate creating channels or platforms to enable direct communication between drone controllers and ordinary citizens/bystanders. This will help build trust and form appropriate social norms over time around drone use that respects citizens.

Third, our results also hinted that the different perceptions and expectations of drone usage are embedded in larger social, cultural, and political contexts. For instance, some of our informants talked about the perceptions of drone usage in the Indian culture. Drone designs should consider the cultures or norms of the country/environment that they operate in. Different pre-defined privacy settings or modes may be used as defaults in certain countries to respect their social norms.

Next, we discuss more specific ideas for designing privacy-friendly drones. Fig. 3 shows example ideas including features from existing privacy-enhancing technologies for drones and other devices as well as suggestions from our informants and ourselves. None of them is a silver bullet to solve all the privacy issues, but collectively they will raise the bar for protecting ordinary citizens' privacy regarding drones. The ideas for drones/controllers may be built by the drone manufacturers and used by drones/controllers.

Designs for drones/controllers. From the standpoint of the drone or controllers, a number of privacy-enhancing techniques can be applied. When a camera is recording, it usually signals the recording with a red light which could be detected by bystanders if the camera is relatively close. However, since drones can fly and record at a distance, people might not detect this signal. Recently, a group of researchers have proposed a system called LightCense that uses flash lights as a drone's ID and people can use their phones to look up the drone by decoding the sequence of lights via a mobile app [40]. While innovative, it does not show what the drone camera sees, which some informants requested. Some of our informants suggested using particular colors to manifest the purpose of a drone (e.g., a police drone). Using standardized color schemes can help ordinary citizens quickly determine, for example, a police drone or a recreational drone.

Besides, there are a large number of citizens with visual impairments who would have difficulties leveraging visual cues. Instead, designers could explore designs that allow people to discover drones on their devices such as smart phones rather than manually searching drones in the sky. For instance, if a drone includes a GPS unit and flies in someone's area, that person can be notified about the drone via an app. If the controller registers with the app, he or she could also provide information about the drone (even including pictures or videos it has taken) and about himself or herself. The app users can approach and interact with the controllers via the app to negotiate their goals and privacy. Drones may also broadcast ultra-frequency sounds (human cannot hear) which encode information about the drones and people's phones or devices can detect these sounds, decode and present these drone information to the citizens/bystanders.

Drone privacy designs should also protect both people's physical privacy and information privacy. The FAA has de-

	Proactive	Reactive	
Drone (Controller)	Signal recording Unique ID Provide drone/recording info Inform/enforce rules Detect sensitive location/info Context-sensitive camera control No fly zone	*Signal ownership/purpose *Broadcast notifications *Privacy-friendly defaults *Communication with citizens *Training/best practices	Encryption Access control Transformation of data Logging and audit (accountability)
Bystander	Request drone/recording info (e.g., what the drone sees)	*Detect drones *Receive notification *Signal opt-out *Communicate to controller	Access to collected data Request data filtering and removal

Fig. 3. Concrete ideas of privacy-enhancing drones for drone controllers and ordinary citizens/bystanders, including both proactive and reactive measures. + denotes suggestions from our informants. * denotes the new suggestions that we propose.

veloped B4UFLY, a mobile app that helps drone controllers “determine whether there are any restrictions or requirements in effect at the location where they want to fly” [19]. Besides, the no-fly zones can help keep drones away from sensitive areas. These mechanisms could help prevent safety and security issues as well as protect people’s physical privacy.

In terms of information privacy, the standard best practices, such as encrypting the content, setting up appropriate access control, redacting sensitive content, logging and auditing would be useful. For instance, drone designs can explore existing access control mechanisms for continuous sensing [54]. We advocate that the drone designers and manufacturers consider incorporating these privacy-enhancing designs in their drone products. While incorporating these privacy features might seem as an increase to the cost, the benefits of making drones more privacy-friendly are also competitive advantages in the drone marketplace.

Since privacy concerns were raised by all of our informants, we suggest that techniques such as facial recognition and sensitive information detection may be incorporated into drones for data obfuscation/filtering purposes. Similar techniques have been introduced by previous studies in the wearable camera context for both bystanders and controllers/owners. For bystanders, Korayem et al. introduced ScreenAvoider, a framework that can help users to manage their privacy by protecting users’ sensitive images and information on computer screen from wearable cameras [38]. Using deep learning techniques, ScreenAvoider can detect and classify sensitive information on computer monitor, then provide users that ability to control the disclosure of these information [38]. For controllers, prior literature shows that controllers are concerned about the bystanders’ privacy [31]. Raval et al. propose PrivacyEye and WaveOff, as their privacy marker system [53]. In this system, sensitive information will be automati-

cally covered by a virtual bounding box from the operators’ device, thus protect the bystanders’ privacy [53]. However, to our knowledge, these techniques have not been adopted by drones. Future work can explore adapting these existing techniques to drones and developing new techniques to address the specific privacy concerns that people have with drones.

We also value Dourish’s perspective that context is not a static representation but is dynamically produced and reproduced in the course of activity [17]. Via our scenario-based questions, we did find that our respondents’ sense of private and public spaces, concerns on privacy, as well as opinions about notification were context-dependent. Taking this perspective, we suggest that automatic location or context detection techniques may be explored for drones. One similar system has been introduced in wearable cameras. Templeman et al. proposed PlaceAvoider, a technique for the wearable cameras to identify the current location [58]. If the current location is considered as sensitive (bedroom, bathroom), the images captured by the cameras will be flagged for further review before made available to other applications. [58] We did not find similar techniques developed for drones.

Considering the extreme mobility of drones, we propose that future work can explore and leverage location and context detection for privacy protections in drones. For instance, drones can implement smart privacy-friendly default settings or privacy-friendly camera modes, such as blurring people’s faces, or abstain from taking pictures / videos in obviously private/personal locations or spaces such as people’s residences. These settings or modes can be applied in recording but also viewing without recording (i.e., controllers can have a live feed of the camera view even when it is not recording).

Since drone technologies are relatively new to the general public, social norms around appropriate use of this technology do not exist. Designers should explore ways to nurture

the formation of these norms. For instance, designs could help strengthen the relationship between citizens and drone controllers so that they can develop trust and expectations for each other. A social platform (e.g., an app) for citizens and drone controllers to meet and mingle could be valuable. For instance, many drone manufacturers already have online forums for drone users/controllers (e.g., forum.dji.com). The manufacturers could extend these platforms into a community that allows drone controllers to provide information about their drones (e.g., where they fly and the purpose of flying/recording), welcome ordinary citizens/bystanders to voice their concerns and support direct communication between controllers and citizens. Similarly, location-based drone picture/video sharing sites can also be extended to help support this type of communication. When best practices and social norms of drones emerge, controllers can be informed or trained about these best practices using educational materials and tools (e.g., games).

Designs for citizens/bystanders. Since it is difficult for citizens to always be able to detect nearby drones and their recording behavior, they should be enabled to pull information about nearby drones. One way to achieve this is to build a database that drone controllers can voluntarily provide information about their drones that citizens can retrieve. If controllers do not provide such information, researchers could also look into ways to help citizens actively detect nearby drones. For instance, the aforementioned mechanism of drones broadcasting information about themselves via sound can be used to allow automatic detection of drones.

In addition, citizens should be able to express their opt out of being recorded by drones. For instance, if users can perform certain pre-defined gesture or the users' devices can broadcast opt-out signals (e.g., color or sound), the drone/camera can potentially capture, interpret and honor the request.

Lastly, citizens should be able to communicate with the drone controllers. For instance, the automatic detection of drones could provide information about the drone controllers (e.g., their email address). The communication platform (e.g., a website or an app) we discussed earlier may also support this communication as well as allowing citizens to request access to the recorded data and request data filtering or deletion. In summary, the key idea is to allow drone controllers and citizens to communicate and negotiate about citizens' privacy.

5.4 Implications for policy

In terms of public policy, both federal regulation and industry self-regulation of drones should take privacy protection as a priority. The FAA drone policy and the code of conduct of drone associations (e.g., Association for Unmanned Vehicle Systems International) barely touch on privacy protection.

However, all informants raised privacy questions without any priming. This finding provides the timely empirical evidence that privacy concerns of drones are real and they need to be addressed.

The FAA launched a required drone controller registration in December 2015 [39]. The registration requires information about a drone controller's name, address, and a credit card, but not any information about the drones that this person owns. But, any registered drone controller should post his or her registration sticker on any drone he or she wants to fly outdoors. By early January 2016 over, 180,000 drone controllers had registered [44]. However, if a citizen is concerned about a drone taking pictures at distance, the citizen is unlikely to see the drone controller information on the sticker/drone and know who the controller is. As such, it is unclear how much this registration enables to protect people's privacy against drones in practice. In February 2016, the FAA announced that they will set up a committee to propose rules to govern how close drones can get to bystanders, mostly for safety reasons [2].

Based on our preliminary review of state-level legislation, 24 states in the US have passed drone-related legislation. We have three important observations. First, there is no consensus on the definition of "drone" among these state drone laws. Some states equate drone with Unmanned Aircraft Vehicle (UAV) such as Oregon, whereas other states separate the two, such as Idaho. Moving forward, a standardized definition is desirable. Second, only some states regulate the drone controllers in addition to the drones. For instance, North Carolina requires drone controllers to pass a knowledge test and get a permit issued by the The Division of Aviation of the Department of Transportation. Given the duality of drone, we advocate for regulations that cover both drones and their controllers. Third, few states have detailed rules on privacy. One exception is Iowa, which has rules on three aspects of privacy: (1) trespassing, (2) invasion of privacy (intrusion upon seclusion, public disclosure of private facts, and sexually motivated privacy invasion), and (3) harassment and stalking. We urge drone laws to include detailed privacy rules.

Some of our informants expected prior consent, however, practically consent would be difficult to implement. Imagine you plan to fly a drone in a park where there are one hundred people. It would be difficult, costly, or even unrealistic to get everybody's permission or consent before flying the drone and/or recording videos in that park. Instead of relying on getting people's prior consent, we suggest considering the ideas of *accountability* and audit. Drone controllers would be held accountable and receive audits for their drone operations.

5.5 Study limitations

Our study is a first step towards a deep understanding of people's perceptions of drones, and it has many limitations.

First, our study scope focused on civilian uses of lightweight drones operated by human controllers. As such, we did not explore military uses of drones. We separated military and civilian uses because they serve distinct purposes, have different implications on society, and thus require separate treatments (e.g., the FAA in the US only regulates civilian use of drones). We also did not study fully autonomous drones (FADs). To our knowledge, FADs are mostly used for military purposes. FADs and military uses may lead to perceptions of drones that are different from what we reported on civilian uses. Since military uses were excluded, our study also did not explore the entanglements between military and civilian uses of drones (e.g., some drone manufacturers have both military and civilian drones).

Second, the list of our scenarios is by no means comprehensive. We chose realistic scenarios that are already happening in the real world because they would be easier for people to understand. All of our scenarios might be perceived as having a "positive" purpose (e.g., searching criminals). We did not have a scenario that has a clearly controversial or "negative" intention (e.g., surveillance or mission creep). Having futuristic and/or "negative" scenarios may solicit different (and presumably more negative) perceptions of drones. In addition, we did not design the scenarios for highlighting the different affordances between drones and other tracking technologies (e.g., camera phones and CCTV). This limitation means that our study may miss some perceived differences between drones and other tracking technologies. Furthermore, each scenario presented a one-off drone use and thus did not highlight the possibility of continuous, repeated or multiple drone uses over an extended period of time. These long-term uses of drones and down-stream data analyses can evoke perceptions that we did not uncover. While some of our scenarios represented organizational uses of drones (e.g., Amazon package delivery, or the police uses drones for searching criminals), our informants seemed more cognizant of individual controllers than what organizations can do with drones and what data they can collect and use over time.

Third, our results are based on a limited sample size and the majority of our informants were university students in the US. This means that our findings might not be generalizable to the general population. University students can be more accepting of new technologies. Therefore, the general population may have even more privacy concerns over drone use than what we reported. In addition, we did not explicitly recruit for informants with varying social-economic status (e.g., minorities, vulnerable populations, or people with low incomes).

People with these backgrounds may have different perceptions or concerns that we did not uncover.

Fourth, while we showed our informants an actual drone, flied it and showed them its live video feed when the weather permitted, it might be still difficult for them to think about this relatively unfamiliar technology. Since most of our informants were not very familiar with drones, they may have been focused more on undesirable aspects (e.g., new technology can bring privacy risks) than their actual perceptions. However, we did ask the perceived benefits of drones at the beginning of the interviews, so our informants were not biased to only consider the risky aspect of drones.

Fifth, we asked our informants to compare drones with camera phones and CCTV, two familiar tracking technologies, as references. We could have included other tracking technologies such as wearable cameras, which may elicit additional insights. However, people are generally less familiar with wearable cameras, making them less ideal as references.

Lastly, we used a specific drone in our study. This might limit our results as other examples of drones may elicit different perceptions.

6 Conclusion

Once a military technology, civilian drones are rapidly moving into the daily lives of people in the US and other countries. Our interview study is a first step towards understanding people's nuanced perceptions of drones. Our informants identified both potential benefits and promising applications of drones, but also safety, security and privacy issues. Our results also suggest that drone is more than just another tracking and recording technology. Its potential for surveillance and impact on people's physical and information privacy is almost unparalleled. The duality of drone implies that, metaphorically, the flying eyes (drones and their controllers) can enter and peek into people's private spaces and lives together. As a result, drone controllers should be held accountable for what they and drones do. Lastly, while the FAA has proposed drone rules to focus primarily on safety and security issues, our study provides timely empirical evidence that people's privacy concerns of drones are real, nuanced, and must be addressed.

7 Acknowledgments

We thank our informants for sharing their insights. We are also grateful to Jason Dedrick, Bryan Semaan, Seda Gürses and anonymous reviewers for their thoughtful feedback on earlier

versions of this paper. This work was supported in part by a Syracuse University internal research grant.

References

- [1] Alessandro Acquisti, Laura Brandimarte, and George Loewenstein. 2015. Privacy and human behavior in the age of information. *Science* 347, 6221 (2015), 509–514.
- [2] Alan Boyle. 2016. FAA panel will decide how drones and people mix. (Feb. 2016). <http://www.geekwire.com/2016/faa-creates-panel-to-come-up-with-rules-for-letting-drones-come-near-innocent-bystanders/>
- [3] Anita Allen. 2015. Privacy and Medicine. In *The Stanford Encyclopedia of Philosophy* (fall 2015 ed.), Edward N. Zalta (Ed.). <http://plato.stanford.edu/archives/fall2015/entries/privacy-medicine/>
- [4] Amazon. 2014. Amazon Prime Air. (2014). <http://www.amazon.com/b?node=8037720011>
- [5] Rebecca Angeles. 2007. An empirical study of the anticipated consumer response to RFID product item tagging. *Industrial Management & Data Systems* 107, 4 (2007), 461–483.
- [6] Melissa Barbee. 2014. Uncharted Territory: The FAA and the Regulation of Privacy via Rulemaking for Domestic Drones. (2014).
- [7] Richard E. Boyatzis. 1998. *Transforming Qualitative Information: Thematic Analysis and Code Development*. SAGE.
- [8] Virginia Braun and Victoria Clarke. 2006. Using thematic analysis in psychology. *Qualitative Research in Psychology* 3, 2 (Jan. 2006), 77–101. DOI: <http://dx.doi.org/10.1191/1478088706qp063oa>
- [9] Daniel J. Butler, Justin Huang, Franziska Roesner, and Maya Cakmak. 2015. The Privacy-Utility Tradeoff for Remotely Teleoperated Robots. In *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction (HRI '15)*. ACM, New York, NY, USA, 27–34. DOI: <http://dx.doi.org/10.1145/2696454.2696484>
- [10] John Travis Butler and Arvin Agah. 2001. Psychological effects of behavior patterns of a mobile personal robot. *Autonomous Robots* 10, 2 (2001), 185–202.
- [11] Ryan Calo. 2011. The drone as privacy catalyst. *Stanford Law Review Online* 64 (2011), 29–33.
- [12] Ann Cavoukian. 2012. *Privacy and drones: Unmanned aerial vehicles*. Information and Privacy Commissioner of Ontario, Canada.
- [13] Reece A Clothier, Dominique A Greer, Duncan G Greer, and Amisha M Mehta. 2015. Risk perception and the public acceptance of drones. *Risk analysis* (2015).
- [14] Consumer Electronics Association (CEA). 2015. *U.S. Consumer Electronics Sales and Forecasts*. Technical Report. Consumer Electronics Association (CEA). <https://www.cta.tech/News/News-Releases/Press-Releases/2015-Press-Releases/New-Tech-to-Drive-CE-Industry-Growth-in-2015,-Proj.aspx>
- [15] Kathleen Bartzan Culver. 2014. From battlefield to newsroom: Ethical implications of drone technology in journalism. *Journal of Mass Media Ethics* 29, 1 (2014), 52–64.
- [16] Tamara Denning, Zakariya Dehlawi, and Tadayoshi Kohno. 2014. In situ with bystanders of augmented reality glasses: Perspectives on recording and privacy-mediating technologies. In *Proceedings of the 32nd annual ACM conference on Human factors in computing systems*. ACM, 2377–2386.
- [17] Paul Dourish. 2004. What we talk about when we talk about context. *Personal and ubiquitous computing* 8, 1 (2004), 19–30.
- [18] Travis Dunlap. 2009. We've got our eyes on you: When surveillance by unmanned aircraft systems constitutes a Fourth Amendment search. *S. Tex. L. Rev.* 51 (2009), 173.
- [19] FAA. 2015a. B4UFLY Smartphone App. (2015). <https://www.faa.gov/uas/b4ufly/>
- [20] FAA. 2015b. Unmanned Aircraft Systems. (2015). <https://www.faa.gov/uas/>
- [21] Adrienne Porter Felt, Serge Egelman, and David Wagner. 2012. I've got 99 problems, but vibration ain't one: a survey of smartphone users' concerns. In *Proceedings of the second ACM workshop on Security and privacy in smartphones and mobile devices*. ACM, 33–44.
- [22] Baruch Fischhoff, Paul Slovic, Sarah Lichtenstein, Stephen Read, and Barbara Combs. 1978. How safe is safe enough? A psychometric study of attitudes towards technological risks and benefits. *Policy Sciences* 9, 2 (April 1978), 127–152. DOI: <http://dx.doi.org/10.1007/BF00143739>
- [23] Erving Goffman. 1959. *The Presentation of Self in Everyday Life* (1 ed.). Anchor.
- [24] Erving Goffman. 1971. Relations in public. microstructure of the public order. *Hannondsworth: Penguin* (1971).
- [25] Edward Twitchell Hall. 1966. The hidden dimension . (1966).
- [26] Steve Hodges, Emma Berry, and Ken Wood. 2011. Sense-Cam: a wearable camera that stimulates and rehabilitates autobiographical memory. *Memory (Hove, England)* 19, 7 (Oct. 2011), 685–696. DOI: <http://dx.doi.org/10.1080/09658211.2011.605591>
- [27] Joachim R Höfllich. 2006. The mobile phone and the dynamic between private and public communication: Results of an international exploratory study. *Knowledge, Technology & Policy* 19, 2 (2006), 58–68.
- [28] Jason Hong. 2013. Considering privacy issues in the context of Google glass. *Commun. ACM* 56, 11 (2013), 10–11.
- [29] Mohammad Alamgir Hossain and Yogesh K Dwivedi. 2014. What improves citizens' privacy perceptions toward RFID technology? A cross-country investigation using mixed method approach. *International Journal of Information Management* 34, 6 (2014), 711–719.
- [30] Roberto Hoyle, Robert Templeman, Denise Anthony, David Crandall, and Apu Kapadia. 2015. Sensitive Lifelogs: A Privacy Analysis of Photos from Wearable Cameras. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. ACM, 1645–1648.
- [31] Roberto Hoyle, Robert Templeman, Steven Armes, Denise Anthony, David Crandall, and Apu Kapadia. 2014. Privacy behaviors of lifeloggers using wearable cameras. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing*. ACM, 571–582.
- [32] Ira Lamcja. 2015. Canada's police forces take to the sky with drones | Metro News. (2015). <http://goo.gl/ITh2Nu> <http://www.metronews.ca/news/canada/2015/03/19/canadas-police-forces-taking-to-the-sky-with-drones.html>

- [33] Jeremy Diamond. 2015. Obama: We need more drone regulations. (2015). <http://www.cnn.com/2015/01/27/politics/obama-drones-fareed/>
- [34] Jonathan Kaiman. 2015. Chinese film star Zhang Ziyi is proposed to – by drone. *The Guardian* (Feb. 2015). <http://goo.gl/kxo4lh> <http://www.theguardian.com/world/2015/feb/09/chinese-film-star-zhang-ziyi-is-proposed-to-by-drone>.
- [35] Jerry Kang. 1998. Information Privacy in Cyberspace Transactions. *Stanford Law Review* 50, 4 (April 1998), 1193–1294. DOI : <http://dx.doi.org/10.2307/1229286>
- [36] Kathy Baxter and Catherine Courage. 2005. *Understanding Your Users: A Practical Guide to User Requirements Methods, Tools, and Techniques* (1 edition ed.). Morgan Kaufmann, San Francisco, CA.
- [37] Yoohwan Kim, Juyeon Jo, and Sanjeeb Shrestha. 2014. A server-based real-time privacy protection scheme against video surveillance by Unmanned Aerial Systems. In *Unmanned Aircraft Systems (ICUAS), 2014 International Conference on*. IEEE, 684–691.
- [38] Mohammed Korayem, Robert Templeman, Dennis Chen, David Crandall, and Apu Kapadia. 2014. Screenavoider: Protecting computer screens from ubiquitous cameras. *arXiv preprint arXiv:1412.0008* (2014).
- [39] Les Dorr and Alison Duquette. 2015. Press Release – Unmanned Aircraft Registration System Takes Flight. (Dec. 2015). http://www.faa.gov/news/press_releases/news_story.cfm?newsId=19874
- [40] LightCense. 2016. LightCense. (2016). <http://www.lightcense.co/>
- [41] Lyn H. Lofland. 1998. *The Public Realm: Exploring the City's Quintessential Social Territory*. Transaction Publishers.
- [42] Aleecia M. McDonald and Lorrie Faith Cranor. 2010. Americans' attitudes about internet behavioral advertising practices. ACM Press, 63. DOI : <http://dx.doi.org/10.1145/1866919.1866929>
- [43] David H. Nguyen and Gillian R. Hayes. 2010. Information Privacy in Institutional and End-user Tracking and Recording Technologies. *Personal Ubiquitous Comput.* 14, 1 (Jan. 2010), 53–72. DOI : <http://dx.doi.org/10.1007/s00779-009-0229-4>
- [44] Jack Nicas. 2016. U.S. Drone Users Number At Least 181,000. (Jan. 2016). <http://blogs.wsj.com/digits/2016/01/06/u-s-drone-users-number-at-least-181000/>
- [45] Helen Nissenbaum. 2004. Privacy as contextual integrity. *Washington law review* 79, 1 (2004), 119–158.
- [46] Helen Nissenbaum. 2009. *Privacy in context: Technology, policy, and the integrity of social life*. Stanford University Press.
- [47] NoFlyZone. 2016. NoFlyZone. (2016). <https://www.noflyzone.org/>
- [48] Wanda J. Orlikowski. 1992. The Duality of Technology: Rethinking the Concept of Technology in Organizations. *Organization Science* 3, 3 (Aug. 1992), 398–427. <http://www.jstor.org/stable/2635280>
- [49] George Orwell. 1949. 1984. Signet Classic.
- [50] Leysia Palen and Paul Dourish. 2003. Unpacking “privacy” for a networked world. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. ACM, Ft. Lauderdale, Florida, USA, 129–136. DOI : <http://dx.doi.org/10.1145/642611.642635>
- [51] Pensacola News Journal. 2015. Drone video: Pensacola Grand Mardi Gras Parade. (2015). <http://www.pnj.com/videos/entertainment/events/mardi-gras/2015/02/15/23444447/>
- [52] Katerina Pramataris and Aristeidis Theotokis. 2009. Consumer acceptance of RFID-enabled services: a model of multiple attitudes, perceived system characteristics and individual traits. *European Journal of Information Systems* 18, 6 (2009), 541–552.
- [53] Nisarg Raval, Landon Cox, Animesh Srivastava, Ashwin Machanavajhala, and Kiron Lebeck. 2014. Markit: privacy markers for protecting visual secrets. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct Publication*. ACM, 1289–1295.
- [54] Franziska Roesner, David Molnar, Alexander Moshchuk, Tadayoshi Kohno, and Helen J Wang. 2014. World-driven access control for continuous sensing. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*. ACM, 1169–1181.
- [55] Bruce Schneier. 2015. Is it OK to shoot down a drone over your house? - CNN.com. (2015). <http://www.cnn.com/2015/09/09/opinions/schneier-shoot-down-drones/index.html>
- [56] H Jeff Smith, Tamara Dinev, and Heng Xu. 2011. Information privacy research: an interdisciplinary review. *MIS quarterly* 35, 4 (2011), 989–1016.
- [57] Daniel J Solove. 2006. A taxonomy of privacy. *University of Pennsylvania law review* (2006), 477–564.
- [58] Robert Templeman, Mohammed Korayem, David Crandall, and Apu Kapadia. 2014. PlaceAvoider: Steering first-person cameras away from sensitive spaces. In *Network and Distributed System Security Symposium (NDSS)*.
- [59] Emily Troshynski, Charlotte Lee, and Paul Dourish. 2008. Accountabilities of presence: reframing location-based systems. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 487–496.
- [60] Janice Tsai, Patrick Kelley, Lorrie Cranor, and Norman Sadeh. 2010. Location-Sharing Technologies: Privacy Risks and Controls. *I/S: A Journal of Law and Policy for the Information Society* 6, 2 (2010), 119–152. <http://ssrn.com/abstract=1997782>
- [61] Blase Ur, Pedro Giovanni Leon, Lorrie Faith Cranor, Richard Shay, and Yang Wang. 2012a. Smart, useful, scary, creepy: perceptions of online behavioral advertising. In *Proceedings of the Eighth Symposium on Usable Privacy and Security (SOUPS '12)*. ACM, New York, NY, USA, 4:1–4:15. DOI : <http://dx.doi.org/10.1145/2335356.2335362>
- [62] Blase Ur, Pedro Giovanni Leon, Lorrie Faith Cranor, Richard Shay, and Yang Wang. 2012b. Smart, useful, scary, creepy: perceptions of online behavioral advertising. In *Proceedings of the Eighth Symposium on Usable Privacy and Security*. ACM, 4–19.
- [63] Samuel Warren and Louis Brandeis. 1890. The Right to Privacy. *Harvard Law Review* 4, 5 (1890), 193–220.
- [64] Alan F. Westin. 2003. Social and Political Dimensions of Privacy. *Journal of Social Issues* 59, 2 (July 2003), 431–453. DOI : <http://dx.doi.org/10.1111/1540-4560.00072>
- [65] David Wright, Rachel Finn, Raphael Gellert, Serge Gutwirth, Philip Schütz, Michael Friedewald, Silvia Venier, and Emilio Mordini. 2014. Ethical dilemma scenarios and emerging

technologies. *Technological Forecasting and Social Change* 87 (2014), 325–336.

- [66] WUSA9. 2015. Drone video goes inside abandoned White Flint Mall. (2015). <http://goo.gl/Akndp2>
<http://bethesda.wusa9.com/news/news/2664031-drone-video-goes-inside-abandoned-white-flint-mall>.

A Interview Questions

A.1 General questions about drones

1. Have you heard of drones? What is the first thing that comes to your mind when you hear about drones?
2. What have you heard about drones?
3. How do you feel about drones?
4. Do you see any benefits or issues of drones?
5. What information do you think drones can collect about you?
6. Did you know that you can record video with drones?
7. Why do you think someone would want to have a drone?
8. How would you compare recording by a drone with recording by a cell phone with a camera? Why?
9. How would you compare recording by a drone with recording by a CCTV camera? Why?
10. How do you feel about being around with a flying drone? Why?
11. Would you want someone who plans to fly a drone near you to ask for your permission before recording a video?

A.2 Context-based questions

12. Are there situations in which you would be more willing to let drone flying round you and recording?
13. For each of the following scenarios, please indicate if you would accept this drone usage. Please explain the reasoning behind your decisions.
 - (a) Imagine you are in a shopping mall where a promotion event is going on, like on the Black Friday. The store owner decides to use a drone to monitor and record this event, and you happen to be in that event.
 - (b) Imagine you are at your friend's party, and your friend decides to use a drone to record the party.
 - (c) Imagine that Amazon decides to use drones to deliver goods that you have bought in its online store to your house.
 - (d) Imagine you are in a parade. Some news agency decides to use a drone to record the parade.

- (e) Imagine that there are some suspects or criminals lurking in your residential area. The police department decides to use drones to search for these people in your residential area.

Do you have any other thought about drones' possible applications?

14. Have any of your expectations changed on drones?
15. Are there any circumstances in which you would NOT like drones to collect data about you?
16. Are you aware of any laws dealing with drones?
17. Do you have any additional comments?

A.3 Expected notification and control

The following questions were inspired by: Drone Aircraft Privacy and Transparency Act of 2013 (proposed but not enacted in the US).

18. Do you expect to have the list of individuals who have the authority to operate or who are operating drones?
19. Do you expect to be notified about the exact locations of the operating drones?
20. Do you expect to be notified about the time periods during which drones can/will be operated?
21. Do you expect to be notified about the types of information that the operating drones might collect?

The following questions were adapted from a RFID user study [29].

22. Do you expect to be asked for any kind of "explicit consent" to allow drones to fly near you? Why?
23. Do you expect to see detailed explanations if a drone takes pictures or videos that can capture you? Why?
24. Do you expect to have any control over your privacy regarding drones operated or owned by others? Why? If you do have such expectations of control, could you give me an example?

Appendix B

Table 12.1: Full list of participants' perceived benefits of smart homes

Group	Participants	Perceived pros/benefits	Perceived cons/risks
1	P1	Convenient (e.g., playing music); cool	Device could go off sometimes (e.g., Amazon Echo started to play music by itself); lack of data control
	P2	Quick access to easy tasks (e.g., turning off living room lights)	Device not functioning properly due to the lack of power (e.g., smart locks would be tough to deal with if the power went off)

P3 Convenient (e.g., checking local weather quickly) Security camera caused significant privacy concerns; no concern if consent was granted, but would have privacy concerns if no consent was provided

2 P4 Providing proof for law enforcement (e.g., Amazon Echo recorded the process of a murder case) The more you put on technology, the more vulnerable you become; lack of regulations on smart home devices usage

P5 Ensuring home safety (e.g., remote access through security camera and record break-ins); convenient; cool Smart home could take over the house and do crazy things (e.g., let burglars come in)

P6 Cool (e.g., interacting with Google Home was very futuristic) Potential data sharing among multiple users of the same devices (e.g., roommate); long term impact on people's acceptance of data collection

3 P7 Convenient (e.g., playing music on Spotify through Amazon Echo) Users' habits were not formed (e.g., she only used Amazon Echo to play music)

P8 Made home more accessible (e.g., proving convenience for people with disability) Could be privacy intrusive for people with disabilities (e.g., blind users could not know the running status of the security camera); lack of awareness in general

P9 Home automation made some tasks easier (e.g., smart coffee maker could be controlled by the phone to make coffee) Security camera could be very intrusive and record every single move; security camera could be hidden

	P10	Convenient (e.g., voice communication with Amazon Echo to check the package status)	Had a long learning curve (e.g., he had to learn how to use the devices through YouTube videos)
4	P11	Ensuring home safety (e.g., making emergency calls through Amazon Echo)	Expensive
	P12	Easy connection with other family members (e.g., using Amazon Show)	Expensive; manufacturers could collect and abuse data
5	P13	Convenient (e.g., making calls through voice assistant); ensuring home safety (e.g., outdoor security camera sending images to her phone)	Device malfunction (e.g., outdoor security cameras got cut off)

	P14	Easy connection with other family members and ensuring in-home safety (e.g., sharing live videos through security camera app so that other family members could see him if he fell)	Privacy concerns (e.g., sharing the videos all the time); Internet connection could be difficult for people who stayed at the senior citizen centers
	P15	Convenient (e.g., let people come into the house using smartphone)	Security risks (e.g., hacking the phone and getting access to the house)
6	P16	Convenient (e.g., playing music)	Privacy concerns (e.g., building profiles)
	P17	Convenient (e.g., making calls, remotely monitoring kids and pets at home)	Lack of trust towards big companies (e.g., skeptical towards Amazon Echo's policy regarding deleting the audios after 45s)

P18	Safer than the Internet; convenient	Expensive; lack of trust towards manufacturers (e.g., smart home devices made by small manufacturers might have less data protection measures)
-----	-------------------------------------	--

Bibliography

- [1] 2016. FAA Drone Registration. (2016). <https://registermyuas.faa.gov/>
- [2] 2016a. Multistakeholder Process: Cybersecurity Vulnerabilities. (2016). <https://www.ntia.doc.gov/other-publication/2016/multistakeholder-process-cybersecurity-vulnerabilities>
- [3] 2016b. Multistakeholder Process: Unmanned Aircraft Systems. (2016). <https://www.ntia.doc.gov/other-publication/2016/multistakeholder-process-unmanned-aircraft-systems>
- [4] 2017. Multistakeholder Process; Internet of Things (IoT) Security Upgradability and Patching. (2017). [ntia.doc.gov/other-publication/2016/multistakeholder-process-iot-security](https://www.ntia.doc.gov/other-publication/2016/multistakeholder-process-iot-security)
- [5] 2018. Did NTIA's Multi-Stakeholder Process Work? Depends On Whom You Ask. (2018). <https://iapp.org/news/a/did-ntias-multi-stakeholder-process-work-depends-whom-you-ask/>
- [6] 2018. Tracking Protection. (2018). <https://support.mozilla.org/en-US/kb/tracking-protection>

- [7] Lisa C. Abrams, Rob Cross, Eric Lesser, and Daniel Z. Levin. 2003. Nurturing interpersonal trust in knowledge-sharing networks. *The Academy of Management Executive* 17, 4 (Nov. 2003), 64–77. DOI:<http://dx.doi.org/10.5465/AME.2003.11851845>
- [8] Alessandro Acquisti, Laura Brandimarte, and George Loewenstein. 2015. Privacy and human behavior in the age of information. *Science* 347, 6221 (2015), 509–514.
- [9] Alessandro Acquisti and Jens Grossklags. 2005. Privacy and Rationality in Individual Decision Making. *IEEE Security & Privacy* 3, 1 (2005), 26–33.
- [10] Anne Adams and Martina Angela Sasse. 1999. Users Are Not the Enemy. *Commun. ACM* 42, 12 (1999), 40–46.
- [11] Muhammad Aleem, Asra Ishtiaq, Salma Hussain Abbasi, and Muhammad Arshad Islam. 2017. User tracking mechanisms and counter measures. *International Journal of Applied Mathematics, Electronics and Computers* 5, 2 (2017), 33–40.
- [12] Hazim Almuhiemedi, Florian Schaub, Norman Sadeh, Idris Adjerid, Alessandro Acquisti, Joshua Gluck, Lorrie Faith Cranor, and Yuvraj Agarwal. 2015. Your location has been shared 5,398 times!: A field study on mobile app privacy nudging. In *Proceedings of the 33rd annual ACM conference on human factors in computing systems*. ACM, 787–796.
- [13] Irwin Altman. 1975. *The Environment and Social Behavior: Privacy, Personal Space, Territory, and Crowding*. (1975).
- [14] Rebecca Angeles. 2007. An empirical study of the anticipated consumer response to RFID product item tagging. *Industrial Management & Data Systems* 107, 4 (2007), 461–483.

- [15] Anirudha Majumdar and Russ Tedrake. 2016. *Funnel Libraries for Real-Time Robust Feedback Motion Planning*. Technical Report. Massachusetts Institute of Technology. http://groups.csail.mit.edu/robotics-center/public_papers/Majumdar16.pdf
- [16] Noah Apthorpe, Dillon Reisman, and Nick Feamster. 2017a. A smart home is no castle: Privacy vulnerabilities of encrypted IoT traffic. *arXiv preprint arXiv:1705.06805* (2017).
- [17] Noah Apthorpe, Dillon Reisman, Srikanth Sundaresan, Arvind Narayanan, and Nick Feamster. 2017b. Spying on the smart home: Privacy attacks and defenses on encrypted iot traffic. *arXiv preprint arXiv:1708.05044* (2017).
- [18] Noah Apthorpe, Yan Shvartzshnaider, Arunesh Mathur, Dillon Reisman, and Nick Feamster. 2018. Discovering Smart Home Internet of Things Privacy Norms Using Contextual Integrity. *arXiv preprint arXiv:1805.06031* (2018).
- [19] Abdullahi Arabo, Ian Brown, and Fadi El-Moussa. 2012. Privacy in the age of mobility and smart devices in smart homes. In *Privacy, Security, Risk and Trust (PASSAT), 2012 International Conference on and 2012 International Confernece on Social Computing (SocialCom)*. IEEE, 819–826.
- [20] Farzaneh Asgharpour, Debin Liu, and L. Jean Camp. 2007. Mental Models of Security Risks. In *Financial Cryptography and Data Security*, Sven Dietrich and Rachna Dhamija (Eds.). Number 4886 in Lecture Notes in Computer Science. Springer Berlin Heidelberg, 367–377.
- [21] Anthony A Atkinson, John H Waterhouse, and Robert B Wells. 1997. A stakeholder approach to strategic performance measurement. *MIT Sloan Management Review* 38, 3 (1997), 25.

- [22] James Auger. 2013. Speculative design: crafting the speculation. *Digital Creativity* 24, 1 (2013), 11–35.
- [23] Sasikanth Avancha, Amit Baxi, and David Kotz. 2012. Privacy in mobile technology for personal healthcare. *ACM Computing Surveys (CSUR)* 45, 1 (2012), 1–54.
- [24] Emmanuel W Ayaburi and Daniel N Treku. 2020. Effect of penitence on social media trust and privacy concerns: The case of Facebook. *International Journal of Information Management* 50 (2020), 171–181.
- [25] Privacy Badger. 2017. Privacy Badger blocks spying ads and invisible trackers. (2017). <https://www EFF.org/privacybadger>
- [26] Natã M Barbosa, Joon S Park, Yaxing Yao, and Yang Wang. 2019. “What if?” Predicting Individual Users’ Smart Home Privacy Preferences and Their Changes. *Proceedings on Privacy Enhancing Technologies* 4 (2019), 1–21.
- [27] Howard Beales. 2010. The value of behavioral targeting. *Network Advertising Initiative* (2010).
- [28] Charles H Bennett, Gilles Brassard, Claude Crépeau, and Ueli M Maurer. 1995. Generalized privacy amplification. *IEEE Transactions on Information Theory* 41, 6 (1995), 1915–1923.
- [29] Andrew Besmer and Heather Richter Lipford. 2010. Moving beyond untagging: photo privacy in a tagged world. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 1563–1572.

- [30] Patrick Biernacki and Dan Waldorf. 1981. Snowball Sampling: Problems and Techniques of Chain Referral Sampling. *Sociological Methods & Research* 10, 2 (Nov. 1981), 141–163.
- [31] Mikhail Bilenko, Matthew Richardson, and Janice Tsai. 2011. Targeted, not tracked: Client-side solutions for privacy-friendly behavioral advertising. In *Privacy Enhancing Technologies Symposium (PETS 2011)*.
- [32] Richard E Boyatzis. 1998. *Transforming qualitative information: Thematic analysis and code development*. sage.
- [33] Danah Boyd. 2010. Social network sites as networked publics: Affordances, dynamics, and implications. In *A networked self*. Routledge, 47–66.
- [34] Virginia Braun and Victoria Clarke. 2006. Using thematic analysis in psychology. *Qualitative Research in Psychology* 3, 2 (Jan. 2006), 77–101. DOI:<http://dx.doi.org/10.1191/1478088706qp063oa>
- [35] Cristian Bravo-Lillo, Lorrie Faith Cranor, Julie Downs, and Saranga Komanduri. 2010. Bridging the gap in computer security warnings: A mental model approach. *IEEE Security & Privacy* 9, 2 (2010), 18–26.
- [36] Pam Briggs and Lisa Thomas. 2015. An Inclusive, Value Sensitive Design Perspective on Future Identity Technologies. *ACM Trans. Comput.-Hum. Interact.* 22, 5, Article 23 (Aug. 2015), 28 pages. DOI:<http://dx.doi.org/10.1145/2778972>
- [37] Rupert Brown and Sam Gaertner (Eds.). 2002. *Blackwell Handbook of Social Psychology: Intergroup Processes*. Wiley-Blackwell, Malden, MA etc.

- [38] Brave Browser. 2017a. Secure, Fast & Private Web Browser. (2017). <https://brave.com/>
- [39] Tor Browser. 2017b. The Tor software protects you by bouncing your communications around a distributed network. (2017). <https://www.torproject.org/projects/torbrowser.html.en>
- [40] AJ Brush, Bongshin Lee, Ratul Mahajan, Sharad Agarwal, Stefan Saroiu, and Colin Dixon. 2011. Home automation in the wild: challenges and opportunities. In *proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 2115–2124.
- [41] Joseph Bugeja, Andreas Jacobsson, and Paul Davidsson. 2016. On privacy and security challenges in smart connected homes. In *Intelligence and Security Informatics Conference (EISIC), 2016 European*. IEEE, 172–175.
- [42] Daniel J. Butler, Justin Huang, Franziska Roesner, and Maya Cakmak. 2015. The Privacy-Utility Tradeoff for Remotely Teleoperated Robots. In *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction (HRI '15)*. ACM, New York, NY, USA, 27–34. DOI:<http://dx.doi.org/10.1145/2696454.2696484>
- [43] M. Ryan Calo. 2011. The Drone as Privacy Catalyst. *Stanford Law Review Online* 64 (Dec. 2011), 29. http://www.stanfordlawreview.org/online/drone-privacy-catalyst?utm_source=publish2&utm_medium=referral&utm_campaign=www.kpbs.org
- [44] L. J. Camp. 2009. Mental models of privacy and security. *IEEE Technology and Society Magazine* 28, 3 (2009), 37–46.

- [45] John Edward Campbell and Matt Carlson. 2002. Panopticon. com: Online surveillance and the commodification of privacy. *Journal of Broadcasting & Electronic Media* 46, 4 (2002), 586–606.
- [46] Jessica R. Cauchard, Jane L. E, Kevin Y. Zhai, and James A. Landay. 2015. Drone & Me: An Exploration into Natural Human-drone Interaction. In *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp '15)*. ACM, New York, NY, USA, 361–365. DOI:<http://dx.doi.org/10.1145/2750858.2805823>
- [47] Antorweep Chakravorty, Tomasz Wlodarczyk, and Chunming Rong. 2013. Privacy preserving data analytics for smart homes. In *Security and Privacy Workshops (SPW), 2013 IEEE*. IEEE, 23–27.
- [48] Farah Chanchary and Sonia Chiasson. 2015. User Perceptions of Sharing, Advertising, and Tracking. In *Proceedings of the Eleventh Symposium On Usable Privacy and Security (SOUPS 2015)*. 53–67.
- [49] Andria Cheng. 2018. Amazon-Marriott Deal Will Make Alexa A Hotel Butler, But The Implications Range Far Wider. (Jun 2018). <https://www.forbes.com/sites/andriacheng/2018/06/19/amazons-marriott-deal-is-way-beyond-alexa-as-your-new-hotel-butler/#22b53db9721e>
- [50] Chia-Fang Chung, Kristin Dew, Allison Cole, Jasmine Zia, James Fogarty, Julie A Kientz, and Sean A Munson. 2016. Boundary negotiating artifacts in personal informatics: Patient-provider

- collaboration with patient-generated data. In *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing*. ACM, 770–786.
- [51] Max E Clarkson. 1995. A stakeholder framework for analyzing and evaluating corporate social performance. *Academy of management review* 20, 1 (1995), 92–117.
- [52] Reece A Clothier, Dominique A Greer, Duncan G Greer, and Amisha M Mehta. 2015. Risk perception and the public acceptance of drones. *Risk analysis* 35, 6 (2015), 1167–1183.
- [53] Julie E Cohen. 2000. Privacy, ideology, and technology: A response to Jeffrey Rosen. *Geo. LJ* 89 (2000), 2029.
- [54] Federal Trade Commission. 2010. *Protecting Consumer Privacy in an Era of Rapid Change*. Technical Report. www.ftc.gov/os/2010/12/101201privacyreport.pdf
- [55] Federal Trade Commission and others. 2009. FTC staff report: Self-regulatory principles for online behavioral advertising, 2009. *Federal Trade Commission, Washington, DC* (2009).
- [56] Federal Trade Commission and others. 2015. Internet of Things: Privacy & security in a connected world. *Washington, DC: Federal Trade Commission* (2015).
- [57] K. J. W. Craik. 1967. *The Nature of Explanation*. Cambridge University Press.
- [58] Lorrie Faith Cranor. 2000. Agents of choice: Tools that facilitate notice and choice about web site data practices. *arXiv preprint cs/0001011* (2000).
- [59] Lorrie Faith Cranor. 2003. P3P: Making privacy policies more useful. *IEEE Security & Privacy* 99, 6 (2003), 50–55.

- [60] Lorrie Faith Cranor. 2012. Necessary but not sufficient: Standardized mechanisms for privacy notice and choice. *J. on Telecomm. & High Tech. L.* 10 (2012), 273.
- [61] Mary J. Culnan and Pamela K. Armstrong. 1999. Information Privacy Concerns, Procedural Fairness, and Impersonal Trust: An Empirical Investigation. *Organization Science* 10, 1 (Jan. 1999), 104–115. DOI:<http://dx.doi.org/10.1287/orsc.10.1.104>
- [62] Mary J Culnan and Robert J Bies. 2003. Consumer privacy: Balancing economic and justice considerations. *Journal of social issues* 59, 2 (2003), 323–342.
- [63] Anupam Das, Martin Degeling, Daniel Smullen, and Norman Sadeh. 2018. Personalized Privacy Assistants for the Internet of Things. (2018).
- [64] Trisha Datta, Noah Apthorpe, and Nick Feamster. 2018. A Developer-Friendly Library for Smart Home IoT Privacy-Preserving Traffic Obfuscation. In *Proceedings of the 2018 Workshop on IoT Security and Privacy*. ACM, 43–48.
- [65] Simon G Davies. 1997. Re-engineering the right to privacy: how privacy has been transformed from a right to a commodity. In *Technology and privacy*. MIT Press, 143–165.
- [66] Tamara Denning, Zakariya Dehlawi, and Tadayoshi Kohno. 2014. In situ with bystanders of augmented reality glasses: Perspectives on recording and privacy-mediating technologies. In *Proceedings of the 32nd annual ACM conference on Human factors in computing systems*. ACM, 2377–2386.
- [67] Disconnect. 2017. Used by over a million people - Disconnect lets you visualize & block the invisible websites that track you. (2017). <https://disconnect.me/disconnect>

- [68] Emily Dixon. 2019. Family finds hidden camera livestreaming from their Airbnb in Ireland. (Apr 2019). <https://www.cnn.com/2019/04/05/europe/ireland-airbnb-hidden-camera-scli-intl/index.html>
- [69] DJI. 2015. DJI Introduces New Geofencing System for its Drones. (2015). <http://www.dji.com/newsroom/news/dji-fly-safe-system>
- [70] Ali Dorri, Salil S Kanhere, Raja Jurdak, and Praveen Gauravaram. 2017. Blockchain for IoT security and privacy: The case study of a smart home. In *Pervasive Computing and Communications Workshops (PerCom Workshops), 2017 IEEE International Conference on*. IEEE, 618–623.
- [71] DuckDuckGo. 2017. The search engine that doesn't track you. (2017). <https://duckduckgo.com>
- [72] Travis Dunlap. 2009. We've got our eyes on you: When surveillance by unmanned aircraft systems constitutes a Fourth Amendment search. *S. Tex. L. Rev.* 51 (2009), 173.
- [73] Anthony Dunne. 2005. Hertzian Tales: Electronic Products, Aesthetic Experience, and Critical Design. (2005).
- [74] Anthony Dunne and Fiona Raby. 2013. *Speculative everything: design, fiction, and social dreaming*. MIT press.
- [75] Janna Lynn Dupree, Richard Devries, Daniel M. Berry, and Edward Lank. 2016. Privacy Personas: Clustering Users via Attitudes and Behaviors Toward Security Practices. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16)*. ACM, New York, NY, USA, 5228–5239.

- [76] Mary Ann Eastlick, Sherry L. Lotz, and Patricia Warrington. 2006. Understanding online B-to-C relationships: An integrated model of privacy concerns, trust, and commitment. *Journal of Business Research* 59, 8 (Aug. 2006), 877–886. DOI:<http://dx.doi.org/10.1016/j.jbusres.2006.02.006>
- [77] Peter Eckersley. 2010. How unique is your web browser?. In *International Symposium on Privacy Enhancing Technologies Symposium*. Springer, 1–18.
- [78] David Eckhoff and Isabel Wagner. 2017. Privacy in the smart city—Applications, technologies, challenges, and solutions. *IEEE Communications Surveys & Tutorials* 20, 1 (2017), 489–516.
- [79] Serge Egelman, AJ Brush, and Kori M Inkpen. 2008. Family accounts: a new paradigm for user accounts within the home environment. In *Proceedings of the 2008 ACM conference on Computer supported cooperative work*. ACM, 669–678.
- [80] Electronic Privacy Information Center (EPIC). 2016. EPIC - Domestic Unmanned Aerial Vehicles (UAVs) and Drones. (2016). <https://epic.org/privacy/drones/>
- [81] Steven Englehardt, Chris Eubank, Peter Zimmerman, Dillon Reisman, and Arvind Narayanan. 2015. OpenWPM: An automated platform for web privacy measurement. Tech Report. Princeton University. (2015).
- [82] Federal Aviation Administration (FAA). 2015a. B4UFLY Smartphone App. (2015). <https://www.faa.gov/uas/b4ufly/>
- [83] Federal Aviation Administration (FAA). 2015b. Unmanned Aircraft Systems. (2015). <https://www.faa.gov/uas/>

- [84] Federal Aviation Administration (FAA). 2016. *Summary of the Small UAS Rule*. Technical Report. https://www.faa.gov/uas/media/Part_107_Summary.pdf
- [85] David Ferraiolo, D Richard Kuhn, and Ramaswamy Chandramouli. 2003. *Role-based access control*. Artech House.
- [86] Joseph L Fleiss, Bruce Levin, and Myunghee Cho Paik. 2013. *Statistical methods for rates and proportions*. John Wiley & Sons.
- [87] Geoffrey A. Fowler. 2015. Talking Toys Are Getting Smarter: Should We Be Worried? (Dec 2015). <https://www.wsj.com/articles/talking-toys-are-getting-smarter-should-we-be-worried-1450378215>
- [88] R Edward Freeman. 2010. *Strategic management: A stakeholder approach*. Cambridge university press.
- [89] R Edward Freeman, Jeffrey S Harrison, Andrew C Wicks, Bidhan L Parmar, and Simone De Colle. 2010. *Stakeholder theory: The state of the art*. Cambridge University Press.
- [90] R Edward Freeman and John McVea. 2001. A stakeholder approach to strategic management. *The Blackwell handbook of strategic management* (2001), 189–207.
- [91] Batya Friedman, Peter H Kahn, and Alan Borning. 2008. Value sensitive design and information systems. *The handbook of information and computer ethics* (2008), 69–101.
- [92] David Frohlich and Robert Kraut. 2003. The social context of home computing. In *Inside the smart home*. Springer, 127–162.

- [93] Andrea Frome, German Cheung, Ahmad Abdulkader, Marco Zennaro, Bo Wu, Alessandro Bissacco, Hartwig Adam, Hartmut Neven, and Luc Vincent. 2009. Large-scale privacy protection in Google Street View. In *2009 IEEE 12th International Conference on Computer Vision*. 2373–2380. DOI:<http://dx.doi.org/10.1109/ICCV.2009.5459413>
- [94] Sidney Fussell. 2019. Airbnb Has a Hidden-Camera Problem. (March 2019). <https://www.theatlantic.com/technology/archive/2019/03/what-happens-when-you-find-cameras-your-airbnb/585007/>
- [95] Future of Privacy Forum, Intel, and PrecisionHawk. 2016. *Drones and Privacy by Design: Embedding Privacy Enhancing Technology in Unmanned Aircraft*. Technical Report. https://fpf.org/wp-content/uploads/2016/08/Drones_and_Privacy_by_Design_FPF_Intel_PrecisionHawk.pdf
- [96] Alan Garnham and Jane Oakhill. 1996. *Mental Models In Cognitive Science: Essays In Honour Of Phil Johnson-Laird*. Psychology Press.
- [97] Simon Gerber, Michael Fry, Judy Kay, Bob Kummerfeld, Glen Pink, and Rainer Wasinger. 2010. PersonisJ: Mobile, Client-Side User Modelling. In *International Conference on User Modeling, Adaptation, and Personalization (Lecture Notes in Computer Science)*, Vol. 6075. Springer Berlin / Heidelberg, 111–122. http://dx.doi.org/10.1007/978-3-642-13470-8_12
- [98] Ghostery. 2017. Ghostery detects and blocks tracking technologies to speed up page loads, eliminate clutter, and protect your data. (2017). <https://www.ghostery.com>

- [99] Barney G. Glaser and Anselm L. Strauss. 2006. *The discovery of grounded theory: strategies for qualitative research*. Transaction Publishers.
- [100] Mohamed Grissa, Attila A Yavuz, and Bechir Hamdaoui. 2016. Preserving the location privacy of secondary users in cooperative spectrum sensing. *IEEE Transactions on Information Forensics and Security* 12, 2 (2016), 418–431.
- [101] Jonathan Grudin. 1994. Computer-supported cooperative work: History and focus. *Computer* 27, 5 (1994), 19–26.
- [102] Jeffrey S Harrison and R Edward Freeman. 1999. Stakeholders, social responsibility, and performance: Empirical evidence and theoretical perspectives. *Academy of management Journal* 42, 5 (1999), 479–485.
- [103] Yangyang He, Paritosh Bahirat, Bart P Knijnenburg, and Abhilash Menon. 2019. A Data-Driven Approach to Designing for Privacy in Household IoT. *ACM Transactions on Interactive Intelligent Systems (TiiS)* 10, 1 (2019), 1–47.
- [104] Eelco Herder and Boping Zhang. 2019. Unexpected and Unpredictable: Factors That Make Personalized Advertisements Creepy. In *Proceedings of the 23rd International Workshop on Personalization and Recommendation on the Web and Beyond*. 1–6.
- [105] Steve Hodges, Emma Berry, and Ken Wood. 2011. SenseCam: a wearable camera that stimulates and rehabilitates autobiographical memory. *Memory (Hove, England)* 19, 7 (Oct. 2011), 685–696. DOI:<http://dx.doi.org/10.1080/09658211.2011.605591>
- [106] Jason Hong. 2013. Considering privacy issues in the context of Google glass. *Commun. ACM* 56, 11 (2013), 10–11.

- [107] Roberto Hoyle, Robert Templeman, Denise Anthony, David Crandall, and Apu Kapadia. 2015. Sensitive Lifelogs: A Privacy Analysis of Photos from Wearable Cameras. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. ACM, 1645–1648.
- [108] Roberto Hoyle, Robert Templeman, Steven Armes, Denise Anthony, David Crandall, and Apu Kapadia. 2014. Privacy behaviors of lifeloggers using wearable cameras. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing*. ACM, 571–582.
- [109] Andreas Jacobsson, Martin Boldt, and Bengt Carlsson. 2016. A risk analysis of a smart home automation system. *Future Generation Computer Systems* 56 (2016), 719–733.
- [110] Andreas Jacobsson and Paul Davidsson. 2015. Towards a model of privacy and security for smart homes. In *Internet of Things (WF-IoT), 2015 IEEE 2nd World Forum on*. IEEE, 727–732.
- [111] Timo Jakobi, Sameer Patil, Dave Randall, Gunnar Stevens, and Volker Wulf. 2019. It Is About What They Could Do with the Data: A User Perspective on Privacy in Smart Metering. *ACM Transactions on Computer-Human Interaction (TOCHI)* 26, 1 (2019), 2.
- [112] Timo Jakobi, Gunnar Stevens, Nico Castelli, Corinna Ogonowski, Florian Schaub, Nils Vindice, Dave Randall, Peter Tolmie, and Volker Wulf. 2018. Evolving Needs in IoT Control and Accountability: A Longitudinal Study on Smart Home Intelligibility. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 2, 4 (2018), 171.

- [113] Michael C Jensen. 2002. Value maximization, stakeholder theory, and the corporate objective function. *Business ethics quarterly* (2002), 235–256.
- [114] Jeremy Diamond. 2015. Obama: We need more drone regulations. (2015). <http://www.cnn.com/2015/01/27/politics/obama-drones-fareed/>
- [115] Li Jiang, Da-You Liu, and Bo Yang. 2004. Smart home research. In *Machine Learning and Cybernetics, 2004. Proceedings of 2004 International Conference on*, Vol. 2. IEEE, 659–663.
- [116] Philip Johnson-Laird, Vittorio Girotto, and Paolo Legrenzi. 1998. Mental models: a gentle guide for outsiders. *Sistemi Intelligenti* 9, 68 (1998), 33.
- [117] Natalie A. Jones, Helen Ross, Timothy Lynam, Pascal Perez, and Anne Leitch. 2011. Mental models: an interdisciplinary synthesis of theory and methods. *Ecology and Society* 16 (March 2011), 1–13.
- [118] Thomas M Jones. 1995. Instrumental stakeholder theory: A synthesis of ethics and economics. *Academy of management review* 20, 2 (1995), 404–437.
- [119] Ruogu Kang, Laura Dabbish, Nathaniel Fruchter, and Sara Kiesler. 2015. “My Data Just Goes Everywhere:” User Mental Models of the Internet and Implications for Privacy and Security. In *Eleventh Symposium On Usable Privacy and Security (SOUPS 2015)*. 39–52.
- [120] Yoohwan Kim, Juyeon Jo, and Sanjeeb Shrestha. 2014. A server-based real-time privacy protection scheme against video surveillance by Unmanned Aerial Systems. In *Unmanned Aircraft Systems (ICUAS), 2014 International Conference on*. IEEE, 684–691.

- [121] Paul Kinlan. 2019. User Location. (2019). <https://developers.google.com/web/fundamentals/native-hardware/user-location/>
- [122] Bart P. Knijnenburg, Alfred Kobsa, and Hongxia Jin. 2013. Dimensionality of information disclosure behavior. *International Journal of Human-Computer Studies* 71, 12 (2013), 1144–1162.
- [123] Bastian Könings, Sebastian Thoma, Florian Schaub, and Michael Weber. 2014. Pripref broadcaster: Enabling users to broadcast privacy preferences in their physical proximity. In *Proceedings of the 13th International Conference on Mobile and Ubiquitous Multimedia*. ACM, 133–142.
- [124] Martin J Kraemer and Ivan Flechais. 2018. Researching privacy in smart homes: A roadmap of future directions and research methods. (2018).
- [125] Sathish Alampalayam Kumar, Tyler Vealey, and Harshit Srivastava. 2016. Security in internet of things: Challenges, solutions and future directions. In *System Sciences (HICSS), 2016 49th Hawaii International Conference on*. IEEE, 5772–5781.
- [126] Ponnurangam Kumaraguru, Yong Rhee, Alessandro Acquisti, Lorrie Faith Cranor, Jason Hong, and Elizabeth Nunge. 2007. Protecting people from phishing: the design and evaluation of an embedded training email system. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. ACM, 905–914.
- [127] Patricia G Lange. 2007. Publicly private and privately public: Social networking on YouTube. *Journal of computer-mediated communication* 13, 1 (2007), 361–380.

- [128] Marc Langheinrich. 2002. A privacy awareness system for ubiquitous computing environments. In *international conference on Ubiquitous Computing*. Springer, 237–245.
- [129] Josephine Lau, Benjamin Zimmerman, and Florian Schaub. 2018. Alexa, are you listening? privacy perceptions, concerns and privacy-seeking behaviors with smart speakers. In *Pro. ACM Human-Computer Interaction CSCW*. ACM.
- [130] Robert S Laufer and Maxine Wolfe. 1977. Privacy as a concept and a social issue: A multidimensional developmental theory. *Journal of social Issues* 33, 3 (1977), 22–42.
- [131] Mathias Lecuyer, Riley Spahn, Yannis Spiliopolous, Augustin Chaintreau, Roxana Geambasu, and Daniel Hsu. 2015. Sunlight: Fine-grained Targeting Detection at Scale with Statistical Confidence. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*. ACM, 554–566.
- [132] Pedro Leon, Blase Ur, Richard Shay, Yang Wang, Rebecca Balebako, and Lorrie Cranor. 2012. Why Johnny can’t opt out: a usability evaluation of tools to limit online behavioral advertising. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 589–598.
- [133] Pedro Giovanni Leon, Blase Ur, Yang Wang, Manya Sleeper, Rebecca Balebako, Richard Shay, Lujo Bauer, Mihai Christodorescu, and Lorrie Faith Cranor. 2013. What matters to users?: factors that affect users’ willingness to share information with online advertisers. In *Proceedings of the Symposium on Usable Privacy and Security*. ACM, 7–26.
- [134] Lawrence Lessig. 1997. What things regulate speech: CDA 2.0 vs. filtering. *Jurimetrics* 38 (1997), 629.

- [135] Yuting Liao, Jessica Vitak, Priya Kumar, Michael Zimmer, and Katherine Kritikos. 2019. Understanding the Role of Privacy and Trust in Intelligent Personal Assistant Adoption. In *International Conference on Information*. Springer, 102–113.
- [136] LightCense. 2016. LightCense. (2016). <http://www.lightcense.co/>
- [137] Huichen Lin and Neil W Bergmann. 2016. IoT privacy and security challenges for smart home environments. *Information* 7, 3 (2016), 44.
- [138] Jialiu Lin, Shahriyar Amini, Jason I. Hong, Norman Sadeh, Janne Lindqvist, and Joy Zhang. 2012. 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)*. ACM, New York, NY, USA, 501–510.
- [139] Joseph Lindley and Paul Coulton. 2015. Game of drones. In *Proceedings of the 2015 annual symposium on computer-human interaction in play*. ACM, 613–618.
- [140] Joseph Lindley, Paul Coulton, and Miriam Sturdee. 2017b. Implications for adoption. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. ACM, 265–277.
- [141] Joseph Galen Lindley, Paul Coulton, Haider Akmal, and Brandin Hanson Knowles. 2017a. Anticipating GDPR in Smart Homes Through Fictional Conversational Objects. (2017).
- [142] John Lofland, David A. Snow, Leon Anderson, and Lyn H. Lofland. 2005. *Analyzing Social Settings: A Guide to Qualitative Observation and Analysis* (4 edition ed.). Cengage Learning, Belmont, CA.

- [143] Nathan Malkin, Julia Bernd, Maritza Johnson, and Serge Egelman. 2018. “What Can’t Data Be Used For?” Privacy Expectations about Smart TVs in the US. *European Workshop on Usable Security (EuroUSEC)* (2018).
- [144] Nathan Malkin, Joe Deatruck, Allen Tong, Primal Wijesekera, Serge Egelman, and David Wagner. 2019. Privacy Attitudes of Smart Speaker Users. *Proceedings on Privacy Enhancing Technologies* 2019, 4 (2019), 250–271.
- [145] Stephen T Margulis. 1977. Conceptions of privacy: Current status and next steps. *Journal of Social Issues* 33, 3 (1977), 5–21.
- [146] Jonathan R Mayer and John C Mitchell. 2012. Third-party web tracking: Policy and technology. In *Security and Privacy (SP), 2012 IEEE Symposium on*. IEEE, 413–427.
- [147] Aleecia McDonald and Lorrie Faith Cranor. 2010. Beliefs and behaviors: Internet users’ understanding of behavioral advertising. Research Conference on Communications, Information and Internet Policy (TPRC).
- [148] Roisin McNaney, John Vines, Jamie Mercer, Leon Mexter, Daniel Welsh, and Tony Young. 2017. DemYouth: Co-Designing and Enacting Tools to Support Young People’s Engagement with People with Dementia. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. ACM, 1313–1325.
- [149] Emily McReynolds, Sarah Hubbard, Timothy Lau, Aditya Saraf, Maya Cakmak, and Franziska Roesner. 2017. Toys that listen: A study of parents, children, and internet-connected toys. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. ACM, 5197–5207.

- [150] William Melicher, Mahmood Sharif, Joshua Tan, Lujo Bauer, Mihai Christodorescu, and Pedro Giovanni Leon. 2016. (Do Not) Track Me Sometimes: Users' Contextual Preferences for Web Tracking. *Proceedings on Privacy Enhancing Technologies* 2016, 2 (2016), 135–154.
- [151] Tiago DP Mendes, Radu Godina, Eduardo MG Rodrigues, João CO Matias, and João PS Catalão. 2015. Smart home communication technologies and applications: Wireless protocol assessment for home area network resources. *Energies* 8, 7 (2015), 7279–7311.
- [152] Andrew Meola. 2016. How the Internet of Things will affect security & privacy. (2016).
- [153] Georg Merzdovnik, Markus Huber, Damjan Buhov, Nick Nikiforakis, Sebastian Neuner, Martin Schmiedecker, and Edgar Weippl. 2017. Block me if you can: A large-scale study of tracker-blocking tools. In *Security and Privacy (EuroS&P), 2017 IEEE European Symposium on*. IEEE, 319–333.
- [154] Abraham H Mhaidli, Yixin Zou, and Florian Schaub. 2019. "We Can't Live Without Them!" App Developers' Adoption of Ad Networks and Their Considerations of Consumer Risks. In *Fifteenth Symposium on Usable Privacy and Security (SOUPS) 2019*.
- [155] Ronald K Mitchell, Bradley R Agle, and Donna J Wood. 1997. Toward a theory of stakeholder identification and salience: Defining the principle of who and what really counts. *Academy of management review* 22, 4 (1997), 853–886.
- [156] Simon Moncrieff, Svetha Venkatesh, and Geoff West. 2007. Dynamic privacy in a smart house environment. In *Multimedia and Expo, 2007 IEEE International Conference on*. IEEE, 2034–2037.

- [157] Deirdre K. Mulligan, Colin Koopman, and Nick Doty. 2016. Privacy is an essentially contested concept: a multi-dimensional analytic for mapping privacy. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 374, 2083 (dec 2016). DOI:<http://dx.doi.org/10.1098/rsta.2016.0118>
- [158] Alena Naiakshina, Anastasia Danilova, Sergej Dechand, Kat Krol, M. Angela Sasse, and Matthew Smith. 2016. Poster: Mental Models-User understanding of messaging and encryption. In *Proceedings of European Symposium on Security and Privacy*. <http://www.ieee-security.org/TC/EuroSP2016/posters/number18.pdf>
- [159] Lisa P Nathan, Batya Friedman, Predrag Klasnja, Shaun K Kane, and Jessica K Miller. 2008. Envisioning systemic effects on persons and society throughout interactive system design. In *Proceedings of the 7th ACM conference on Designing interactive systems*. ACM, 1–10.
- [160] Lisa P Nathan, Predrag V Klasnja, and Batya Friedman. 2007. Value scenarios: a technique for envisioning systemic effects of new technologies. In *CHI'07 extended abstracts on Human factors in computing systems*. ACM, 2585–2590.
- [161] National Telecommunications and Information Administration. 2016. *Voluntary Best Practices for UAS Privacy, Transparency, and Accountability*. Technical Report. https://www.ntia.doc.gov/files/ntia/publications/voluntary_best_practices_for_uas_privacy_transparency_and_accountability_0.pdf
- [162] Nest. 2018. Nest Cam. (2018). <https://nest.com/cameras/nest-cam-indoor/overview/>

- [163] David H. Nguyen and Gillian R. Hayes. 2010. Information Privacy in Institutional and End-user Tracking and Recording Technologies. *Personal Ubiquitous Comput.* 14, 1 (Jan. 2010), 53–72. DOI:<http://dx.doi.org/10.1007/s00779-009-0229-4>
- [164] Helen Nissenbaum. 2004. Privacy as contextual integrity. *Washington law review* 79, 1 (2004), 119–158.
- [165] Helen Nissenbaum. 2009. *Privacy in context: Technology, policy, and the integrity of social life*. Stanford University Press.
- [166] Helen Nissenbaum. 2011. A contextual approach to privacy online. *Daedalus* 140, 4 (2011), 32–48.
- [167] NoFlyZone. 2016. NoFlyZone. (2016). <https://www.noflyzone.org/>
- [168] Anthony J. Onwuegbuzie and Nancy L. Leech. 2006. Validity and Qualitative Research: An Oxymoron? *Quality & Quantity* 41, 2 (May 2006), 233–249.
- [169] Siani Pearson. 2013. Privacy, security and trust in cloud computing. In *Privacy and security for cloud computing*. Springer, 3–42.
- [170] James Pierce, Sarah Fox, Nick Merrill, Richmond Wong, and Carl DiSalvo. 2018. An Interface without A User: An Exploratory Design Study of Online Privacy Policies and Digital Legalese. In *Proceedings of the 2018 on Designing Interactive Systems Conference 2018*. ACM, 1345–1358.
- [171] Adblock Plus. 2017. Surf the web without annoying ads! (2017). <https://adblockplus.org>

- [172] Request Policy. 2017. RequestPolicy is an extension for Mozilla browsers that increases your browsing privacy, security, and speed by giving you control over cross-site requests on pages you visit. (2017). <https://requestpolicycontinued.github.io>
- [173] Erika Shehan Poole, Marshini Chetty, Rebecca E Grinter, and W Keith Edwards. 2008a. More than meets the eye: transforming the user experience of home network management. In *Proceedings of the 7th ACM conference on Designing interactive systems*. ACM, 455–464.
- [174] Erika Shehan Poole, Christopher A Le Dantec, James R Eagan, and W Keith Edwards. 2008b. Reflecting on the invisible: understanding end-user perceptions of ubiquitous computing. In *Proceedings of the 10th international conference on Ubiquitous computing*. ACM, 192–201.
- [175] Kristin Purcell, Joanna Brenner, and Lee Rainie. 2012. Search engine use 2012. Pew Internet & American Life Project. (2012).
- [176] Zhengrui Qin, Shanhe Yi, Qun Li, and Dmitry Zamkov. 2014. Preserving secondary users' privacy in cognitive radio networks. In *IEEE INFOCOM 2014-IEEE Conference on Computer Communications*. IEEE, 772–780.
- [177] Emilee Rader. 2014. Awareness of Behavioral Tracking and Information Privacy Concern in Facebook and Google. In *Symposium on Usable Privacy and Security (SOUPS)*. 51–67.
- [178] Matt Ratto. 2011. Critical Making: Conceptual and Material Studies in Technology and Social Life. *The Information Society* 27, 4 (2011), 252–260. DOI:<http://dx.doi.org/10.1080/01972243.2011.583819>

- [179] Ian Reay, Scott Dick, and James Miller. 2009. A large-scale empirical study of P3P privacy policies: Stated actions vs. legal obligations. *ACM Transactions on the Web (TWEB)* 3, 2 (2009), 6.
- [180] Jingjing Ren, Ashwin Rao, Martina Lindorfer, Arnaud Legout, and David Choffnes. 2015. Recon: Revealing and controlling privacy leaks in mobile network traffic. *arXiv preprint arXiv:1507.00255* (2015).
- [181] Robin W Roberts. 1992. Determinants of corporate social responsibility disclosure: An application of stakeholder theory. *Accounting, organizations and society* 17, 6 (1992), 595–612.
- [182] Franziska Roesner, Tadayoshi Kohno, and David Wetherall. 2012. Detecting and defending against third-party tracking on the web. In *Proceedings of the 9th USENIX conference on Networked Systems Design and Implementation*. USENIX Association, 12–12.
- [183] Franziska Roesner, David Molnar, Alexander Moshchuk, Tadayoshi Kohno, and Helen J Wang. 2014. World-driven access control for continuous sensing. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*. ACM, 1169–1181.
- [184] Timothy J Rowley. 1997. Moving beyond dyadic ties: A network theory of stakeholder influences. *Academy of management Review* 22, 4 (1997), 887–910.
- [185] Roy G. D’Andrade. 1995. *The Development of Cognitive Anthropology*. Cambridge University Press.

- [186] Iskander Sanchez-Rola, Xabier Ugarte-Pedrero, Igor Santos, and Pablo G Bringas. 2017. The web is watching you: A comprehensive review of web-tracking techniques and countermeasures. *Logic Journal of the IGPL* 25, 1 (2017), 18–29.
- [187] Elizabeth B-N Sanders and Pieter Jan Stappers. 2008. Co-creation and the new landscapes of design. *Co-design* 4, 1 (2008), 5–18.
- [188] Ravi S Sandhu, Edward J Coyne, Hal L Feinstein, and Charles E Youman. 1996. Role-based access control models. *Computer* 29, 2 (1996), 38–47.
- [189] Susanne G Scott and Vicki R Lane. 2000. A stakeholder approach to organizational identity. *Academy of Management review* 25, 1 (2000), 43–62.
- [190] Phoebe Sengers, Kirsten Boehner, Shay David, and Joseph 'Jofish' Kaye. 2005. Reflective Design. In *Proceedings of the 4th Decennial Conference on Critical Computing: Between Sense and Sensibility (CC '05)*. ACM, New York, NY, USA, 49–58. DOI:<http://dx.doi.org/10.1145/1094562.1094569>
- [191] Katie Shilton. 2009. Four billion little brothers? Privacy, mobile phones, and ubiquitous data collection. *Commun. ACM* 52, 11 (2009), 48–53.
- [192] H Jeff Smith, Tamara Dinev, and Heng Xu. 2011. Information privacy research: an interdisciplinary review. *MIS quarterly* 35, 4 (2011), 989–1016.
- [193] Daniel J Solove. 2005. A taxonomy of privacy. *U. Pa. L. Rev.* 154 (2005), 477.
- [194] Sarah Spiekermann, Jens Grossklags, and Bettina Berendt. 2001. E-privacy in 2Nd Generation E-commerce: Privacy Preferences Versus Actual Behavior. In *Proceedings of the 3rd ACM*

- Conference on Electronic Commerce (EC '01)*. ACM, New York, NY, USA, 38–47. DOI :
<http://dx.doi.org/10.1145/501158.501163>
- [195] Marc Steen, Menno Manschot, and Nicole De Koning. 2011. Benefits of co-design in service design projects. *International Journal of Design* 5, 2 (2011).
- [196] Steven Englehardt and Arvind Narayanan. 2016. *Online tracking: A 1-million-site measurement and analysis*. Technical Report. Princeton University.
http://randomwalker.info/publications/OpenWPM_1_million_site_tracking_measurement.pdf
- [197] Biljana L Risteska Stojkoska and Kire V Trivodaliev. 2017. A review of Internet of Things for smart home: Challenges and solutions. *Journal of Cleaner Production* 140 (2017), 1454–1464.
- [198] Dianna L Stone and Eugene F Stone-Romero. 1998. A multiple stakeholder model of privacy in organizations. In *Managerial Ethics*. Psychology Press, 45–70.
- [199] William H Stufflebeam, Annie I Antón, Qingfeng He, and Neha Jain. 2004. Specifying privacy policies with P3P and EPAL: lessons learned. In *Proceedings of the 2004 ACM workshop on Privacy in the electronic society*. ACM, 35–35.
- [200] Andrew Thatcher and Mike Greyling. 1998. Mental models of the Internet. *International journal of industrial ergonomics* 22, 4 (1998), 299–305.
- [201] The New York Times. 2018. Is Alexa Listening? (2018). <https://www.nytimes.com/2018/05/25/business/amazon-alexa-conversation-shared-echo.html>

- [202] Eran Toch, Yang Wang, and Lorrie Faith Cranor. 2012. Personalization and privacy: a survey of privacy risks and remedies in personalization-based systems. *User Modeling and User-Adapted Interaction* 22, 1-2 (2012), 203–220.
- [203] Stefano Traverso, Martino Trevisan, Leonardo Giannantoni, Marco Mellia, and Hassan Metwalley. 2017. Benchmark and comparison of tracker-blockers: Should you trust them?. In *Network Traffic Measurement and Analysis Conference (TMA), 2017*. IEEE, 1–9.
- [204] Joseph Turow, Jennifer King, Chris Jay Hoofnagle, Amy Bleakley, and Michael Hennessy. 2009. Americans reject tailored advertising and three activities that enable it. *Technical report, Annenberg School for Communications, University of Pennsylvania* (2009). http://repository.upenn.edu/asc_papers/137/
- [205] Blase Ur, Pedro Giovanni Leon, Lorrie Faith Cranor, Richard Shay, and Yang Wang. 2012. Smart, useful, scary, creepy: perceptions of online behavioral advertising. In *Proceedings of the Eighth Symposium on Usable Privacy and Security*. ACM, 4–19.
- [206] Froukje Sleeswijk Visser, Pieter Jan Stappers, Remko Van der Lugt, and Elizabeth BN Sanders. 2005. Contextmapping: experiences from practice. *CoDesign* 1, 2 (2005), 119–149.
- [207] James Vlahos. 2019. Smart talking: are our devices threatening our privacy? (Mar 2019). <https://www.theguardian.com/technology/2019/mar/26/smart-talking-are-our-devices-threatening-our-privacy>
- [208] Junjue Wang, Brandon Amos, Anupam Das, Padmanabhan Pillai, Norman Sadeh, and Mahadev Satyanarayanan. 2018. Enabling Live Video Analytics with a Scalable and Privacy-Aware Framework. *ACM Transactions on Multimedia Computing, Communications, and*

Applications (TOMM) 14, 3s (2018), 64.

- [209] Sijun Wang, Sharon E. Beatty, and William Foxx. 2004. Signaling the trustworthiness of small online retailers. *Journal of Interactive Marketing* 18, 1 (2004), 53–69. DOI: <http://dx.doi.org/10.1002/dir.10071>
- [210] Yang Wang, Pedro Giovanni Leon, Alessandro Acquisti, Lorrie Faith Cranor, Alain Forget, and Norman Sadeh. 2014. A field trial of privacy nudges for facebook. In *Proceedings of the SIGCHI conference on human factors in computing systems*. ACM, 2367–2376.
- [211] Yang Wang, Huichuan Xia, and Yun Huang. 2016. Examining American and Chinese Internet Users' Contextual Privacy Preferences of Behavioral Advertising. In *Proceedings of ACM Conference on Computer-Supported Cooperative Work and Social Computing (CSCW 2016)*. 539–552.
- [212] Samuel D Warren and Louis D Brandeis. 1890. The right to privacy. *Harvard law review* (1890), 193–220.
- [213] Rick Wash. 2010. Folk models of home computer security. In *Proceedings of the Sixth Symposium on Usable Privacy and Security*. ACM, 11:1–11:16.
- [214] Wendy L Weinstein. 1971. The private and the free: A conceptual inquiry. (1971).
- [215] Alan F Westin. 1968. Privacy and freedom. *Washington and Lee Law Review* 25, 1 (1968), 166.
- [216] Alan F. Westin. 2003. Social and Political Dimensions of Privacy. *Journal of Social Issues* 59, 2 (July 2003), 431–453. DOI:<http://dx.doi.org/10.1111/1540-4560.00072>

- [217] Willett Kempton. 1986. Two theories of home heat control. *Cognitive Science* 10 (1986), 75–90.
- [218] Craig E Wills and Doruk C Uzunoglu. 2016. What Ad Blockers Are (and Are Not) Doing. In *Hot Topics in Web Systems and Technologies (HotWeb), 2016 Fourth IEEE Workshop on*. IEEE, 72–77.
- [219] Richmond Y Wong and Vera Khovanskaya. 2018. Speculative Design in HCI: From Corporate Imaginations to Critical Orientations. In *New Directions in Third Wave Human-Computer Interaction: Volume 2-Methodologies*. Springer, 175–202.
- [220] Richmond Y. Wong and Deirdre K. Mulligan. 2019. Bringing Design to the Privacy Table: Broadening ”Design” in ”Privacy by Design” Through the Lens of HCI. In *CHI Conference on Human Factors in Computing Systems (CHI 2019)*. DOI:<http://dx.doi.org/10.1145/3290605.3300492>
- [221] Richmond Y Wong, Deirdre K Mulligan, Ellen Van Wyk, James Pierce, and John Chuang. 2017a. Eliciting values reflections by engaging privacy futures using design workbooks. *Proceedings of the ACM on Human-Computer Interaction* 1, CSCW (2017), 111.
- [222] Richmond Y Wong, Ellen Van Wyk, and James Pierce. 2017b. Real-Fictional Entanglements: Using Science Fiction and Design Fiction to Interrogate Sensing Technologies. In *Proceedings of the 2017 Conference on Designing Interactive Systems*. ACM, 567–579.
- [223] Peter Worthy, Ben Matthews, and Stephen Viller. 2016. Trust me: doubts and concerns living with the Internet of Things. In *Proceedings of the 2016 ACM Conference on Designing Interactive Systems*. ACM, 427–434.

- [224] David Wright and Paul De Hert. 2011. *Privacy impact assessment*. Vol. 6. Springer Science & Business Media.
- [225] David Wright, Rachel Finn, Raphael Gellert, Serge Gutwirth, Philip Schütz, Michael Friedewald, Silvia Venier, and Emilio Mordini. 2014. Ethical dilemma scenarios and emerging technologies. *Technological Forecasting and Social Change* 87 (2014), 325–336.
- [226] Mike Wuerthele. 2018. Here's how Apple protects your privacy in Safari with Intelligent Tracking Protection 2.0. (2018). <https://appleinsider.com/articles/18/06/20/heres-how-apple-protects-your-privacy-in-safari-with-intelligent-tracking-protection-20>
- [227] Heng Xu, Robert E Crossler, and France BéLanger. 2012. A value sensitive design investigation of privacy enhancing tools in web browsers. *Decision support systems* 54, 1 (2012), 424–433.
- [228] Yang Wang, Huichuan Xia, Yaxing Yao, and Yun Huang. 2016. Flying Eyes and Hidden Controllers: A Qualitative Study of People's Privacy Perceptions of Civilian Drones in the US. *Proceedings on Privacy Enhancing Technologies (PoPETS)* 3 (2016), 172–190.
- [229] Yaxing Yao, Justin Reed Basdeo, Smirity Kaushik, and Yang Wang. 2019a. Defending My Castle: A Co-Design Study of Privacy Mechanisms for Smart Homes. (2019).
- [230] Yaxing Yao, Justin Reed Basdeo, Oriana Rosata Mcdonough, and Yang Wang. 2019b. Privacy Perceptions and Designs of Bystanders in Smart Homes. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019), 1–24.

- [231] Yaxing Yao, Davide Lo Re, and Yang Wang. 2017a. Folk Models of Online Behavioral Advertising. In *Proceedings of the Computer Supported Cooperative Work (CSCW)*. ACM.
- [232] Yaxing Yao, Huichuan Xia, Yun Huang, and Yang Wang. 2017b. Free to Fly in Public Spaces: Drone Controllers' Privacy Perceptions and Practices. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. ACM, 6789–6793.
- [233] Yaxing Yao, Huichuan Xia, Yun Huang, and Yang Wang. 2017c. Privacy Mechanisms for Drones: Perceptions of Drone Controllers and Bystanders. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. ACM, 6777–6788.
- [234] Kenji Yoshigoe, Wei Dai, Melissa Abramson, and Alexander Jacobs. 2015. Overcoming invasion of privacy in smart home environment with synthetic packet injection. In *TRON Symposium (TRONSHOW), 2015*. IEEE, 1–7.
- [235] Eric Zeng, Shirang Mare, and Franziska Roesner. 2017. End User Security & Privacy Concerns with Smart Homes. In *Symposium on Usable Privacy and Security (SOUPS)*.
- [236] Bo Zhang and Heng Xu. 2016. Privacy nudges for mobile applications: Effects on the creepiness emotion and privacy attitudes. In *Proceedings of the 19th ACM conference on computer-supported cooperative work & social computing*. ACM, 1676–1690.
- [237] Serena Zheng, Marshini Chetty, and Nick Feamster. 2018. User Perceptions of Privacy in Smart Homes. *arXiv preprint arXiv:1802.08182* (2018).
- [238] Martina Ziefle, Carsten Rocker, and Andreas Holzinger. 2011. Medical technology in smart homes: exploring the user's perspective on privacy, intimacy and trust. In *Computer Software*

and Applications Conference Workshops (COMPSACW), 2011 IEEE 35th Annual. IEEE, 410–415.

[239] Michael Zimmer. 2010. “But the data is already public”: on the ethics of research in Facebook. *Ethics and information technology* 12, 4 (2010), 313–325.

[240] Verena Zimmermann, Merve Bennighof, Miriam Edel, Oliver Hofmann, Judith Jung, and Melina von Wick. 2018. ‘Home, Smart Home’–Exploring End Users’ Mental Models of Smart Homes. *Mensch und Computer 2018-Workshopband* (2018).

Curriculum Vitae

Yaxing Yao, Ph.D.

Post-Doctoral Fellow
Institute for Software Research, School of Computer Science
Carnegie Mellon University, Pittsburgh, PA 15217

Phone: 206.747.8867
Email: yaxingyao@cmu.edu
Website: <http://www.yaxingyao.com>

- RESEARCH INTERESTS** Understanding the privacy implications when users interact with a computing system; Designing for multiple stakeholders' privacy needs; Developing and evaluating various privacy-enhancing tools.
- Keywords:** Human-Computer Interaction; Computer-Supported Cooperative Work; Usable Privacy; Internet of Things; Multi-Stakeholders; Design
- ACADEMIC APPOINTMENT** **Carnegie Mellon University, Pittsburgh, PA** 01/2020 to Now
- Post-Doctoral Fellow (Supervisor: Dr. Norman Sadeh)
- EDUCATION** **Syracuse University, Syracuse, NY** 05/2020
- Ph.D. in Information Science and Technology (Advisor: Dr. Yang Wang)
 - Dissertation: Designing a Privacy Awareness System for the Multi-Stakeholder Environment in Smart Homes
 - Committee: Yang Wang, Jason Dedrick, Joon S. Park, Bryan Semaan, Josh Introne, Florian Schaub (University of Michigan)
- University of Washington, Seattle, WA** 07/2014
- Master of Science in Information Management
- Harbin Institute of Technology, China** 07/2012
- Bachelor of Business Administration
- GRANTS** Syracuse iSchool Research Grant (2019, \$6,500)
Syracuse iSchool Research Fellowship (2018, \$12,000)
National Science Foundation I-Corp Grant, (2018, \$50,000, Entrepreneurial Lead)
Syracuse iSchool Research Grant (2017, \$6,000)
Katzer Research Grant (2016, \$6,000)
- AWARDS & RECOGNITIONS** Best Paper Honorable Mention Award, ACM CSCW 2019
Best Paper Nomination, HICSS 2019
USENIX Student Travel Award (2019, \$2,360)
The Graduate Student Office Travel Award (2018, \$300)
The University of Houston Privacy and Security Workshop Award (2018, \$1,400)
National Science Foundation Travel Award (2017, \$175)
PoPETS Student Travel Award (2016, \$900)
Katzer Research Fund (2016, \$6,000)
Graduate School Fund for Excellence and Innovation (2014, \$400)
Mary Hotchkiss Endowed Fellowship (2014, \$1,000)
- PUBLICATIONS** **Journal Papers**
- J5. **Yao, Y.,** Basdeo, J. R., Mcdonough, O. R., Wang, Y. Privacy Perceptions and Designs of Bystanders in Smart Homes. In Proceedings of the ACM on

-  Human-Computer Interaction, Vol 3. CSCW 2019. **(Best Paper Honorable Mention Award)**
- J4. Huang, Y., Sang, Y., Wu, Q., Yao, Y. Higher Education Check-Ins: Exploring the User Experience of Hybrid Location Sensing. In Proceedings of the ACM on Human-Computer Interaction, Vol 3. CSCW 2019
- J3. Barbosa, N. M., Park, J. S., Yao, Y., Wang, Y. “What if?” Predicting Individual Users’ Smart Home Privacy Preferences and Their Changes. In *Proceedings on Privacy Enhancing Technologies (PoPETS)*. 2019.4. (Acceptance Rate: 19%)
- J2. Wang, Y., Xia, H., Yao, Y., Huang, Y. Flying Eyes and Hidden Controllers: A Qualitative Study of People's Privacy Perceptions of Civilian Drones in the US. *Proceedings on Privacy Enhancing Technologies (PoPETS)*. 2016 (3):1–19. (Acceptance Rate: 19%)
- J1. Scholl, H. J., Wang, K., Wang, Y., Woods, G., Xu, D., Yao, Y., & Krcmar, H. (2014). Top soccer teams in cyberspace: Online channels for services, communications, research, and sales. *Journal of Marketing Analytics*, 2(2), 98-119.

Conference Proceedings

(In computer science and HCI fields, conferences are highly selective and intended for archival papers only)

- C8. Yao, Y., Basdeo, J., Kaushik, S., and Wang, Y. Defending My Castle: A Co-Design Study of Privacy Mechanisms for Smart Homes. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*. ACM.
- C7. Yao, Y., Huang, Y., & Wang, Y. (2019, January). Unpacking People's Understandings of Bluetooth Beacon Systems-A Location-Based IoT Technology. In *Proceedings of the 52nd Hawaii International Conference on System Sciences*. **(Best Paper Nomination)**
- C6. Yao, Y., Xia, H., Huang, Y., & Wang, Y. (2017, May). Privacy mechanisms for drones: Perceptions of drone controllers and bystanders. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (pp. 6777-6788). ACM.
- C5. Yao, Y., Xia, H., Huang, Y., & Wang, Y. (2017, May). Free to fly in public spaces: Drone controllers' privacy perceptions and practices. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (pp. 6789-6793). ACM.
- C4. Yao, Y., Lo Re, D., & Wang, Y. (2017, February). Folk models of online behavioral advertising. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing* (pp. 1957-1969). ACM.

- C3. Huang, Y., Huo, S., **Yao, Y.**, Chao, N., Wang, Y., Grygiel, J., & Sawyer, S. (2016, June). Municipal police departments on Facebook: What are they posting and are people engaging? In *Proceedings of the 17th International Digital Government Research Conference on Digital Government Research* (pp. 366-374). ACM.
- C2. Guajardo, V., **Yao, Y.**, Bayo Urban, I. and Gomez, R. (2014). CasaCare.org: A sociotechnical platform for women immigrant workers in the home care industry. *11th Community Informatics Conference CIRN*, Prato, Italy, October 2014.
- C1. Jurisch, M., Krcmar, H., Scholl, H. J., Wang, K., Wang, Y., Woods, G., Xu, D., & **Yao, Y.** (2014, January). Digital and social media in pro sports: analysis of the 2013 UEFA top four. In *Proceedings of the 47th Hawaii International Conference on System Sciences* (pp. 3073-3082). IEEE.

Workshop Organized/Papers

- W6. **Yao, Y.**, Chouhan, C., Wong, R., Emami-Naeini, P., Merrill, N., Page, X., Wang, Y., Wisniewski, P. Ubiquitous Privacy: Research and Design for Mobile and IoT Platforms. To appear at *Companion of the 2019 ACM Conference on Computer Supported Cooperative Work and Social Computing*. ACM. Austin, TX, November 2019.
- W5. Ahmed, T., Barbosa, N. M., Calandrino, J., Coventry, L., Lerner, A., Marsh, A., Wang, Y., **Yao, Y.** 4th Workshop on Inclusive Privacy and Security (WIPS) – Privacy and Security for Everyone, Anytime, Anywhere. In *the 15th Symposium on Usable Privacy and Security (SOUPS 2019)*, Santa Clara, CA, August 2019.
- W4. Badillo-Urquiola, K., **Yao, Y.**, Ayalon, O., Knijnenurg, B., Page, X., Toch, E., Wang, Y., & Wisniewski, P. J. (2018, October). Privacy in Context: Critically Engaging with Theory to Guide Privacy Research and Design. In *Companion of the 2018 ACM Conference on Computer Supported Cooperative Work and Social Computing* (pp. 425-431). ACM. Jersey City, NJ, November 2018.
- W3. **Yao, Y.** Exploring a Speculative Design Approach for Inclusive Privacy and Security. *3rd Workshop on Inclusive Privacy and Security (WIPS)*, Baltimore, MD, August 2018.
- W2. Wang, Y., **Yao, Y.**, Whose Privacy? The Case of Drone Controllers and Bystanders. *Networked Privacy Workshop of the ACM Conference on Computer-Supported Cooperative Work and Social Computing (CSCW 2017)*, Portland, OR, February 2017.
- W1. **Yao, Y.**, Personalized Privacy Assistant to Protect People’s Privacy in Smart Home Environment. *Networked Privacy Workshop of the ACM Conference on Human Factors in Computing Systems (CHI 2018)*, Montreal, QC, Canada, April 2018.

Extended Abstract/Posters

- P8. Huang, Y., Sang, Y., Wu, Q., & **Yao, Y.** (2019, April). Studying User Experience of a Hybrid Location Sensing System. In *Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems*. ACM.
- P7. Huang, Y., Wu, Q., & **Yao, Y.** (2018, October). Bluetooth Low Energy (BLE) Beacons Alone Didn't Work!. In *Proceedings of the 2018 ACM International Joint Conference and 2018 International Symposium on Pervasive and Ubiquitous Computing and Wearable Computers* (pp. 62-65). ACM.
- P6. **Yao, Y.**, Bort, J., & Huang, Y. (2017, May). Understanding Danmaku's Potential in Online Video Learning. In *Proceedings of the 2017 CHI conference extended abstracts on human factors in computing systems* (pp. 3034-3040). ACM.
- P5. **Yao, Y.**, Huang, Y., Wang, Y. Unpacking People's Understandings of Bluetooth Beacon. *The 14th Symposium on Usable Privacy and Security*, Baltimore, MD, August 2018.
- P4. **Yao, Y.**, Xia, H., Kaushik, S., Wang, Y. Design and Evaluation of an Information-Based Web Tracking Blocking Mechanism. *The Federal Trade Commission PrivacyCon*, Washington D.C., February 2018.
- P3. **Yao, Y.**, Xia, H., Kaushik, S., Wang, Y. Design and Evaluation of an Information-Based Web Tracking Blocking Mechanism. *Great Lake Security Day*, Rochester, NY, September 2017.
- P2. Huang, Y., **Yao, Y.**, Bort, J., Wu, Q., Danmaku: Understanding its Usage in China and its Broader Potential. *CSCW in China and Beyond Workshop of the ACM Conference on Computer-Supported Cooperative Work and Social Computing (CSCW 2017)*, Portland, OR, February 2017.
- P1. Wang, Y., Xia, H., **Yao, Y.**, Huang, Y. Flying Eyes and Hidden Controllers: A Qualitative Study of People's Privacy Perceptions of Civilian Drones in the US. *The 12th Symposium on Usable Privacy and Security (SOUPS2016)*, Denver, CO, June 2017.

Magazine Articles

- M1. Walker, A., **Yao, Y.**, Geeng, C., Hoyle, R., and Wisniewski, P. (October 2019). Moving beyond 'one size fits all': research considerations for working with vulnerable populations. *Interactions* 26, 6 (October 2019), 34-39. DOI: <https://doi.org/10.1145/3358904>

Invited/Conference Talks

- T12 "Privacy Perceptions and Designs of Bystanders in Smart Homes." ACM CSCW, Austin, TX, November 2019

- T11. **INVITED:** “‘What if?’ Predicting Individual Users’ Smart Home Privacy Preferences and Their Changes.”
Federal Trade Commission (FTC) PrivacyCon, Washington D.C., June 2019
- T10. “Defending My Castle: A Co-Design Study of Privacy Mechanisms for Smart Homes.”
ACM CHI, Glasgow, UK, May 2019
- T9. “Explore Privacy in Smart Homes from a Multi-Stakeholder Perspective.”
iConference, Baltimore, MD, March 2019
- T8. **INVITED:** “Defending My Castle: A Co-Design Study of Privacy Mechanisms for Smart Homes.”
Syracuse University Graduate Seminar, Syracuse, NY, October 2018
- T7. “Exploring a Speculative Design Approach for Inclusive Privacy and Security.”
Inclusive Privacy and Security Workshop, SOUPS, Baltimore, MD, August 2018
- T6. “Privacy mechanisms for drones: Perceptions of drone controllers and bystanders.”
ACM CHI, Denver, CO, May 2017
- T5. “Free to fly in public spaces: Drone controllers' privacy perceptions and practices.”
ACM CHI, Denver, CO, May 2017
- T4. “Folk models of online behavioral advertising.”
ACM CSCW, Portland, OR, February 2017
- T3. “Personalized Privacy Assistant to Protect People’s Privacy in Smart Home Environment.”
Networked Privacy Workshop, ACM CHI, Montreal, QC, May 2018
- T2. “Whose Privacy? The Case of Drone Controllers and Bystanders.”
Networked Privacy Workshop, ACM CSCW, Portland, OR, February 2017
- T1. “Digital and social media in pro sports: analysis of the 2013 UEFA top four.”
HICSS, Big Island, HI, January 2014

Doctoral Colloquium

- D2. ACM Conference on Computer-Supported Cooperative Work and Social Computing (CSCW), Austin, TX, November 2019

D1. iConference, College Park, MD, March 2019

TEACHING
EXPERIENCE

Teaching Assistant

- IST 800 Privacy Policy, Law, and Technology, Doctoral-level, Spring 2018
Facilitated in-class discussion
- IST 649 Human Interaction w/Computers, Master-level, Fall 201
Gave guest lecture on prototyping; led in-class paper discussion
- IST 719 Information Visualization, Master-level, Spring 2016
Gave guest lecture on creating interactive plots using R and D3.js; led in-class discussion and exercise
- IST 687 Applied Data Science, Master-level, Fall 2015
Gave guest lecture on data cleaning using R; led in-class exercises; hold weekly office hours

Students Supervision

- Smirity Kaushik, *Master student in Information Management, Syracuse University (currently a Privacy Analyst at Earnest & Young)*
- Charlotte Emily Price, *Master student in Library Science, Syracuse University*
- Justin Reed Basdeo, *Undergraduate student in Industrial Design, Syracuse University*
- Oriana Rosata Mcdonough, *Undergraduate student in Informatics, Syracuse University*
- Jordan Hayes, *Independent researcher, Syracuse University*
- Casey Sawyer, *Undergraduate summer intern, University of the Pacific*
- Dylan Wiki Rogers, *Undergraduate summer intern, Bucknell University*

INDUSTRY
EXPERIENCE

Taggle. Inc, Software Engineer, Seattle, WA 06/2014-05/2015

- Supported the development of an iPhone application, Prose, to encourage users to sharing their readings and writings
- Successfully implemented a social network web application individually to serve more than 10,000 users

SERVICE

iSchool Service at Syracuse University

- Doctoral Program Committee Member (2015-2016)
- Faculty Committee Member (2016-2017)
- Faculty Search Committee Member (2018-2019)
- Syracuse University School of Information Studies Dean Search Committee (2018-2020)

Organizing Committee

- Co-chair, Publicity, SOUPS 2020
- Lead organizers, Networked Privacy Workshop, ACM CSCW 2019
- Co-organizers, Inclusive Privacy and Security Workshop, SOUPS 2019
- Co-organizers, Networked Privacy Workshop, ACM CSCW 2018

Program Committee

- ACM CHI 2021
- ACM GROUP 2020

- ACM CHI Late-Breaking Work 2017, 2019
- ACM DIS Provocations and Works-in-Progress 2019
- AAAI Fall Symposium on Privacy and Language Technologies (PLT) 2016, 2018
- Workshop on Usable Security and Privacy (USEC) 2019
- AAAI Spring Symposium on Privacy-Enhancing AI and Language Technology (PAL) 2019

Invited Reviewer

- ACM CHI 2016, 2017, 2018, 2019, 2020
- ACM CHI Late-Breaking Work 2017, 2018, 2019
- ACM CSCW 2016, 2017, 2018, 2019
- iConference 2016, 2017, 2018, 2019
- PoPETS 2016
- SOUPS 2016, 2017, 2018
- IEEE Security & Privacy 2018
- ICIS 2018
- USEC 2019
- IMWUT 2019
- ACM GROUP 2020
- JASIST

Student Volunteer

- ACM CHI 2017, 2018

REFERENCES

Dr. Yang Wang (Advisor)

Associate Professor
 School of Information Sciences
 School of Computer Science (affiliate)
 University of Illinois at Urbana-Champaign
 Email: yvw@illinois.edu

Dr. Jason Dedrick (Committee member)

Professor
 School of Information Studies
 Syracuse University
 Email: jdedrick@syr.edu

Dr. Bryan Semaan (Committee member)

Assistant Professor
 School of Information Studies
 Syracuse University
 Email: bsemaan@syr.edu

Dr. Pamela J. Wisniewski (Collaborator)

Assistant Professor
 College of Engineering and Computer Science
 University of Central Florida
 Email: pamwis@ucf.edu