An Introduction to Privacy for Technology Professionals

Executive Editor
Travis D. Breaux, CIPT

Contributing Authors
Lujo Bauer
Chris Clifton
Lorrie Faith Cranor, CIPT
Simson L. Garfinkel, CIPP/US
David Gordon
David James Marcos, CIPM, CIPT
Aaron Massey
Florian Schaub, CIPP/US, CIPT
Stuart S. Shapiro, CIPP/G, CIPP/US
Manya Sleeper
Blase Ur

An IAPP Publication

Copy editor and proofreader: Julia Homer

Indexer: Hyde Park Publishing Services

ISBN: 978-1-948771-90-0

Library of Congress Control Number: 2019953068

Contents

About the IAPP

The International Association of Privacy Professionals (IAPP) is the largest and most comprehensive global information privacy community and resource, helping practitioners develop and advance their careers and organizations manage and protect their data.

The IAPP is a not-for-profit association founded in 2000 with a mission to define, support and improve the privacy profession globally. We are committed to providing a forum for privacy professionals to share best practices, track trends, advance privacy management issues, standardize the designations for privacy professionals and provide education and guidance on opportunities in the field of information privacy.

The IAPP is responsible for developing and launching the only globally recognized credentialing programs in information privacy: the Certified Information Privacy Professional (CIPP®), the Certified Information Privacy Manager (CIPM®) and the Certified Information Privacy Technologist (CIPT®). The CIPP, CIPM and CIPT are the leading privacy certifications for thousands of professionals around the world who serve the data protection, information auditing, information security, legal compliance and risk management needs of their organizations.

In addition, the IAPP offers a full suite of educational and professional development services and holds annual conferences that are recognized internationally as the leading forums for the discussion and debate of issues related to privacy policy and practice.

Acknowledgments

In the years since *Introduction to IT Privacy: A Handbook for Technologists* was published, we've seen the rapid evolution of privacy technology and the growing importance of technologists in both ensuring compliance and creating a culture of privacy within their organizations. This book, *An Introduction to Privacy for Technology Professionals*, is the new edition of *Introduction to IT Privacy* and reflects these changes in the field. The IAPP is delighted to present this book in support of the Certified Information Privacy Technologist (CIPT) credential.

We would like to thank the many professionals who contributed their time and expertise to the development of this comprehensive resource. It's a collaborative effort, and without them, it would not be possible.

The training advisory board members provide ongoing support of our training products. We rely on their knowledge and guidance to develop the highest quality resources. Thank you to these members, past and present. Current members include:

Francesco Banterle, CIPP/E
Punit Bhatia, CIPP/E, CIPM
Machiel Bolhuis, CIPP/E, CIPM, FIP
Michaela Buck
Duncan Campbell, CIPP/US
Ionela Cuciureanu
Evan Davies, CIPP/E
Karen Duffy, CIPP/E
Marjory Gentry, CIPP/E, CIPP/US, CIPM
Promila Gonsalves, CIPP/C
Ryan Hammer, CIPP/E, CIPP/G, CIPP/US, CIPM, CIPT, FIP
Missi Hart-Kothari, CIPP/US
Richard Ingle
Laura Kiviharju, CIPM

Henri Kujala, CIPP/E, CIPM, FIP
Viviane Maldonado
Ana Monteiro, CIPP/E, CIPM, CIPT, FIP
Michelle Muthiani, CIPP/E, CIPP/US
James Park, CIPP/E, CIPT
Anna Pateraki
Cassandra Porter, CIPP/US, CIPM, FIP
Stephen Ramey
Brandon Schneider, CIPP/G, CIPT, FIP
Thea Sogenbits
Tiina Suomela, CIPP/E, CIPM, FIP
Liisa Thomas
Maaike van Kampen-Duchateau, CIPP/E, CIPT, FIP
Emily Wall, CIPP/US, CIPM
Ben Westwood, CIPP/E, CIPP/US, CIPM, FIP
Christin Williams, CIPP/E, CIPP/US
Brett Wise, CIPP/US, CIPT, FIP
Matthew Woldin, CIPP/US, CIPM, FIP
Laurel Yancey, CIPP/E, CIPP/US, CIPM
Philip Yang, CIPM

We are grateful that Travis D. Breaux, CIPT, served as executive editor for both *Introduction to IT Privacy* and *An Introduction to Privacy for Technology Professionals*. For both editions of the book, Travis assembled a team of top thought leaders in the field and, in addition to contributing to the text himself, led the books through all stages of content development. Other contributors to these books include Lujo Bauer; Chris Clifton; Lorrie Faith Cranor, CIPT; Simson Garfinkel, CIPP/US; David Gordon; Malcolm Harkins; David Hoffman, CIPP/US; David James Marcos, CIPM, CIPT; Aaron Massey; Florian Schaub, CIPP/US, CIPT; Stuart S. Shapiro, CIPP/G, CIPP/US; Manya Sleeper and Blase Ur. Thank you to Travis and all of the contributors for their dedication and continued support of our CIPT program and textbook.

Many thanks to Anne Christman; Biju Dasappan, CIPT; Jonathan Fox, CIPP/US, CIPM; Javier Salido, CIPP/US, CIPT, FIPP and Thea Sogenbits for reviewing the draft manuscript and providing thoughtful feedback that helped shape the final text. We are grateful for Julia Homer's meticulous eye in both copy editing and proofreading the text. Thank you to Hyde Park Publishing Services for creating the book index.

We are so appreciative for the many professionals who contributed to this textbook. We hope that you will find it to be both an indispensable resource in preparing for your certification as well as a practical resource for your professional career.

Marla Berry, CIPT
Training Director
International Association of Privacy Professionals

Preface

Since the first edition of this book, *Introduction to IT Privacy*, was published in 2014, we have observed significant new advances in information technology and public policy that affect privacy. In many ways, this eight-year period tells a story where technologies that historically lived in research laboratories have seen wider commercial adoption as new forms of automation. These technologies include autonomous and semiautonomous vehicles, voice-activated assistants, smart devices and biometrics. While the conveniences afforded by some forms of automation are still being understood, the deployment is already raising new privacy challenges. Google Glasses, the eyeglasses that could snap photos and later record video, were introduced in 2012 and quickly raised questions about reasonable expectations of privacy in public spaces. While this technology did not become mainstream with the public, voice-activated smart speakers, which consumers deploy in their homes and use to check the weather and play music, are quite popular, with multiple manufacturers competing for market share, including Amazon, Apple and Google.

The privacy risks introduced by these technologies are varied, including new ways to acquire personal data, new ways to build richer personal profiles, and new challenges in discerning the truth about a person from mere fabrication. Machine learning and robotics, for example, have led to commercial drones that allow individuals to more easily acquire overhead video and sensor data, smart televisions that detect which content is being viewed and then share this information with device manufacturers, and health sensors that track real-time fitness and blood pressure. These advances in many ways integrate with and extend the smartphone revolution that led to the novel collection of real-time location by mobile apps and advertisers. Additionally, increased deployment of sensors, often to enable new consumer products and safety features, allow for the creation of richer personal profiles. In 2015, telematics devices introduced by insurance companies were capable of recording driving distances, braking patterns and speed traveled for the purpose of informing insurance rates. In 2016, a popular social media site deployed an algorithm to classify its users by their political preferences,

which contributes to other factors in behavioral profiles, such as affinities for specific racial groups. Lastly, advances in generative machine learning now allow one to create fake images and fake video, called deep fakes because they rely on deep neural networks to generate the content. This includes an app that was designed to "auto-undress" photos of women as well as an app that allows a person to speak into a camera and record audio that is transformed into the video and audio of a public personality. These technologies raise new questions about the veracity of information and the use of data to misrepresent a person's character or beliefs, potentially poisoning the public record.

Finally, in recent years, we've seen major advances in regulatory practices aimed at addressing privacy in a world where data is increasingly shared globally. In Europe, the General Data Protection Regulation (GDPR) replaces the EU Directive 95/46/EC with a major change that requires companies to obtain an individual's consent before they can build user profiles. In addition, the National Institute for Standards and Technology (NIST) conducted a series of workshops with U.S. companies, government agencies and the public to develop a framework for privacy engineering. The framework aims to guide companies in how to select actions for reducing privacy risk within their enterprise. Going forward, regulators face great challenges in ensuring that regulation keeps pace with emerging technology as it shapes how we define and promote privacy. For example, how should companies treat machine learning models trained on data of EU citizens who choose to be forgotten?

In this new edition, we enshrined advances in technology, policy and practice in updates to all the existing chapters. This includes advances in cryptographic algorithms in Chapter 2 and database reconstruction attacks and new deployments of biometric authentication services in Chapter 4. In Chapter 6, we have new material on sensor-based surveillance, due in part to the emergence of the internet of things (IoT). Chapter 7 includes new material on deep fakes as well as on fairness and bias in machine learning, and Chapter 8 has been updated to focus more on enterprise privacy for cloud computing. Finally, we added two new chapters: Chapter 5, on how to make privacy usable given that users play an increasingly large role in managing their privacy preferences; and Chapter 9, on cybersecurity and how security frameworks support protecting privacy. While these are a few highlights for the new edition, I believe you'll find this updated volume offers a single source on the topic of IT privacy that is simply unparalleled in terms of breadth and depth.

Travis D. Breaux, CIPT
Executive Editor

Introduction

A universal aspect of being a technology professional is that change is not only expected but guaranteed. Change ensures the opportunities to learn never wane—that's what keeps many of us in the profession. However, in recent years, the pace of change is both increasing dramatically and expanding to affect the broad population. People have come to expect the same quality of engagement with day-to-day technology regardless of channel or industry.

On my own journey, finding myself in the privacy industry now feels reminiscent of the early days of information security. There wasn't a roadmap on the work, just passionate, tech-savvy people who cared about making sure company assets were safe from emerging security threats.

Privacy engineering is today's equivalent. While privacy laws and guidelines have been with us for decades, to date the approach has been about terms and conditions and contractually holding suppliers and service providers accountable for company behavior. With the ever-increasing presence, dependence, and personal implications of technology on individuals, the contractual is no longer enough.

The stakes are increasing. The prevalence of internet-of-things (IoT) devices like smart watches, smart homes and smart cities is increasing both our digital footprint and its perceived value and usage. Technology has made it possible to capture and track people's movements, interests and personal information and aggregate it for marketing insights or nefarious intentions.

While technology has become a convenience for many—making it easy to order food and clothes, track our exercise or health, and keep in touch with friends and family—it has also become a necessity. Our personal and work lives frequently rely on today's technologies to simply meet day-to-day expectations. People want the convenience but are recognizing that giving up our personal data is becoming annoying, creepy or the source of lost time, data or money because of the frequency of data breaches. Additionally, incidents of malware, ransomware and identity theft are fostering distrust of technology. More and more, people are becoming fatigued by the trade-off.

Recently, however, the laws have evolved to require organizations to manage this within the technology ecosystem. Specifically, the General Data Protection Regulation (GDPR) introduced the requirement of data protection by design and by default, radically changing the paradigm from the contractual to the automatic. For technology professionals, this means a new way to design, build, test and maintain our applications, infrastructure and processes. These guidelines place the individual's needs front and center, and require protection of their rights and interests to be built in our systems, processes and infrastructure.

Today's technology professional recognizes the urgency to quickly adapt to the pace of change. The emergence of user-centric design, the wide adoption of DevOps and Agile practices, and the increased commitment to diversity and inclusion in building software today are all evidence of this recognition. In line with the now table-stakes expectations of building secure solutions, technology professionals need to adapt to the growing demand for privacy.

This book offers technology professionals a better understanding of privacy issues and their implications for today's technology solutions, processes and infrastructure. Any person in the technology solutions ecosystem will benefit from the concepts, guidelines, tools and frameworks available in this book to ensure privacy by design in their solutions. Further, the book will give tech-savvy professionals a common language with which to speak with lawyers, compliance officers, business stakeholders and others involved in the definitions needed to design, build, test and maintain solutions with privacy in mind.

Cathleen R. Scerbo
Vice President and CIO
International Association of Privacy Professionals

CHAPTER 5

Usable and Useful Privacy Interfaces

Florian Schaub, CIPP/US, CIPT; Lorrie Faith Cranor, CIPT

The design of a system or technology, in particular its user experience (UX) design, affects and shapes how people interact with it. Privacy engineering and UX design frequently intersect. Privacy laws and regulations require that data subjects are informed about a system's data practices, asked for consent, provided with a mechanism to withdraw consent and given access to their own data—including options for data portability and data deletion. To satisfy these requirements and address users' privacy needs, most services offer some form of privacy notices, privacy controls or privacy settings through a website or through the user interface of an app or other piece of software.

However, too often privacy notices are not readable, people do not understand what they consent to, and people are not aware of certain data practices or the privacy settings or controls available to them. The challenge is that an emphasis on meeting legal and regulatory obligations is not sufficient to create privacy interfaces that are *usable* and *useful* for users. Usable means that people can find, understand and successfully use provided privacy information and controls. Useful means that privacy information and controls align with users' needs with respect to making privacy-related decisions and managing their privacy. This chapter provides insights into the reasons why it can be difficult to design privacy interfaces that are usable and useful. It further provides guidance and best practices for user-centric privacy design that meets both legal obligations and users' needs. Designing effective privacy user experiences not only makes it easier for users to manage and control their privacy, but also benefits organizations by minimizing surprise for their users and facilitating user trust. Any privacy notice and control is not just a compliance tool but rather an opportunity to engage with users about privacy, to explain the rationale behind practices that may seem invasive without proper context, to make users aware of potential privacy risks, and to communicate the measures and effort taken to mitigate those risks and protect users' privacy.

Privacy laws, privacy technology, and privacy management are typically centered on information—how information is collected, processed, stored and transferred; how it can and must be protected; and how to ensure compliance and accountability. To be

effective, designing privacy user experiences requires a shift in focus: While information and compliance are of course still relevant, user-centric privacy design focuses on people, their privacy needs and their interaction with a system's privacy interfaces.

Why is it important to pay attention to the usability of privacy interfaces? How do people make privacy decisions? What drives their privacy concerns and behavior? We answer these questions in this chapter and then provide an introduction to UX design. We discuss common usability issues in privacy interfaces and describe a set of privacy design principles and a user-centric process for designing usable and effective privacy interfaces, concluding with an overview of best practices.

5.1 Why User-Centered Privacy Design?

If you have ever tried to adjust your privacy settings in a mobile app or a social media site, or tried to figure out how to exercise privacy choices on a website, chances are you encountered a privacy interface that was not all that easy to use. Developers often add privacy interfaces to comply with regulations or company policy, but seldom do they engage user experience experts to design a usable privacy interface or evaluate its usability. However, some privacy laws now include usability requirements, which means that simply providing a choice somewhere in the user interface may not be sufficient to meet legal obligations. Furthermore, people who encounter privacy notices and controls that are difficult to find, use or understand may be surprised when they read a news report and find out about how an organization actually uses their data. Thus, even well-intended organizations may end up disappointing or surprising their customers, which in turn may decrease their trust in the organization. By taking a user-centered approach to privacy design, organizations can meet legal obligations while also minimizing user surprise and facilitating trust.

5.1.1 Meeting Legal Obligations

Many privacy laws include requirements for transparency and choice that include usability requirements. For example, the EU General Data Protection Regulation (GDPR) requires "clear and plain language" and transparency that is "concise, transparent, intelligible and easily accessible."[1] A number of U.S. laws require privacy policies to be "clear and conspicuous." Legal requirements related to privacy notice usability focus on text that is readable to users, easy for users to find, and contains information users need to know. Some laws also require that users be presented with easily accessible choices about the use of their data, ways to access the data an organization has collected about them, and mechanisms for requesting that their data be deleted. While following the legal obligations

and requirements is important, it is often not enough to ensure that privacy notices and controls are actually useful and usable for users.

5.1.2 Minimizing Surprise

When users are surprised at how their data is used by an organization, they are likely to become angry and lose trust in that organization.[2]

5.1.2.1 Aligning User Expectations and System Behavior

Users often have preconceived notions about how a system will use and share their personal information. As they use the system, the UX may further shape their expectations. Users may be unaware of data collection that occurs automatically or of inferences that are made based on their behavior. They may assume that companies will use data only to fulfill user requests and not for unrelated purposes. Any unexpected data use or sharing may come as an unpleasant surprise to users when they learn about it through news reports, social media, or because they are affected by how the data is used. For example, in 2018, it was revealed that Cambridge Analytica was able to obtain information about a huge number of Facebook users because Facebook apps could access information not only about those Facebook users who had installed an application but also about the friends of those users, without their knowledge. This data sharing with third parties by Facebook was met with shock and surprise by consumers, and led to the largest regulatory penalty for a privacy violation to date.[3]

Companies can prevent these surprises by helping users understand how their data will be used and shared, and by drawing the user's attention to data practices most likely to be surprising. For example, a company might highlight potentially surprising practices in a privacy policy and include just-in-time notices where appropriate throughout the system's UX.

Likewise, companies should make an effort to design privacy-related controls so that users can fully understand what they do. When users check a box to opt-out, they should know exactly what they are opting out of; for example, are they opting out of seeing advertising, ads that are targeted to them, or the actual tracking of their behavior for ad targeting? Most advertising-related opt-outs fail to clarify the actual effect of the opt-out.[4] Likewise, access control settings for social media and calendars are often difficult to understand, and users may inadvertently grant access to more people than intended. Therefore, some calendar programs assist users who change the visibility of a calendar entry by displaying an explanation of how this will impact who can see that entry. The user interface should also make users aware of any side effects that may occur as a result of the settings they choose. For example, allowing an app access to location information might not only enable mapping features, but also enable location-based advertising.

5.1.2.2 Reducing Opportunities for User Regret

Sometimes users take actions that may have an adverse impact on their privacy that they may later come to regret. This may occur when users misunderstand settings, select the wrong settings by mistake, or don't understand or misremember important details. For example, social media users may forget who is in their friends or followers group and thus may be surprised when their parent, boss or coworker sees a post not intended for them, causing regret.[5] Some systems use "nudges" or time delays to reduce opportunities for user regret. Researchers have developed and tested nudges for Facebook that remind users of who is in the audience for their posts and provide a short count-down timer before posting.[6] Facebook itself has implemented a nudge to address the problem of users changing their audience setting to public for one post and forgetting to change it back to a more privacy-protective setting next time they post.[7] Gmail offers users the opportunity to set a delay of 5-30 seconds before sent messages are actually sent, giving users a short opportunity to recall them.[8] Features such as these that reduce opportunities for user regret can help users help themselves and protect their privacy.

5.1.3 Facilitating Trust

Users are more likely to trust companies that they believe will treat them fairly, communicate clearly about their data practices and privacy protections and offer them meaningful choices and controls.

5.1.3.1 Communicating Privacy Protections

Users are not likely to trust privacy information or notices that appear to users as just a bunch of legalese meant to protect the organization. Furthermore, when users—or the media—find that a company or product engages in data practices that are surprising and privacy invasive, it can destroy trust and reputation. Often practices are unexpected because they were not communicated clearly or conspicuously. For example, in 2019, multiple companies producing smart speakers and smart assistants were criticized because they had employees listen to and manually annotate people's voice commands. While such manual annotation may be necessary to improve speech recognition, a statement in the privacy policy that "your data may be used to improve our services" does not properly convey that people might be listening to what you tell your smart speaker in the privacy of your home. To help build trust, organizations should communicate about privacy in a clear and straightforward way that will be meaningful to users. Privacy communication does not just serve compliance purposes—it provides opportunities for organizations to communicate their efforts for protecting their users' privacy, which otherwise might be invisible to the user.

5.1.3.2 Providing Meaningful Choices and Controls

Besides communicating clearly about privacy, organizations should offer meaningful choices and controls. A user who fully understands an organization's practices but is not offered any meaningful way to control them may not have a positive view of the organization. On the other hand, privacy controls that enable users to manage their privacy settings in ways that align with their privacy preferences and needs provide a sense of agency and control. This can help foster trust in the organization's ability and intent to properly protect its users' privacy.

5.2 Privacy Decision-Making, Behavior and Concerns

Transparency and *user rights* are core concepts of privacy legislation and guidelines globally, ranging from the Organisation of Economic Co-operation and Development (OECD) privacy guidelines, to the U.S. Federal Trade Commission's (FTC's) fair information practice principles, to Europe's GDPR, and privacy legislation in many other countries.[9] While specific requirements may vary, companies that collect or process personally identifiable information (PII) typically have to be transparent about their data practices and to inform data subjects about their rights and options for controlling or preventing certain data practices. This is often known as the *notice and choice* model.

Given such transparency and respective choices, consumers are supposed to be able to make informed privacy and consent decisions. However, in practice, privacy decision-making is not quite as simple or straightforward. Understanding how people make privacy decisions and what shapes their privacy concerns is essential for being able to design usable and useful privacy interfaces that effectively make people aware of data practices and support informed privacy decision-making.

5.2.1 Privacy Preferences versus Behavior

The notice and choice model is based on rational choice theory: A rationally acting person ingests information about companies' data practices and uses this information to engage in a risk-benefit analysis and make rational decisions about which services to adopt or which companies to trust with their personal information. Such decisions are assumed to be consistent with the person's privacy preferences in order to optimize their desired level of privacy. They are also assumed to be stable and consistent across decisions and time. This kind of rational decision-making about privacy is also referred to as the *privacy calculus*.[10] However, in practice people's privacy decisions and behavior are rarely rational or predictable but rather highly context dependent and malleable.

5.2.1.1 Privacy Paradox

Even though research studies and surveys frequently find that people are concerned about their privacy and want to do more to protect their privacy,[11] people share copious personal details on social media,[12] express regrets about sharing too much information online,[13] are frequently surprised by the data practices of services they use,[14] and are often unaware of privacy controls and protections available to them.[15]

This contradiction of people being concerned and complaining about data practices while freely providing and sharing personal information is known as the *privacy paradox*: People express certain privacy preferences or intentions but act contrary to them.[16] For example, people may be sensitive about sharing their health information, yet may use fitness trackers, smart watches or consumer genetic testing kits because they provide them insights on their health, even though by doing so, they also end up giving detailed health information to the companies behind those gadgets or services.

5.2.1.2 Self-Censorship and Chilling Effects

While people seem to be sharing information freely despite privacy concerns, research also suggests that people engage in self-censorship because of privacy concerns. For example, some people may opt not to share personal opinions, political leanings or certain updates about their lives online in order to avoid arguments or to manage their self-presentation.[17]

Increased awareness about invasive data practices may also cause *chilling effects*. For instance, Edward Snowden's 2013 revelations about the U.S. government's online surveillance practices have been linked to a significant decline in searches for terrorism-related topics (e.g., Al Qaeda, jihad, chemical weapons) both on Wikipedia and Google Search.[18]

5.2.2 Factors Affecting Privacy Decision-Making and Behavior

How can the privacy paradox, over-sharing, self-censorship, and chilling effects all exist at the same time? Table 5-1 distinguishes between a set of related yet distinct concepts regarding privacy decision-making, providing a lens through which to view and understand seemingly contradictory privacy behavior.

Table 5-1: Common Concepts in Privacy Decision-Making and Behavior

Concept	Nature	Description
Privacy attitude	Normative	The data subject's general predisposition regarding privacy
Privacy preference	Aspirational	What the data subject prefers to happen
Privacy concern	Hypothetical	What the data subject fears will happen
Privacy expectation	Anticipatory	What the data subject thinks will happen
Privacy decision	Intentional	What the data subject decides or intends to do
Privacy behavior	Factual	What the data subject actually does
Data practice	Factual	The data processing that actually occurs
Privacy harm	Factual	A negative impact of data processing on the data subject's privacy
Privacy regret	Retrospective	The data subject realizing behavior and expectations were harmfully misaligned

A person's *privacy preferences* are aspirational—they describe what the person would prefer to happen in a certain situation, which may be informed by one's general *attitudes* towards privacy, sociocultural norms, and prior experiences. *Privacy concerns* are privacy risks or potential *privacy harm* a person is aware of or worried about; they may range from minor concerns to worst-case scenarios. *Privacy expectations*—what data processing or privacy infringements a person anticipates will occur in a given situation—are shaped by the person's privacy preferences, concerns and awareness of data practices. Such expectations may be consistent with the actual *data practices*, or may be inaccurate—either assuming a lower privacy risk than exists or expecting privacy risks that are unwarranted.

Privacy preferences, expectations and concerns all factor into a person's privacy decision-making process, both consciously and subconsciously. People rarely make fully rational privacy decisions in which they seek out all available information and optimize their decision to align with their preferences—indeed, in many situations, this would be quite time consuming and may not even be possible. *Privacy decisions* are subject to both external influences (e.g., incomplete information, context) and internal influences (e.g., bounded rationality, experience). As a result, a person's privacy decision in a given situation may not be a consistent reflection of their privacy preference or might even contradict them. The implementation of a privacy decision (*privacy behavior*) may further be affected by external factors, such as choice architecture, in the given situation. We discuss these factors in more detail next.

People experience *privacy regret* when they realize that their privacy behavior or the actual data practices of a system or an organization were misaligned with their privacy expectations in a way that negatively affected them. Data processing can negatively impact the data subject's privacy (*privacy harm*), especially when data practices and the data subject's expectations are misaligned. Sometimes people observe tangible harm, but often they are not aware of the harm. Privacy harms may also not manifest immediately but rather much later. For example, a person's social media posts made years ago as a college student on a long-forgotten social media service may be ingested by a data broker, which sells the information to a job-candidate vetting service, which in turn downgrades the person's job application, resulting in the person not being invited for a job interview. The person might never learn that not being invited to the job interview was a privacy harm.

In light of these many factors influencing privacy decisions and behavior, the often-observed misalignment between people's privacy preferences and actual behavior (the privacy paradox) is less surprising. In particular, *incomplete information, context and bounded rationality* affect privacy decisions and behavior in ways that can result in substantial deviations in a person's privacy-related behavior from their privacy preferences.[19]

5.2.2.1 Incomplete Information and Uncertainty

Informed consent and informed decision-making both assume that a person has fully considered all available information to make an informed decision. However, more often than not, people have to make decisions based on incomplete information. In the case of privacy, this might mean that a person is not aware of all data collection, data processing or data transfer practices associated with a transaction. This is particularly the case when information is seemingly collected for one purpose but also used for secondary purposes. For example, a navigation app needs the user's location information to display the correct part of the map and provide navigation support, but the app provider might also use location information to identify roads with traffic congestion or to infer where the user lives and works or what businesses they frequent.

Even if a person is aware of a data practice, they may not be aware of the associated privacy risks or implications. For instance, in the navigation app example, it may not be apparent to all users that inference of their home and work locations not only makes it easier to quickly navigate to those places but may also reveal information about their socioeconomic status or family situation, and that such information could negatively affect other parts of their lives, ranging from what advertisements are shown to them to what interest rate they are being offered on a home or car loan.

People may also have misconceptions about how their privacy is protected in a given context. They might assume that companies securely transfer and store data, but frequent data breaches show that is not necessarily the case. They might assume that they are protected by privacy law, limiting what data can be collected or how it can be used. In practice, legal privacy protections can vary substantially between countries and even between industry sectors in countries with sectoral privacy laws (e.g., the United States). This can lead to misconceptions such as the common assumption that when a website has a link to a privacy policy, it does not share information about its users with third parties, even though the existence of a privacy policy alone does not guarantee privacy-protective data processing.[20]

Misconceptions may also exist in the other direction: Because of incomplete information or uncertainty about data practices, consumers may assume that a company or product is insecure or engaging in privacy-invasive practices, such as selling their information to third parties, even if it isn't. Thus, helping users develop an accurate understanding of data practices, privacy protections and privacy controls not only provides transparency about data practices but also creates awareness about the measures taken to protect privacy and limit data use, as well as privacy controls and choices available to the user. In fact, research has shown that users may provide and share more information when they feel they have appropriate privacy controls to determine and adjust who has access to their information. This phenomenon has been dubbed the *control paradox*, as perceived control over privacy may lead to increased sharing, which in turn may increase privacy risks.[21]

5.2.2.2 Bounded Rationality

Generally, humans are limited in their ability and time to acquire, memorize and process all information relevant to making a fully informed and rational decision. Behavioral economists call this deviation from the ideal of a fully rational actor *bounded rationality*.[22] To compensate for the inability and impracticality of considering all potential outcomes and risks, humans rely on heuristics in their decision-making to reach a satisfactory solution rather than an optimal one. However, decision heuristics can lead to inaccurate assessments of complex situations.[23] Rational decision-making is further affected by cognitive and behavioral biases—systematic errors in judgment and behaviors.

Bounded rationality also persists in privacy decision-making. Decision heuristics and biases have been shown to substantially affect people's privacy decisions and behavior.[24] Prominent decision heuristics and biases that can affect privacy decisions and behavior include the following:[25]

- *Availability heuristic*—Due to uncertainty about privacy risks, people may look for other available cues to judge the probability of risks and guide their behavior. For example, rather than reading an online store's privacy policy, people rely on readily available cues, such as the store's visual design, the presence of a privacy policy, the vendor's reputation, or even just the company name to make judgments about the privacy risks associated with providing personal contact and payment information to the online store.
- *Representativeness heuristic*—People may perceive privacy intrusions as low-probability events because they rarely encounter privacy intrusions online. However, privacy intrusions, such as behavioral tracking and targeting, may occur frequently or continuously but may just not be visible to the individual.
- *Anchoring*—Available information creates a reference point for future decisions. For example, information or judgments about others' disclosure behavior informs how one reasons about one's own disclosure behavior.[26] Anchoring also manifests in ordering effects: Survey participants disclose more information when a survey starts with intrusive questions and gradually reduces in sensitivity compared with a survey that increases in sensitivity.[27]
- *Loss aversion*—Individuals dislike losses more than they like gains. In the context of privacy, loss aversion helps explain why people report being concerned about companies collecting information about them but are reluctant to pay for better privacy protection (if it is offered).[28]
- *Hyperbolic discounting*—Providing or sharing information often affords immediate gratification of some kind, e.g., the ability to use a needed app/service or social interactions (e.g., messaging or "likes"), whereas privacy implications and risks are often not immediate, hypothetical and shrouded by uncertainty. Furthermore, privacy considerations are rarely the primary motivating factor for wanting to use a service or to enter into a transaction. Thus, even if people claim to care about privacy, they may discount privacy risks in the moment in favor of immediate gratification.[29]
- *Optimism bias*—People systematically underestimate the likelihood of being affected by a negative event. For example, even though people recognize that identity theft is a risk after a data breach, they underestimate the likelihood of personally having their identity stolen and, as a result, may not take sufficient protective action.[30]

- *Status quo bias*—People have a general affinity for default choices. People often keep default settings even if they are privacy invasive, because they are not aware of the setting and/or its privacy implications, because of associated transaction costs (e.g., finding the setting, figuring out how it works, actual costs), or because it is assumed that the default settings are set to protect them.[31]

In practice, multiple such heuristics and biases can affect a person's decisions and behavior at the same time, further exacerbating inconsistencies in behavior and between privacy preferences and actions.

5.2.2.3 Context Dependence

Privacy preferences and concerns can vary from situation to situation. What information someone considers appropriate to share varies based on contextual factors, such as the nature or source of the information, the activity or transaction as part of which information may be shared, as well as the people involved. For instance, you might reveal more about your personal life to a colleague at an after-work event than at work; you adjust how and what you talk about with your friends depending on where you are and who else is around. At the same time, you may also discuss sensitive topics in public spaces, e.g., at a restaurant or on a train, even though others might overhear the conversation. Similarly, people might reveal personal details to a stranger on a plane, their taxi or rideshare driver, or their hairdresser as part of making small talk. Privacy preferences are highly contextual or, as Acquisti et al. state, "the same person can in some situations be oblivious to, but in other situations be acutely concerned about, issues of privacy."[32]

Privacy preferences and expectations are shaped and affected by contextual factors. Gary Marx argues that privacy violations can always be traced back to the breaching of one or more of the following boundaries:[33]

- *Natural borders*, such as walls, closed doors, clothing, sealed envelopes and encryption protect information by limiting observation by others
- *Social borders* are assumptions or expectations of social norms about confidentiality and respect of one's privacy, such as confidential relationships with doctors, lawyers or priests; the integrity and confidentiality of personal correspondence; trust in colleagues, friends and family members to not rifle through one's personal effects; or the assumption that information is not retained longer than required or used for other purposes
- *Spatial or temporal borders*, such as physical distance or the passing of time, separate information from different periods or aspects of a person's life

- *Ephemeral or transitory borders* are based on assumptions that certain interactions or communication only exist in the moment and are not recorded permanently

Each of these borders is increasingly weakened by advances in technology. The proliferation of smart home devices, most of which continuously exchange information with the manufacturer's cloud back end, weakens the reliability of natural boundaries to protect privacy. Social, spatial and temporal borders break down in the face of networked publics and social media. Context collapse—the overlapping and intermixing of previously delineated facets of one's life (e.g., family life, work, different circles of friends, hobbies)—is now a common phenomenon in online interactions, especially interactions over social media.[34] Assumptions of information being ephemeral or transitory run contrary to the vast digital traces people's interactions with technology create, as well as consumer tracking and profiling efforts by companies.

Helen Nissenbaum's framework of *contextual integrity* ties privacy expectations to context-dependent norms of information flow.[35] Information collection, processing and transmission practices that are in accordance with those norms are likely to be perceived as acceptable, whereas practices that do not follow those norms are perceived as privacy violations because they violate contextual integrity.

Contextual integrity can be used to analyze data practices for their appropriateness by identifying the contextual factors of a situation, the informational norms associated with them and whether an actual practice adheres to those norms or violates them. Context-relative privacy norms may be codified, e.g., in laws, rules or procedures, but frequently such norms are implicit as part of conventions, moral values and people's expectations of privacy.

As people go about their daily lives, they transition through many, sometimes overlapping contexts, and in doing so engage in a continuous and dynamic *boundary regulation process* to manage their privacy.[36] People's context-specific privacy preferences are affected by *external changes*, such as changes in context, actors, information or transmission principles, as well as *internal changes*, such as changing privacy attitudes, preferences, prior experiences or new knowledge. People make adjustments to their behavior in an attempt to negotiate an achieved level of privacy that aligns with their privacy preferences. As already discussed, this process is subject to uncertainty and bounded rationality, as well as external constraints, such as the inability to prevent a certain data practice.

An important aspect of this dynamic boundary regulation process is the intention to balance privacy preferences, expectations and achieved privacy. If the achieved privacy is less than what is preferred, a privacy invasion occurs. If the achieved privacy

is more than desired, it results in social isolation. Thus, privacy regulation is not just about keeping information private, but rather about finding—and helping users find—the desired balance between privacy and information sharing for a given context.[37] This pertains both to managing one's privacy toward a data processor and managing information in contexts that involve interpersonal interactions. When multiple people or entities have access to certain information, privacy management becomes a collective task that relies on the negotiation of mutually acceptable privacy rules about further processing or sharing of that information. *Private information turbulence*—misalignment of or disagreement about what is appropriate or invasive—occurs when such rules are violated or have not been effectively negotiated in the first place.[38]

5.2.3 Manipulation of Privacy Behavior

Privacy-related decisions and behavior, such as what information someone is willing to disclose, can be manipulated through system design. A system's *choice architecture*—what and how choices are presented to a user—has a substantial impact on users' information disclosure and sharing behavior. Dark patterns are interface or system designs that purposefully exploit cognitive and behavioral biases in order to get people to behave a certain way regardless of whether that behavior aligns with their preferences. Some common privacy dark patterns include the following:[39]

- *Default settings*—Default settings frequently exploit status quo bias. Most people do not review or change default settings, which in the case of privacy means that opt-out choices, such as opting-out of targeted advertising or marketing use of contact information, are rarely exercised even though people may not like those practices. Similarly, preselecting a certain option nudges users towards accepting that choice.
- *Cumbersome privacy choices*—Making it more difficult, arduous and lengthy to select a privacy-friendly choice compared with a privacy-invasive one deters users from privacy-friendly options.
- *Framing*—How a choice is described and presented can affect behavior. An emphasis on benefits, a de-emphasis on risks or the presentation of trust cues may lead to people making riskier privacy decisions than they would with a neutral presentation of choices.[40]
- *Rewards and punishment*—Users are enticed to select a service's preferred choice with rewards or are deterred from privacy-friendlier choices with punishments. For instance, granting a mobile app access to location data

may enable certain services that are not available otherwise. Rewards and punishments are particularly problematic when they are not directly related to the choice, i.e., when choosing a privacy-friendly option poses stronger constraints than is necessary. Rewards, in particular, such as a meter that goes up when you add more information to your online profile, leverage the human need for immediate gratification.[41]

- *Forced action*—Users must accept a data practice or privacy choice in order to continue to a desired service (hyperbolic discounting), regardless of whether they actually agree with the practice. They are forced to act against their privacy preference.
- *Norm shaping*—Other people's observed information-sharing behavior, say, on a social media service, shapes the perceived norms of information sharing on a platform and individuals' own disclosure behavior. For example, a controlled experiment showed that people who see more revealing posts in the feed of a photo-sharing service tend to consider such photos more appropriate and are more likely to share more revealing information themselves than those who see less revealing photos.[42] Thus, the algorithm for determining news feed content, which might purposefully highlight or suppress certain content to activate such anchoring effects, has a lot of power in steering users' behavior, regardless of whether the displayed posts are actually representative of user behavior.
- *Distractions and delays*—Even small distractions or delays can create a distance between awareness of privacy risks and behavior that can cancel out the effects of a privacy notice.[43]

Such manipulations of privacy behavior are unethical, as they constrain people in their self-determination and agency over their privacy. Moreover, manipulations of privacy behavior further exacerbate the misalignment between people's privacy preferences, expectations and their actual behavior and the data practices to which they are subject.

Furthermore, tricking people into agreeing to data practices that do not meet their preferences or behaving in ways that go against their privacy concerns and preferences can result in increased regrets by users. Such regrets may intensify the perception that their privacy has been violated, and ultimately can negatively affect people's adoption of technology or result in outcries in social media as well as bad press.

Even if your intentions are good, the user interface of your privacy notices and controls can unintentionally mislead your users, cause surprise about unexpected data practices, get people to agree to data practices or set privacy settings in ways that

do not align with their preferences, or accidentally promote oversharing resulting in regret. Therefore, UX design plays an important role in making privacy information and controls understandable and usable.

5.3 Usability and User Experience

The key to creating usable privacy notices and controls is UX design, sometimes also referred to as customer experience design. Well-designed privacy interfaces support users in forming an accurate understanding of a system's data practices and provide them with the right tools and support for managing their privacy. Depending on the size of your organization, you might have dedicated UX designers, UX researchers or UX engineers who are responsible for the user-facing side of a product, or your organization might outsource some UX research or design work. However, UX professionals are not necessarily trained or experienced in designing privacy user experiences. Therefore, it is important for privacy professionals to work together with UX professionals to ensure that privacy requirements and users' privacy needs are taken into account. It is useful for privacy professionals to be aware of the UX design process and general methods in order to make informed decisions about when and how to work with UX professionals on the design of privacy interfaces. Even just posting a privacy policy on your website is a UX design decision, and too often one that is made without properly considering the UX and how it impacts the user.

5.3.1 What is UX?

UX design has one major goal: designing systems and processes to meet users' needs. This includes fostering understanding of how the system works, building systems that are useful for people and ensuring that people are able to use the system as intended. Ideally, people should also enjoy using the system.

5.3.1.1 Usability

At the core of UX is usability, i.e., how easy a system or interface is to use. ISO 9241-11 defines usability as the "extent to which a system, product or service can be used by specified users to achieve specified goals with effectiveness, efficiency and satisfaction in a specified context of use."[44] Jakob Nielsen, one of the pioneers of human-computer interaction, distinguishes the following components that determine a system's usability:[45]

- *Learnability*—How easy is it for users to accomplish basic tasks the first time they encounter the system?

- *Efficiency*—Once users have learned the system, how quickly can they perform tasks?
- *Memorability*—When users return to the system after a period of not using it, how easily can they reestablish proficiency?
- *Errors*—How many errors do users make, how severe are these errors and how easily can they recover from the errors?
- *Satisfaction*—How pleasant is it to use the system?

Consider, for example, the task of unsubscribing from a promotional email. The effectiveness of an email opt-out mechanism would be determined by whether a person can find the opt-out in the first place and whether they can understand what steps to undertake to complete the opt-out (learnability), execute those steps in the right order and achieve the intended result (success/errors) and later check or change their setting (memorability). Efficiency would be determined by the number of steps and the amount of time it takes to complete the opt-out. Satisfaction would be determined by the perceived ease or difficulty of completing the task and whether it was a pleasant and supportive process or one that was arduous and confusing.

5.3.1.2 Utility

A concept closely related to usability is utility.[46] Utility is about functionality. Does the system support users in satisfying their needs and accomplishing their goals? An interface can be very usable, but it is useless if it does not align with users' actual needs and expectations.

For example, the unsubscribe mechanism might be easy, fast and pleasant to use (great usability), but only gives users the option to unsubscribe from all of an organization's communication or none of them, even though some users might want to unsubscribe from marketing emails but continue receiving important notifications about their account activity. As a result, people may not use the very usable opt-out mechanism because it is not useful for them.

A system with high utility meets the exact needs of users. A useful system has both good utility and good usability.

5.3.1.3 User Experience

As the discussion of utility and usability suggests, usability is important but, on its own, often not sufficient to characterize what constitutes a good or bad experience for users. UX design takes a more holistic perspective that places users and their needs at the center and "encompasses all aspects of the end-user's interaction with the company,

its services, and its products."[47] This might include the actual product; terms of service and privacy policies tied to the product; the product's purchase, unboxing and sign-up experience as well as customer support, documentation, manuals, privacy settings and so on.

A system's UX encompasses the extent to which a system meets users' needs (utility), usability, aesthetics and simplicity, and the joy, emotional reactions and fulfillment a system provides. UX design therefore integrates user interface design, engineering, visual design, product design, marketing and branding as well as business models and related considerations.

5.3.2 User-Centered Design Process

Good UX design does not just happen. Similar to engineering practice, UX design follows a principled and systematic process. At the center of the design process are users—typically a set of anticipated user populations and stakeholders. While different methodologies follow slightly different steps, generally the user-centered design process consists of three phases—research, design, evaluation—with projects typically going through multiple iterations of the user-centered design process to iteratively refine designs and better align them with user needs.

5.3.2.1 UX Research: Understanding Context of Use and User Needs

UX design starts with research and analysis. The goal of this phase is to understand the context in which a certain system and its UX will operate and function. An important aspect of this is identifying a system's user populations and then analyzing their specific characteristics and needs as they relate to both the system and its context of use. This often includes learning about people's mental models of systems or processes in order to identify and understand potential misconceptions.

Common user research methods and activities include semistructured interviews, diary studies, contextual inquiry, survey research and usability tests (with the current system or related competitors). User research is typically complemented by desk research, involving competitive analysis, heuristic evaluation and review of relevant academic research and literature. Personas, scenarios, user journeys and affinity diagrams are tools used for synthesizing findings into higher-level insights that can be used and referenced throughout the design process.[48]

Based on the research findings, the requirements for the UX design are defined. This includes specifying which user needs and goals have to be addressed to be successful, including respective metrics. Design requirements have to consolidate potentially contradicting user needs, business goals and regulatory requirements.

Conducting UX research is quite comparable to conducting a privacy impact assessment (PIA), despite the different focus (user versus privacy). Like a PIA, UX research requires planning to ensure specific activities are generating relevant information and are comprehensive.[49] Both a PIA and UX research require identification and consultation of relevant stakeholders. A PIA includes mapping out a system's information flows; UX research involves mapping out users' journeys and interactions with a system. A PIA assesses and identifies privacy risks; UX research identifies user needs and usability issues. A PIA provides recommendations for mitigating privacy risks; UX research produces design requirements to meet user needs and mitigate usability issues. The parallels between PIAs and user-centered design go even further. Just as a PIA is not just "done," but is better understood as a continuous or recurring process in which PIA documentation is updated and adapted according to changes to the system, UX research is also iterative, accompanying development and deployment of systems. These similarities provide opportunities for integrating UX research aspects into PIAs and vice versa.

5.3.2.2 UX Design: Meeting User Needs

The requirements identified in the research phase inform the design phase. UX design aims to find and create solutions that meet user needs, as well as other system requirements. Good UX design takes users' cognitive and physical characteristics into account. UX design is highly iterative and user centric. Solutions are developed in an iterative process that aims to put potential solution designs into the hands of users early on and throughout the refinement process in order to ensure that designs are properly supporting user needs. Designs typically start with very basic prototypes, often sketches and paper prototypes, which look far from the final product but make it possible to simulate and test interaction flows before investing time, effort and resources in further development. This facilitates creative exploration of the solution space—including unconventional, novel ideas—in relatively little time and at low costs. The fidelity of prototypes increases throughout the design process as ideas from different designs get consolidated or trimmed away based on feedback gained through user testing, until the final design is implemented fully.

Good practice in UX design is to leverage established best practices—design patterns—when possible and appropriate. Design patterns encapsulate reusable solutions for common UX challenges. Leveraging established design patterns or language can substantially ease learnability because people are already familiar with a certain interaction. However, design patterns are patterns rather than off-the-shelf solutions, which means they are concepts that typically have to be adapted to a system's specific context of use and UX requirements.

While we discuss elements of privacy design patterns in Section 5.4, unfortunately, we don't yet have a good set of tested design patterns for privacy interfaces that we can recommend. There are certainly many design patterns for privacy interfaces that are widely replicated across many products and websites, but some of these actually suffer from fundamental usability problems.[50] For example, many cookie banners ask users to accept a website's use of cookies without actually offering a choice. This could actually be considered a forced-action dark pattern. Better cookie banners either provide an informational notice without requiring user consent or offer users a meaningful choice. So, while we suggest looking for privacy patterns that you might replicate, choose your patterns carefully. Once you have designed and tested a privacy interface within your organization, try to use the same pattern when a similar interface is needed for other products and services your organization offers.

5.3.2.3 UX Evaluation: Validating That User Needs Are Met

Throughout the design process, designs should be iteratively evaluated with current or future users of a system. As such, the design phase and the evaluation phase are closely interlinked. The purpose of UX evaluation is to validate that the system's designs and prototypes indeed meet the user needs and requirements identified in the research phase. Evaluation methods are often the same or similar to the user research methods mentioned in Section 5.3.2.1, with the addition of A/B testing and production deployment of developed solutions. UX validation may include both quantitative assessments (e.g., log file analysis, interaction times, success rates) and qualitative assessments (e.g., perceived usability, perceived cognitive load, joy of use, comprehension), with both providing important insights to evaluate and further refine designs and potentially also the design requirements. We discuss common user research methods for evaluating privacy user experiences in Section 5.5.

5.3.3 Value-Sensitive Design

An important realization for designing privacy interfaces and user experiences that meet people's privacy needs and address their privacy concerns is that privacy constitutes a *value*, i.e., privacy is considered important by individuals and society, but also competes with other values and norms. In addition, privacy preferences, concerns and expectations are not uniform but rather context-specific, malleable and fraught with uncertainty, as we discussed in Section 5.2. Privacy can therefore be more difficult to specify as a design requirement and measure whether a design achieves the desired privacy compared to other usability or UX aspects, such as efficiency or error rates.

Value-sensitive design is a design approach that accounts for ethical values, such as privacy, in addition to usability-oriented design goals.[51] Value-sensitive design methods help to systematically assess the values at play in relation to a specific technology and respective stakeholders and the ways the technology might meet or violate those values. They also help to iteratively develop designs that are sensitive to and respectful of those values. Friedman et al. recommend the following steps to integrate value sensitivity into design practice.[52] As you will note, many of the steps are also relevant in PIAs as well as in UX research and design in general. The difference is that value-sensitive design places at the center of the design process not just people and their needs but also values important to them.

1. *Clarify project values*—What values do a project and the project team aim to support? What do privacy, informed consent, transparency or other privacy-related values mean for this project and team?
2. *Identify direct and indirect stakeholders*—A value-sensitive approach to stakeholder analysis aims to identify stakeholders directly impacted by technology (e.g., data subjects, data processors) as well as those that are indirectly affected (e.g., bystanders incidentally recorded by the sensors of a self-driving car passing by or other household members that may use a smart home device or be recorded by its sensors). Privacy needs and expectations may also vary based on stakeholders' characteristics and group identities. Individuals may be part of multiple stakeholder groups.
3. *Identify benefits and harms for stakeholders*—What are the potential benefits or harms for each identified stakeholder group? Benefits and harms should be considered on an individual, societal and environmental level. This analysis can include both conceptual as well as empirical investigations, e.g., semistructured interviews or other user research methods. In interviews, a simple but illuminating practice is to ask why when people express positive or negative sentiment about a system or design in order to more deeply understand their reasoning and motivations or concerns.
4. *Identify and elicit potential values*—Identified benefits and harms are a starting point for identifying corresponding values. This mapping can be straightforward, e.g., an unanticipated data-sharing practice affects privacy. But they can also be indirect, e.g., surveillance practices may have a chilling effect, curtailing people's self-expression.

5. *Develop working definitions of key values*—Define what constitutes a specific value and spell out the components that make up the value in order to operationalize it for a specific context. For instance, informed consent is composed of, on the one hand, discovery, processing and comprehension of information and, on the other hand, voluntariness, comprehension and agreement.[53] Rather than inventing new definitions for important values, it is useful to rely on prior definitions and research in the literature to more deeply understand a value and its components.
6. *Identify potential value tensions*—Values do not exist in isolation and frequently conflict with each other as well as other requirements. However, value tensions rarely pose binary tradeoffs. (For example, "you can have security or privacy, but not both.") Rather, they may place constraints on potential designs. (For example, "how might we satisfy the security requirements while also respecting the privacy requirements.")
7. *Value-oriented design and development*—Value-sensitive design focuses not just on the design of technology but rather on the co-evolution of technologies and social structures. In the case of privacy, this means considering the interplay of technological solutions, regulatory solutions and organizational solutions and their social impact when trying to resolve identified value tensions. The implications of technologies on values can be investigated early on in the design process through value-oriented mockups and prototypes that are evaluated or potentially deployed with users. The ability to iteratively address and refine technologies and user experiences to align with stakeholders' values can be supported by building flexibility into the technical architecture so necessary changes can be made easily.

5.4 Design of Privacy Interfaces

While the previous section provided an overview of UX design more generally, we now focus on the design of user experiences specifically for privacy. We first discuss prevalent types of privacy interfaces, followed by common usability issues of privacy interfaces, before describing principles for usable privacy design and a process for implementing those principles in design and engineering practice.

5.4.1 Types of Privacy Interfaces

From an information-processing perspective, the data life cycle consists of multiple phases: data collection, data processing, data retention, data transfer and data deletion.[54] From a UX perspective, these phases need to be accompanied with privacy interfaces that implement established privacy principles and user rights:

- *Transparency* about data collection, processing and sharing practices, as well as privacy and security measures taken to protect user data, and information about who in an organization is responsible for privacy oversight and how to file complaints. Transparency is typically addressed through privacy notices.
- *Choice* regarding (nonessential) data collection, processing and sharing practices, including secondary uses of data. Choice is typically provided through consent interfaces.
- *Control* over the use and exposure of data, e.g., to other users on a platform. User control is typically provided through privacy settings.
- *Access* to the data an organization has about the user in order to facilitate inspection, correction, deletion and export of that data. User access is typically provided via privacy dashboards.

Next, we provide an overview of common privacy interface types for each category.

5.4.1.1 Privacy Notices

Privacy notices aim to provide transparency about an organization's data practices and other privacy-related information, such as measures taken to ensure security and privacy of users' information. Privacy notices need to be provided to users—typically before a data practice takes place—and explain what information about data subjects is being collected, processed, retained or transferred for what purpose. Laws and regulations in different countries may pose specific transparency requirements in terms of what information needs to be provided, when it needs to be provided and how it needs to be provided. Privacy notices can take different shapes and forms:

- *Privacy policies*—Privacy policies are probably the most common type of privacy notice. A privacy policy holistically documents an organization's data collection, processing and transfer practices and also includes other privacy-related information. While most common, privacy policies are also among the most ineffective privacy user interfaces. Most people do not read privacy policies, as they have little incentive to do so. Privacy policies are typically long documents. Aleecia McDonald and Lorrie Faith Cranor estimated in 2008

that reading all privacy policies of the websites an American internet user visits in a year would take about 244 hours.[55] Thanks to the proliferation of mobile apps, wearables and internet of things (IoT) devices since then, the number of privacy policies one would have to read and the actual time required for that are likely much higher today. Privacy policies further require advanced reading skills because of heavy use of legal jargon and often remain at an abstract level, especially when organizations use a single privacy policy to cover a multitude of services or products with many different data-processing practices. Thus, even if a user reads a privacy policy, they may not understand how data about them is actually collected, used or transferred in the context of a specific transaction. Finally, the effort required for reading and understanding a privacy policy is rendered useless by the fact that most privacy policies also contain provisions that the policy may be updated at any time.[56]

- *Informational privacy resources*—Organizations often complement privacy policies with informational resources that summarize or highlight important data practices and aim to communicate the value and role privacy plays for the organization to its users. These resources serve the purpose of educating users about how their data is used and protected, as well as what privacy choices and settings are available to them. The content should be consistent with the organization's privacy policy but presented more concisely and often with visual aids or videos to make the information more accessible and easier to understand for a general audience. While better at informing users about data practices than privacy policies, these informational privacy resources still need to be actively sought out by users, which means most users won't see them.
- *Integrated privacy notices*—Privacy policies and privacy information resources are often decoupled from a user's interactions with a service and product—a user must seek them out to learn the pertinent information. Integrated privacy notices are instead presented in a relevant part of the service or product's UX. For instance, an account creation process may include short explanations about how requested information (e.g., email address) will be used and protected.
- *Privacy indicators*—Privacy information can also be conveyed via privacy indicators—cues in a user interface or on a device. Privacy indicators are particularly useful for conveying either the state of a data practice (e.g., an LED indicating when a video camera is recording) or the state of a privacy setting (e.g., a small icon in a social media site's message posting interface indicating the post's audience).

- *Privacy reminders*—Organizations may choose or be required to remind people about data practices they are subject to. For instance, financial institutions in the United States are required to provide an annual privacy notice to their customers under the Gramm-Leach-Bliley Act (GBLA). However, privacy reminders can also take a more proactive shape and make users aware of data practices they had previously agreed to or nudge them to check and update their privacy settings.

5.4.1.2 Consent Interfaces

User consent is an established way to legitimize a data practice. To be valid, consent needs to be a freely given, specific, informed and unambiguous indication of an individual's agreement.[57] This means consent should pertain to a single data practice rather than be bundled together with consent for multiple data practices. Individuals need to be provided with sufficient information to make a consent decision, and it must be equally possible to agree or disagree. Consent interfaces typically need to enable users to provide or deny initial consent as well as check their consent status and possibly change it, which may mean revoking consent.

Consent interfaces generally support two kinds of consent: opt-in consent and opt-out consent. In the case of an opt-in, consent is obtained before a data practice takes place, i.e., the data processing the consent request pertains to does not start until the data subject explicitly consents. In the case of an opt-out, the respective data practice occurs as if the data subject had provided consent. A person's use of a service or product is assumed to be an expression of agreement with the respective data practices. An opt-out then is an option for the individual to stop a certain data practice. Opt-in consent can serve to make users aware of a data practice for which they have a choice they might otherwise not be aware of, and they might want to object to. Opt-outs carry the risk of a user not being aware of a data practice and being surprised or angry when they find out the data practice had been going on without their awareness or agreement.

Consent interfaces can be implemented in many different ways—most privacy settings are also consent decisions of some kind—but there are some general types of consent interfaces that can be distinguished:

- *Integrated consent prompt*—Opt-in consent requests are typically integrated into a system's UX. Examples include requests to accept a product's terms of service or privacy policy before use, or checkboxes to opt into data use for nonessential purposes during account creation. Requiring users to uncheck prechecked boxes to opt out of a data practice during setup is a discouraged dark pattern. The risk is that people may overlook the option or see it as a recommendation

and continue without properly considering it. While this might mean that more people keep a data practice (i.e., they do not uncheck the prechecked box), it is useless as a documentation of informed consent because it cannot be guaranteed that a user made an actual decision to keep the practice.

- *Decoupled opt-out*—Opt-outs might be decoupled from a user's interaction with a system, e.g., when they are described in a privacy policy, or in separate advertising or cookie policies, or when they are part of privacy settings. The challenge with decoupled opt-outs is that people have to seek them out and may not be aware of their existence or of the data practice they pertain to.
- *Integrated opt-out*—Certain opt-outs are integrated with the UX. One example is an Unsubscribe link in email communication. The advantage of integrated opt-outs is that they are present in the context in which people might need and want to use them.
- *Delegated consent*—Sometimes consent is not directly obtained by the first party but rather by a different service or platform provider. One example for delegated opt-in consent is a mobile permission. Apps do not directly ask users for access to resources on the smartphone (e.g., location, contacts, text messages), but instead programmatically declare required permissions to the smartphone. The smartphone operating system then generates a respective permission prompt, asking the user to accept or deny a permission request. An advantage of this model is that the presentation and functionality of permission requests are consistent regardless of the underlying app, which facilitates learnability of the interface and reduces cognitive load. An example of delegated opt-outs are opt-outs for ad targeting, which are often provided by third parties—either a specific ad company (e.g., Google, Facebook) or an industry association such as the Digital Advertising Alliance (DAA), Network Advertising Initiative (NAI) and the Interactive Advertising Bureau (IAB). A challenge with such delegated opt-outs is that the first party has little control over the usability of the consent mechanisms.

5.4.1.3 Privacy Settings

Privacy settings typically aggregate the privacy choices and controls available to a user of a given product or service in one place. Privacy settings are typically available from within account settings or referenced in the privacy policy, ideally both. Privacy settings interfaces can integrate a multitude of controls, including controls to inspect and change previously made consent decisions, settings to regulate the granularity of information,

settings to regulate the visibility of one's information (e.g., what information is public, visible to friends, or private), and settings to manage the sharing of information with users or with other apps or services. We distinguish two types of privacy settings:

- *First-party privacy settings*—Privacy settings commonly refer to settings made available by the service provider or product manufacturer. The service provider controls what settings are made available, how they are made available and how they are integrated with the overall UX.
- *Platform privacy settings*—What information an app or service has access to is also often controlled by the privacy settings of the respective platform used to provision the service, such as a mobile operating system for mobile apps, a smart assistant platform for voice-based skills or actions, or a web browser. Platform privacy settings may be specific to a certain service or app (e.g., app X can access location; app Y can't) or may allow users to set general privacy settings (e.g., whether third-party cookies are accepted or tracking protections are enabled).

5.4.1.4 Privacy Dashboards

Privacy dashboards typically give users access to the data an organization has about them. This may be in the form of activity timelines, data summaries or access to the actual data. Privacy dashboards should further provide support for other mandated user/data subject rights, such as enabling users to correct data or request rectification of inaccurate data. Users should further be able to export their data as well as delete some or all of it.

Often, privacy dashboards not only facilitate data access, edit, deletion and support but also include all of an organization's privacy settings, and possibly privacy information resources and privacy policies, in order to create a single destination for users for all things privacy.

5.4.2 Common Usability Issues of Privacy Interfaces

While privacy interfaces for transparency, control and support of different user rights exist, many privacy interfaces still fail to help people make informed decisions with respect to the collection and processing of their personal information.[58] This is problematic not just for consumers but also for companies, as it may lead to users being surprised by unexpected data practices and dissatisfied by frustrating user experiences, and may foster general distrust in the company. Many usability problems of privacy interfaces can be traced back to several common issues.

5.4.2.1 Conflating Compliance and User Needs

Frequently the creation of privacy interfaces, including privacy notices, controls and privacy settings or dashboards is driven by the goal of achieving compliance with respective legal and self-regulatory requirements. As a result, privacy interfaces, and in particular privacy policies, are created in a way that will demonstrate compliance to regulators. In turn, regulators such as data protection authorities or consumer protection agencies rely on an organization's privacy notices and controls in their investigation and enforcement of regulatory compliance.[59] Unfortunately, users' privacy needs are often neglected in the process. Privacy notices are written by lawyers for lawyers, resulting in lengthy privacy policies or terms of service that are necessarily complex because the respective laws, regulations and business practices are complex.[60] In addition, organizations may create policy statements that are purposefully vague to create leeway for potential future uses of collected data, which makes notices even more difficult for users to understand.[61] Privacy notices are not written as tools for creating awareness and transparency for users. Similarly, privacy controls are implemented according to regulatory requirements, but may not be sufficient to meet the actual privacy control needs of users. Yet, those privacy notices and controls are treated as if they were providing transparency and control to users.

Designing for users' privacy needs does not require neglecting or ignoring regulatory requirements. Rather, it is essential to distinguish between the information that must be provided for compliance reasons and the information that users might need to make informed privacy decisions—they are sometimes the same, but often not. As users are unlikely to read privacy policies, relevant information needs to be provided to them through additional means in order to reduce surprise. Similarly, privacy controls and choices required by regulation must be provided in a usable and useful manner. Furthermore, additional privacy controls may be necessary to ensure that users can control privacy in alignment with their context-specific privacy preferences and expectations.

5.4.2.2 Lack of Meaningful Choices

Too often, users are forced to make impossible choices: Either accept everything in the terms of service and privacy policy or do not use the service or product. Such take-it-or-leave-it choices are bad for both consumers and organizations. Consumers might accept the terms of service or privacy policy but are unlikely to have read or understood them. This fosters surprise about unexpected data practices and privacy harms, which in turn can result in outcries on social media and in the press and lead to consumer distrust, damaged reputation and possibly investigations by regulators for deceptive practices or

violations of transparency requirements if it is not a reasonable expectation that users would be aware of a practice.

Additionally, forcing users to accept data practices they disagree with, because they otherwise cannot use a service or product, fosters digital resignation.[62] If you rely on consent as the legal basis for processing, it is important to ensure that consent is specific, informed and freely given.

5.4.2.3 Poor Integration with UX

Privacy notices and consent prompts are often shown at inopportune times, without regard for the user's primary task. Providing users with all privacy-relevant information when downloading an app or signing up for a service is common practice, but futile from a UX perspective. When a consumer decides to download an app or sign up for a service, they do so because they want to use the app or service, and anything delaying use—such as asking to accept a long privacy policy or configuring privacy settings—is likely to be dismissed or accepted quickly without much scrutiny. Privacy management is almost always a secondary task, which means privacy interfaces need to be designed to blend in and augment the primary UX rather than blocking it.

In addition, repeated exposure to seemingly irrelevant privacy notices or dialogs results in habituation. This means that people dismiss the notice or prompt automatically without its content even registering in their cognitive processing system.[63]

5.4.2.4 Poor Discoverability

Often privacy notices, opt-outs and other controls are decoupled from a system's primary UX. This dissociation of privacy interfaces from a system has severe usability consequences. Without being exposed to privacy interfaces as part of the system's UX, users require substantial digital literacy and an advanced mental model of how a certain technology works in order to anticipate a system's data practices and what privacy information or privacy controls might be available to them.[64]

Furthermore, if a user wants to seek out privacy information or controls, they must know where to look. However, where privacy information and controls are provided is not consistent across services and products. Some might have account settings that include privacy settings, some might have separate privacy settings, others might have privacy settings but also provide additional privacy-related controls under other settings or in policy documents.[65] As a result, people may have to look in multiple places, search in the wrong places, or even worse, assume that they found all relevant privacy controls without realizing that additional controls are provided in another place.

5.4.2.5 Confusing Interfaces

When privacy interfaces are not designed carefully, they may be confusing and suffer from usability issues that could have been uncovered through user testing. Some common issues include wording or signage (e.g., an icon) that is confusing or ambiguous, privacy choices and opt-outs whose effects are unclear, or privacy controls that behave contradictory to expectations. [66]

5.4.3 Privacy Design Principles

Following a set of privacy design principles and a systematic process makes it possible to design privacy user experiences that are usable, useful and meet users' needs for information and control while also being compliant with privacy laws and regulation.[67]

5.4.3.1 User Centric

Privacy legislation, privacy technology and privacy management are centered on information. What information flows and practices are permissible? How can information be protected, de-identified or processed to enhance privacy? What information does the organization collect, store, process and transfer, as well as what safeguards and processes are in place to ensure accountability and compliance with internal and external policy? Designing usable and useful privacy user experiences, however, requires a shift in focus: While information is of course still relevant, the focus needs to be on users and their interaction with a system's privacy interfaces. Privacy design must be user centric to be effective.

A UX is successful when it meets users' needs. For privacy interfaces, that means first understanding what the privacy needs of different stakeholders and user populations are, both their informational needs and their control needs. Identifying those needs requires investigating and understanding people's privacy preferences, concerns and expectations with respect to a specific information system. It further requires understanding users' mental models of the technology, its data practices and its privacy protections. Such insights help determine what information may be necessary to help users make informed decisions or create awareness of otherwise unexpected data practices. Furthermore, it is important to understand users' privacy control needs in terms of managing privacy toward the organization and toward other users on multiuser platforms or to regulate the visibility of their data.

When designing privacy notices and controls, it is helpful to be aware of how humans receive, process and react to information. Wogalter's communication-human information processing (C-HIP) model explains how humans perceive, process and react to warnings.[68] The Human in the Loop (HILP) model adapts the C-HIP model

for security (and privacy).[69] These models describe a similar progression of steps in human information processing, which, if not considered properly, can present a hurdle for a privacy interface in achieving its goal. The key steps of the information processing model as they pertain to privacy are as follows:

1. *Communication*—Human information processing begins with how information is communicated to a person. A user interface aims to communicate something to the user, e.g., certain privacy information or the availability of a privacy control. This communication may be affected by environmental stimuli that compete with the privacy interface; for example, the user's primary task or other interface elements or system features may make it difficult to notice the privacy control. Interference may prevent users from even receiving the communication; for instance, being distracted by other activities, people nearby, or ambient noise when privacy information is shown (e.g., in a video) may cause a user to miss it.
2. *Attention*—Users need to notice the privacy interface before they can process it. Thus, effective communication has to get users to switch attention to the privacy interface and maintain attention long enough to process it. Environmental stimuli, interference and characteristics of the communication (how it is presented and delivered) affect attention switch and maintenance. Habituation also impacts attention, i.e., the effect of a stimulus decreases over time. For example, presenting too many or repetitive privacy notices to users will likely result in them ignoring the notice content.
3. *Comprehension*—Once users notice an interface, they need to understand its purpose. What is it there for and what is it communicating to the user? A person's digital literacy and privacy awareness may affect how they interpret provided privacy information or controls. Are they aware of the privacy risks and implications associated with the options presented in a consent prompt? How accurate is their mental model of the system's data practices and privacy protections?
4. *Intention*—Based on a user's comprehension of the interface, they may or may not be motivated to act upon the provided information or engage with a privacy control. The user's intention is further affected by their privacy preferences, concerns and expectations.
5. *Behavior*—As discussed in Section 5.2, actual privacy behavior can deviate from the user's respective privacy decision. Usability issues with the

presentation of choices as privacy settings and controls can result in users making errors.

User-centric design of privacy interfaces should consider how to best support each of these phases. When trying to analyze usability issues of privacy interfaces, it is further helpful to isolate which information processing phases are affected in order to address them properly.

5.4.3.2 Relevant

Privacy policies are ineffective at supporting privacy decision-making because they aim to provide a complete and comprehensive overview of an organization's data practices. While such an overview is useful for regulators and researchers to analyze an organization's data practices, it holds little value for most consumers. Consumers are asked to complete a complex cognitive task in trying to map a policy's abstract descriptions of data practices to the specific context and transaction in which they are engaging with the respective organization.

To be useful for privacy decision-making and behavior, privacy information and controls need to be relevant to the user's specific context or transaction. For example, for someone visiting a news website for the first time, there is very little use in reading the news site's privacy policy completely. At this point, it would not be relevant to them how the news site uses contact information provided during account setup, how financial information provided when signing up for a subscription is processed and secured or what data is transferred when one enters into a sweepstakes the news site might offer together with a marketing partner. Instead, when visiting the news website without an account, it would be most relevant to the user whether the news site tracks what articles are being read, how that information is used (e.g., for site analytics, ad targeting), whether and which other third-party trackers are present on the news site and whether there is an option to opt out of such tracking. However, finding these pieces of information in a privacy policy is tedious, because the information is often organized around data practice categories (e.g., collection, use, sharing) rather than common types of interaction contexts (e.g., news reading without account, account creation, site use with account, subscription, sweepstakes).

While organizing information in a privacy policy according to typical user contexts can be helpful, rather than hoping that users might find the relevant statement in the privacy policy and then arrive at the accurate interpretation of the statement, the information or controls that are relevant to a particular context should be provided within the respective UX. This does not mean that every part and step of a system's user interface should be covered in privacy policy text. Instead, relevant information and

controls should be easily accessible from within a context and organized according to the steps in user journeys typical for different user populations.

Furthermore, the information and controls provided to users should support their understanding of specific data practices. Rather than providing abstract descriptions, privacy notices should be explicit about what information is specifically being collected, processed or shared in a given context or transaction, why this data practice is necessary and how it benefits the data subject (if at all), and what controls are available regarding the practice. Ideally, privacy notices and controls would not just be specific to a data practice but also to the individual user. For example, Harbach et al. proposed improving mobile permission dialogs by including personal examples from the user's phone of what an app would gain access to if a request were granted (e.g., a location request would show the current location, a photo access request would show thumbnails of a random selection of the user's photos).[70]

Users should also be informed about potential privacy risks associated with certain data practices. While it might seem counterintuitive to emphasize risks associated with one's system—organizations like to highlight the benefits of their products, not associated risks—openly communicating risks can have multiple advantages. First, some users might be thinking of risks anyway, but may overestimate them. Open risk communication can help users calibrate their risk assessments. Second, users who are not aware of the risks might be surprised and angry if they find out about the risks later or are affected by them. Open risk communication can help mitigate surprise and help shape user's mental model of a system. Third, openly communicating risks is also an opportunity to actively communicate the protective measures in place to mitigate the risk and protect the user's privacy.

5.4.3.3 Understandable

When presenting privacy information or providing privacy controls, it is important that the information and controls are understandable by the users they are targeting. This UX best practice is starting to find its way into privacy regulation. For instance, the GDPR requires that "any information and communication, where processing is addressed to a child, should be in such a clear and plain language that the child can easily understand."[71]

In practice, understanding privacy policies frequently requires advanced reading skills. In 2019, as part of their Privacy Project, *The New York Times* analyzed the privacy policies of 150 popular websites and apps.[72] They found that most analyzed policies required college-level reading skills. To put this into perspective, they compared the reading skills required for privacy policies with those of popular books. Most of the analyzed policies required higher reading skills than are necessary for understanding Stephen Hawking's *A Brief History of Time*—a book about the space-time continuum. Some policies even

exceeded the reading level required to understand Immanuel Kant's *Kritik der reinen Vernunft (Critique of Pure Reason)*—a seminal work in Western philosophy. Privacy policies are also too long to be useful for consumers. A longitudinal analysis of the 500 most popular websites in each EU member state (over 6,750 websites in total) found that privacy policies were on average 3,044 words long after the GDPR went into effect in May 2018— 42 percent longer than in 2016 when the GDPR was ratified.[73]

Clearly, users cannot be reasonably expected to fully read the privacy policy of each service, website or app they use. Instead, by focusing on providing user-centric information that is relevant to a user's specific context, transaction or activity, the provided information can be much shorter and more specific, and thus also be more understandable. Such information would complement the privacy policy rather than replace it, which would also mean the privacy policy could be written specifically for regulators and privacy experts without negatively affecting users.

User-centric privacy notices, such as integrated privacy notices and informational privacy resources, should be written in simple language that is understandable without advanced reading skills, avoiding legal and privacy jargon. For reference, in healthcare, it is recommended that patient educational materials should be written for sixth- to eighth-grade reading skills.[74] The readability of documents, including privacy notices, can be tested with online readability checkers, such as Readable.com.

An important aspect in achieving understandable privacy notices and controls is to find the optimal level of abstraction that best aligns with users' privacy needs. For example, cookies and cookie consent notices (commonly known as cookie banners) have received significant attention because of European legislation. However, providing notices about cookies places the emphasis on the data collection mechanism, whereas people tend to make decisions based on the purposes of data practices. People may not care about cookies specifically, but they might care whether they are being tracked regardless of the specific tracking technology, suggesting that the consent notices may be more effective if they focused on the purposes of cookie use.

5.4.3.4 Actionable

To be useful, information needs to be actionable.[75] It is typically not advisable to give users a privacy notice without an associated action or choice they can take. While privacy information can support users in making more informed privacy decisions, as well as help correct potentially misaligned mental models or expectations, without something to do with that information, there is not much point in the increased awareness. Users have little incentive to engage with provided information unless there are also options for them to realize their privacy decisions. Therefore, privacy information and privacy controls should go hand in hand whenever possible.

The choices offered also need to be meaningful. Asking users to accept a privacy policy that contains many different data practices is not a meaningful privacy choice. Rather than providing users with agency over their privacy, it asks them to either not use the system or accept even if they disagree with statements in the privacy policy.

Instead, choices should not be bundled, and any kind of feature limitations due to the user's choice should be constrained to the immediate context of that choice. A reasonable example of this is mobile permissions in recent mobile operating systems. Users are presented with a concise prompt when an app wants to access a phone resource for the first time (e.g., location, contacts, photos). Users have the option to allow or deny the request and can continue to use the app even if they select "deny," possibly with the exception of certain features. Ideally, a privacy choice or prompt asks users to make a meaningful decision by providing them with options to choose among. Prompts that ask for confirmation only increase the risk of habituation—users clicking the prompt away without reading it—which in turn makes it difficult to interpret their action as an expression of freely given and informed consent.

5.4.3.5 Integrated

Privacy information and controls should be integrated with a system's primary UX rather than added on. Relevant information should be provided at points in a system's UX where users actually make privacy decisions, or where it might be important to make users aware of data practices or privacy risks in order to help them develop an accurate mental model of the system and its information flows. Similarly, privacy controls and indicators (e.g., to see and control a social media post's audience) should be available at interaction points when they matter (e.g., when writing the post).

Whenever possible, privacy interfaces should use the same interaction methods as the system's primary UX. This ensures that interacting with privacy interfaces feels natural in a system's context and is as easy as any other user interaction with the system. For example, when designing privacy controls for a voice-controlled system such as a smart assistant, privacy controls like "mute" should also be available as voice commands, because the user otherwise needs to make a shift in interaction modality (walk to device and press a physical mute button instead of speaking to the device) that deviates from their typically interaction with the system and may therefore be rarely used.[76]

Furthermore, privacy information and controls that are integrated into the UX can be more concise—both in terms of presented information and choices—than privacy policies and privacy settings because they can rely on the user's interaction context to contextualize the presented information for the user. The proper integration of privacy interfaces into the overall UX can help prevent overloading users with information or presenting privacy prompts at inopportune times.

However, while integrated privacy interfaces are important, a system should still provide a privacy policy, informational privacy resources, privacy settings and privacy dashboards in a centralized place so that users (and regulators) can gain a more comprehensive overview of data practices or verify and adjust their privacy settings in a central place when necessary.

5.4.4 Privacy Design Process

To put the privacy design principles into practice, it helps to follow a systematic privacy design process. We discuss a general process that combines UX, PIA, and value-sensitive design.[77] This process consists of six steps:

1. *Build on privacy assessment, privacy management and privacy engineering practice* to systematically identify a system's user rights and transparency requirements
2. *Identify users and their privacy needs* by identifying stakeholders and eliciting their privacy expectations and privacy concerns as well as their privacy information needs and privacy control needs
3. *Identify unexpected data practices*, which are those that users are unaware of or might be surprised by, to help prioritize which data practices and controls to highlight
4. *Integrate privacy interfaces into system's UX* by determining which privacy notices and controls are most relevant to a user at which points in the UX
5. *Leverage the available design space* for privacy notices and controls to develop user-centric privacy interfaces that work within a system's constraints
6. *Conduct user testing* to evaluate the usability and usefulness of developed privacy interfaces

Next, we describe each step in detail and provide practical guidance. The process can and should be adapted to the needs and circumstances of your organization, systems and users.

5.4.4.1 Build on Privacy Assessment, Privacy Management and Privacy Engineering Practice

The essential precursor for designing usable and useful privacy user experiences is a solid understanding of the system's information flows and data practices, as well as the associated privacy implications. Organizations should be conducting PIAs, risk assessments and compliance checks (with internal policies and external regulation) as

part of their internal privacy programs. Such assessments can provide a comprehensive overview of a system's data practices. The recommendations resulting from a PIA may include proposed changes to data collection practices, opportunities for data minimization or a requirement for opt-in consent for a certain data-sharing practice. Data minimization helps reduce the risk of using data in ways that deviate from users' expectations—you cannot misuse data you do not collect or retain. Diligent data minimization can further reduce the number or complexity of data practices that need to be communicated to users.

PIAs and compliance checks can also help systematically identify the user rights and transparency requirements pertaining to the system, which are important for the privacy design process. The outcome of a PIA is typically internal documentation of information flows, data inventories and a comprehensive privacy policy for a product or the organization. While the privacy policy is likely not read by users, it can serve as input for designing user-centric privacy notices and controls. Importantly, PIAs and the privacy policy are just the starting point for designing user-centric privacy interfaces.

Considering the privacy implications of a system and associated transparency and privacy control needs early on in a system's development also provides an opportunity to consider and improve constraints the system might pose for privacy interfaces. For example, recognizing that a smart home camera is collecting potentially sensitive information, the device designers may decide to include a recording indicator and a physical lens shutter in the device's hardware design rather than trying to provide notices and controls solely through software at a later stage.

5.4.4.2 Identify Users and Their Privacy Needs

A common part of PIAs is the identification and consultation of stakeholders whose privacy may be potentially impacted by a system. This can be combined with identifying a system's different types of users and eliciting their privacy expectations and privacy concerns as well as their privacy information needs and privacy control needs. For the design of privacy user experiences, it is important to distinguish at least three groups of users:

- *Primary users* are a system's intended audience. The primary user would make the decision to use a certain system and complete an enrollment process (e.g., account setup or configuration wizard) and would be able to configure the system. This may include activating or consenting to certain optional practices (e.g., activating face recognition for a smart home camera), opting out of certain practices (e.g., analytics), and adjusting privacy settings. For websites and mobile apps, the primary user would be the person using the browser or downloading the app; for IoT devices, it would be the person setting up the device.

- *Secondary users* are other people who may use a system in addition to the primary user. Secondary users might be aware of a system and even actively use it, but they may have an incomplete or inaccurate mental model of the system and its data practices and less control over it compared with the primary user. Examples of secondary users would be a child using a parent's smartphone to play a game or stream a video, or a household member using a smart speaker that's associated with the device owner's account or a smart thermostat that has been configured by someone else. Secondary users may need to be informed about data practices, or even offered controls, in addition to and in ways different from the primary user.
- *Incidental users* might inadvertently and unwittingly become data subjects of a system's data collection and processing. For example, a smart door lock might collect information about all family or household members, including guests.[78] House guests may be required to use a smart speaker to operate smart lights in the house.[79] Doorbell cameras, drones or self-driving cars might incidentally record people passing by. Social media and mobile applications enable users to share information with and about others, e.g., by uploading a geotagged photo of someone else without their knowledge. Incidental users may not always be identifiable by a system, which constrains the potential means for providing privacy notices and controls to them.

Depending on the system, other or additional user groups may need to be considered. There may also be specific regulatory requirements for data collection, notice and user rights regarding protected user groups, such as children.

System designers and privacy professionals need to understand how the privacy of each identified user group may be affected by the system. Often a system's user groups are affected by the same data practices. For example, a smart home device may record the same sensor information about the primary user and other household members. However, sometimes user groups may be affected by different data practices, or certain groups may only be affected by a specific data practice. For example, a smartphone's data practices affect mainly the primary user and possibly secondary users, whereas incidental users may only be affected if they are being recorded by the device, for instance, when a primary or secondary user takes a photo or video.

Next, it is important to understand each group's privacy preferences, concerns, expectations, mental models and control needs regarding the system and its data practices. This information can be gained by conducting user studies, which we discuss in more detail in Section 5.5, and potentially through reviewing academic literature that has investigated privacy preferences of specific user groups in specific contexts.

5.4.4.3 Identify Unexpected Data Practices

A particular focus when seeking to understand users' mental models and privacy preferences should be to identify which of the system's data practices people are unaware of and might be surprised about. Knowledge about unexpected practices helps determine what information and controls need to be presented more prominently in the UX. Data practices that are consistent with the transaction context (and therefore with the user's expectations) might not require immediate notice. However, practices that are unexpected or violate the transaction context should be highlighted. For example, Gluck et al. studied which of a fitness watch's data practices people expected. One finding was that almost all participants (94 percent) expected the watch to record steps, but only 31 percent expected the watch to also record location information. Recording steps is a key purpose of such a device and may not need to be specifically highlighted, whereas collection of location data was unexpected for many, indicating a need for user-centric privacy notices focusing on this aspect.[80] Making people aware of unexpected practices and ideally asking for consent (opt-in) reduces surprise and can potentially increase user trust.

Making users aware of an unexpected data practice is also an opportunity to explain the reasons for the practice, its benefits and the associated privacy risks.[81] It also provides an opportunity to communicate what measures are put in place to limit associated privacy risks and protect user privacy. Privacy and security engineering often happens under the hood—privacy UX design is an opportunity to make concrete the care and effort an organization puts into privacy and data protection.

5.4.4.4 Integrate Privacy Interfaces into System's UX

Once the information has been gathered about the system's data practices, its user groups and the practices that might be more or less expected by different user groups, the next step is to decide when and where to integrate privacy notices, consent prompts and other privacy interfaces into the system's UX. Showing lots of information and controls at once in a single interface is rarely effective. Instead, the goal is to map specific privacy information and controls onto those points in the system's UX where they are most relevant to the user and where an opportunity exists for the privacy interface to gain the user's attention. Such privacy interfaces can be more concise and require less interpretation by the user because they are integrated into the user's interaction with the system, which therefore makes them less disruptive. This might mean providing certain information and controls at multiple interaction points and with varying levels of detail.

The level of detail provided in a specific notice or control must be appropriate for the respective context and user group. However, notices and controls can be *layered*. A short notice may highlight a practice, fact or control to gain the user's attention and provide a

link to additional information or more controls, which in turn might consist of multiple layers that users can reveal and explore as needed. In UX design, this design pattern is called *details on demand*—providing an overview or summary first, with options for users to retrieve details.[82] Thus, rather than trying to provide a single notice or control that does everything, it's better to craft a privacy UX that is composed of many complementary privacy notices and controls at different levels of detail tailored to the respective user group and context.

A *privacy UX* combines privacy interfaces shown at different times, using different modalities and channels, and varying in terms of content and granularity in a structured approach. With such an integrated and multilayered approach, individual users still receive information and choices for data practices that pertain to them but won't be confronted with data practices that are irrelevant for them until they are using a respective feature of the system. Users should still have the option to read the full privacy policy and explore all privacy settings whenever they want.

Deciding what information or controls to include in a brief integrated notice presented as a top layer is a crucial aspect at this stage, because most users will not seek out more detailed privacy interfaces, even if they are referenced and only a click away. Thus, if such a short notice does not capture all relevant aspects of a data practice, it may hide information and impair transparency.[83] This is especially an issue for unexpected data practices. Therefore, privacy interfaces should be hierarchically structured in such a way that the initial notice, prompt or control either already captures the main aspects of the data practice or serves primarily to draw the user's attention to more expressive privacy interfaces. For instance, Almuhimedi et al. showed that just notifying smartphone users about how often their location has been accessed by different apps is an effective nudge for users to check and adjust their permission settings.[84] In contrast, a privacy policy that only shows section headings and requires users to click each heading to read the text is a bad design, because it requires users to make many more clicks and prevents keyword search.[85] A more effective approach in this context would be to summarize a section's key data practices and respective choices and provide a "learn more" button to reveal a detailed description.

Creating and maintaining a privacy user experience composed of multilayered privacy interfaces does incur engineering and management costs, but it has multiple benefits. Users' privacy needs and the system's data practices can be better aligned. Where user consent is required, this approach can yield properly informed consent. Reducing surprise and explaining privacy protections is likely to facilitate user trust.

5.4.4.5 Leverage the Available Design Space

Once it has been decided what information or controls need to be provided at which points in a system's UX, the next step is the actual design of the user-centric privacy interfaces. Multiple aspects can impose constraints on the design of the privacy user experience. A system has certain input and output modalities, which can affect how different user groups can or cannot be provided with privacy information or express their privacy choices.[86] For instance, many IoT devices have no screens or only very small ones, which may not be suitable for displaying privacy information or controls. However, a system and its context may also offer unconventional interaction opportunities for privacy interfaces. Identifying constraints and considering the full design space for privacy interfaces helps generate different design ideas for a privacy interface and identify viable solutions. Based on an extensive analysis of the design space for privacy notices and controls, Schaub et al. identified four design dimensions that should be considered in privacy interface design.[87] The design space with its main dimensions and their design opportunities is shown in Figure 5-1. We provide a short overview of each design dimension and respective key considerations.

Figure 5-1: The Design Space for Privacy Notices and Controls

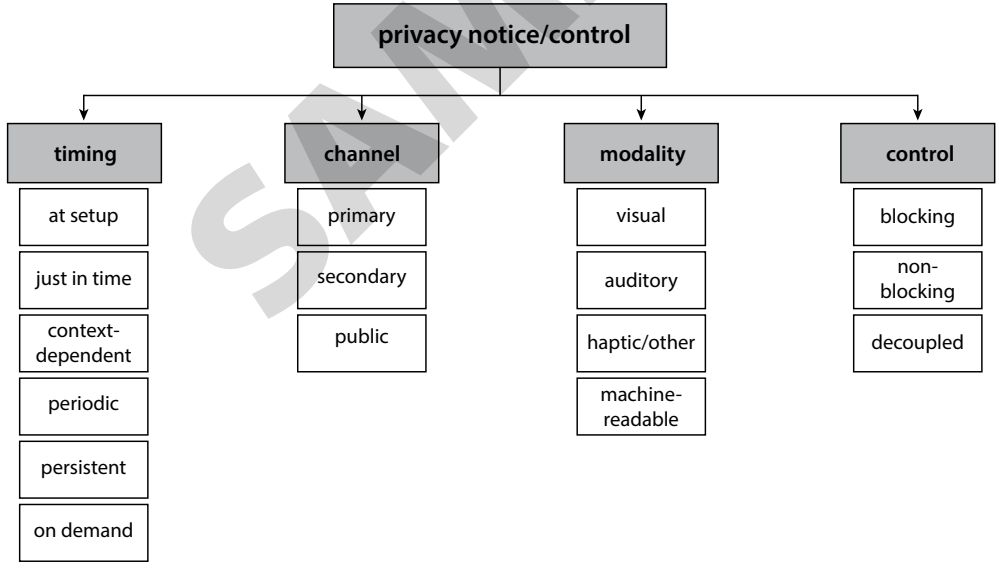

Timing. The point in time at which privacy information or controls are presented to the user has a substantial impact on a privacy interface's usability and utility. Users

may ignore privacy interfaces shown at inopportune times, but may seamlessly interact with a consent prompt that appears exactly when it is needed. Timing is a key factor in ensuring that a privacy interface is relevant and aligned with users' needs.

- *At setup* interfaces are often shown on initial use. However, only information and choices that are truly essential before use should be communicated at setup because users' attention is typically focused on the primary UX at this point. Those privacy information and choices that must be presented at setup should be properly integrated into the setup process (e.g., a configuration wizard or account enrollment) in a way that provides value and meaning to users rather than nudging them to accept a privacy policy or terms of service they have not read.
- *Just-in-time* interfaces are shown in the same transactional context as the data practice they pertain to, which supports reasoning in the moment and means they can be specific and short, communicating only the most relevant information and choices.
- *Context-dependent* privacy interfaces are triggered by certain aspects of the user's context.[88] For instance, being in physical proximity to an IoT sensor may cause the device to announce its presence (e.g., flashing an LED, beeping, sending a description of its data practices to the user's phone). Other context factors might be someone accessing previously uploaded information or analyzing the user's previous privacy settings behavior to warn about potentially unintended privacy settings. For example, Facebook displays a message warning that it is about to post publicly when the user's last post was public, but they are typically not in the habit of posting publicly.
- *Periodic* reminders are useful to remind users about data practices that they agreed to previously and to renew consent if the practice is still ongoing, especially if the data practice occurs in the background, invisible to the user.
- *Persistent* privacy indicators are shown whenever a data practice is active. For instance, cameras often have lights to indicate when the camera is recording. Persistent indicators can provide an unobtrusive cue about especially critical data practices, but there's also a risk that the indicator will not be not noticed.[89]
- *On demand* privacy information and controls allow users to seek out and review privacy information or their privacy settings and opt-outs any time. On-demand interfaces should be made available in a well-known or easily findable location (e.g., a website should have a "privacy" subdomain or "/privacy/" folder and provide links to privacy controls from their privacy policy and in relevant menus).

Channel. Privacy interfaces can be delivered through different communication channels.

- *Primary channel* is the primary device or system the user is interacting with, such as a smartphone or computer. It is the system that is also primarily involved in the respective data practice, e.g., collecting, storing or processing the data.
- *Secondary channels* are useful when the primary system is too constrained in its input and output capabilities to provide for meaningful interaction with the privacy interface. Typically, a secondary channel presents information on a different device or in a different context. For instance, privacy information and settings for a fitness tracker or smart speaker might be available through a companion app or website. Emails, mobile notifications and even letters can also serve as secondary channels for privacy communication.
- *Public channels* can be used when the system's users or data subjects cannot be reached individually or are not known to the system. For example, public signs can be posted to inform about camera use or other sensing technology in a certain environment.

Modality. Privacy interfaces can be presented in different ways, with different interaction modalities.

- *Visual* privacy interfaces are most common. Privacy notices are presented as text or with icons and illustrations; privacy controls rely on graphical user interfaces. Many aspects can affect the usability of visual privacy interfaces, including color, font, white space, combinations of icons and text, the layout of interface elements and how the state of settings is conveyed. Privacy concepts are difficult to represent as icons only; combining icons with explanatory text or using more expressive visualizations and videos is preferable.
- *Auditory* privacy interfaces use sounds or voice to convey privacy information, or voice commands to enable privacy controls. Certain sounds can convey data collection practices very effectively. Take, for example, a camera's shutter click, which is being replicated on digital cameras and smartphones to indicate when a photo is taken, even though those devices may not have a physical shutter.[90] Support hotlines often play an auditory notice before one is being connected with a human operator. (For example, "This call may be recorded for quality assurance purposes.") Voice-based interfaces, such as smart assistants, should offer voice commands to control privacy.[91]

- *Haptic and other* modalities may also be leveraged as privacy interfaces. For instance, device vibration could be used as an indicator for data collection. Olfactory displays could be used to inform about privacy risks with different scents (e.g., lavender scent when visiting a privacy-friendly website; sulphur when visiting a privacy-invasive one). Ambient lights could also serve as privacy indicators. Could taste or skin conduction be used in privacy interfaces? Although it might not be immediately apparent how less conventional modalities could be used for privacy interfaces, the important point is to creatively explore even unconventional design opportunities. Such exploration often leads to helpful insights that can inform practical solutions.
- *Machine-readable* specifications of privacy notices and controls enable the consistent presentation and aggregation of privacy information and controls from different systems or apps.[92] Mobile apps have to declare their permission requests in a machine-readable format, and the mobile operating system is responsible for providing permission prompts and managing the user's consent. IoT devices that lack screens and input capabilities could broadcast their machine-readable privacy notices to smartphones of nearby users, which then present the respective privacy information and choices to the user.[93]

Control. User choices, consent dialogs and privacy settings can be delivered in different ways that affect how users interact with them.

- *Blocking* privacy controls force users to interact with the privacy interface in order to be able to proceed. Blocking controls are useful when the user must make a choice, e.g., when consent is needed. However, how choices are presented affects whether the interaction is actually an expression of the user's preference, and so should be considered carefully. For example, presenting a complex privacy policy and providing only the options to accept or not use the app is not suitable for eliciting consent, as users are likely to click the warning away without reading. Preferable are prompts that are specific to a single data practice and provide options to both allow or deny the practice (e.g., mobile permission requests). All choices should be equally easy for the user to exercise.
- *Nonblocking* privacy controls do not interrupt the interaction flow but are rather integrated as user interface elements into the UX. For example, social media apps might provide an audience selector (e.g., private, friends, public) within the interface for creating a post. The control is available but does not have to be used and, at the same time, reminds the user of their current privacy settings.

- *Decoupled* privacy controls are not part of the primary UX. They are useful to provide the user the opportunity to inspect their privacy settings or the data the system has collected about the user. Common examples are privacy settings and privacy dashboards. The advantage of decoupled privacy controls is that they can be more comprehensive and expressive than integrated controls; the downside is that users need to actively seek them out. Good practice is to provide decoupled privacy controls at a central place and then point to them from other, more concise privacy notices and controls where appropriate.

Frequently, multiple kinds of privacy notices and controls will need to be combined in a multilayered privacy user experience in order to properly address privacy information and control needs of a system's user groups.

5.4.4.6 Conduct User Testing

Privacy user experiences should be developed in a user-centric and iterative design process. Potential users should be involved in the design process from the beginning to account for user needs and to evaluate whether developed privacy notices or controls are usable. Privacy interfaces can be evaluated individually (e.g., by asking whether users understand a particular consent prompt), but the privacy user experience should also be evaluated holistically to gauge whether the combination of privacy interfaces appropriately serves users' needs and how the overall privacy user experience affects user trust in the system and its handling of users' personal information.

5.5 Usability Testing and User Studies for Privacy

Usability testing and user studies are a staple of both research and product development. Many books and guides have been written about how to conduct user studies. In this section, we provide a brief overview of the what, why and how of user testing and user studies as they relate to privacy. Our goal is to convey the importance of user studies for the design of privacy interfaces and give privacy professionals the necessary knowledge to effectively collaborate with UX professionals to design appropriate privacy studies. At least in the beginning, it is useful to work with UX professionals to ensure that a study actually tests what it is intended to test.

5.5.1 Why Test Usability?

Usability testing and user studies play several important roles: they can help to *assess needs, examine tradeoffs, evaluate designs* and *find root causes* of problems.

At the beginning of a design process, exploratory user studies can inform design requirements by providing feedback on which user needs are not being met by existing systems, as well as identify users' pain points and points of confusion. For example, before designing a privacy dashboard, a user study can help designers gain insights into what data users are most interested in seeing and at what level of granularity, as well as what types of controls users are most interested in exercising. Furthermore, designers may learn whether they are likely to satisfy most of their users with a single interface, or whether special interfaces might be needed for certain populations such as advanced or novice users.

Designers often face tradeoffs, and user studies can help provide insights into the pros and cons of each approach under consideration. For example, designers may realize that more granular privacy controls may increase the complexity of privacy choices. Through usability testing, they can gather data that will allow them to weigh the benefits of providing more granular controls against this added complexity. They can observe whether users actually take advantage of the added granularity to adjust settings that best match their personal privacy preferences, and they can observe how much longer it takes users to configure the controls and how often users give up because the controls are too complicated.

While following established usability guidelines and principles is likely to improve usability, the only way to be sure that a system is actually usable by its intended audience is to conduct usability tests with representative users. Thus, usability testing is critical for determining that a design actually meets its stated requirements. Before declaring that their privacy settings allow users to exercise privacy choices quickly and easily, organizations should run usability tests that assess whether users can find privacy settings (discoverability), use them to select the settings that match their personal preferences and understand the effect of the settings they selected.

Finally, after discovering a usability deficiency, usability testing is useful to determine the underlying cause of the problem, which will hopefully lead to solutions. For example, a company may receive complaints from users that their privacy settings are not working. After determining that there is no underlying technical problem, the company may conduct user studies to understand how users believe the privacy settings are supposed to work and where their mental models differ from what the system actually does. This should provide insights that will allow the company to redesign its privacy settings to better align with users' expectations and privacy control needs or to communicate its current function more clearly.

5.5.2 What to test?

The details of what to test depend on the type of system being tested and where in the design process the test is being conducted. Exploratory testing early in the design process may be less about testing a user interface and more about understanding users' preferences, expectations, concerns and mental models. Once a prototype or complete interface has been developed, testing will generally focus on some or all of the following: awareness and attention, discoverability, comprehension, utility and behavior. However, testing may also investigate learnability, efficiency, memorability, errors, satisfaction and other aspects of UX and information processing phases discussed in Sections 5.3.2 and 5.4.3.

5.5.2.1 Preferences, Expectations, Concerns and Mental Models

System designers benefit from an understanding of users' privacy preferences, expectations and concerns, as well as their relevant mental models. As discussed earlier, preferences, expectations and concerns collectively impact a user's privacy decision-making. In turn, they may impact the types of privacy features and controls users seek out or the types of explanations or assurances that may be needed to ensure that users' concerns are addressed. Surveys, semistructured interviews, and focus groups can all provide insights into users' privacy concerns and needs in a specific context. It can be valuable to conduct these exploratory studies early in the design process.

Numerous papers in the academic research literature explore users' privacy preferences, expectations, concerns, and mental models generally, and in the context of specific technologies (e.g. behavioral advertising, health technology, IoT).[94] Such research papers can provide insights to inform the design of privacy interfaces, supplementing or possibly even eliminating the need for exploratory studies in a specific design process.

5.5.2.2 Awareness and Attention

As we discussed earlier in this chapter, following the C-HIP and HILP models, our first big hurdle is getting the user's attention. A just-in-time privacy notice, a privacy-related warning, or an ambient privacy indicator are all useless to users who don't notice them and don't shift their attention to them even briefly. User studies can test whether users notice privacy-related notices, indicators or controls as they carry out a primary task. For example, Balebako et al. developed a mobile game app and created versions that differed only in when the privacy notice was presented to users. They asked study participants to test the app and answer questions about it, including questions about the privacy notice. The timing of the privacy notice presentation impacted whether users paid attention to and thus could correctly answer questions about the privacy notice.[95]

5.5.2.3 Discoverability

Users who want to obtain information about data practices or make privacy choices need to be able to find the relevant interfaces. User studies can test the discoverability of privacy interfaces by evaluating whether users are able to find them at all, and if so, whether they can do so easily and efficiently.

5.5.2.4 Comprehension

Users may notice privacy indicators or find privacy interfaces, but they may not fully understand what they mean or how to use them. User studies can test comprehension of wording, symbols, interface components and controls, either in isolation or in the context of relevant tasks. Comprehension testing is often iterative, as designers identify comprehension problems and then test alternative wording or symbols.

Leon et al. conducted a study to determine what users understood about the AdChoices behavioral advertising icon and associated taglines. They found that comprehension rates were fairly low overall, but that some of the taglines they tested resulted in significantly better comprehension of the icon.[96]

5.5.2.5 Utility

Privacy notices may be discoverable and comprehensible, but they may still leave out important information or settings that users want. For a privacy notice or settings interface to provide utility, it must be aligned with user needs. User studies can help assess the utility of privacy interfaces. Exploratory user studies may also help identify desired utility early in the design process. For example, Emami-Naeni et al. asked IoT device owners about their device purchase experience as well as their privacy and security concerns, and showed them early prototypes of privacy and security labels for IoT devices in order to explore whether the labels' information would address their needs. They found that participants were especially interested in information about their privacy choices and wanted to see a summary rating as well as detailed privacy and security information.[97]

5.5.2.6 Behavior

Users who discover and comprehend privacy indicators or warnings may still fail to act for a variety of reasons. Users may lack motivation, feel that their actions will be ineffective, have concerns that privacy protective steps will be inconvenient or result in their inability to use desired features or services, or be uncertain about what to do. Even those who do act may still make mistakes (e.g., press the wrong button or forget to press a final "submit" button), rendering their action ineffective. Likewise, users presented

with privacy controls may also make mistakes as they attempt to configure the controls to match their personal preferences.

Ed Hutchins, Jim Hollan and Don Norman identified two common underlying challenges that lead to configuration mistakes: understanding the state of the system (*gulf of evaluation*) and taking action to accomplish a specific goal (*gulf of execution*).[98] When users fail to understand the current state of the system, their decisions about what settings to change are based on faulty assumptions. An example of the gulf of evaluation may be a user examining a mobile app privacy setting and incorrectly interpreting the on/off toggle, thus believing that location tracking has been turned off when in fact it is on. Likewise, when users fail to take the correct actions, they will not achieve their desired goals. Thus, a user who clicks an opt-out button on a website but does not scroll to the bottom of the page and click the submit button will fail to achieve the desired goal of opting out due to the gulf of execution.

User studies can test both whether users take privacy actions and whether they execute these actions correctly and effectively to achieve their desired goals. For example, Egelman et al. tested several approaches to providing privacy indicators on websites to determine which had the largest impact on user behavior. They found that privacy indicators had the largest impact on user behavior when they appeared alongside search results rather than on an interstitial or at the top of a web page.[99] In addition, when Leon et al. tested whether users could set up behavioral advertising opt-out tools to achieve a desired result, they found that most users struggled to configure these tools. They recommended a number of design improvements to address the problems they observed, including removing jargon, adding more privacy-protective defaults and providing feedback about what the tool was doing.[100]

5.5.3 How to Test?

UX professionals employ many different testing techniques that vary considerably in their cost, time and effort to complete, and types of insights and conclusions they produce. The most appropriate testing may depend on where you are in the design process, available time and resources, and the reason for testing. (For example, is testing being done primarily to improve an interface, to validate that an interface improves on a prior version or a similar product, to address regulatory requirements or to produce generalizable knowledge that can be published in academic research papers?)

5.5.3.1 Formative and Summative Evaluations

Usability evaluation can be categorized as either *formative* or *summative*.[101] Formative evaluations are used to gain insights into which aspects of a prototype or product could use improvements, while summative evaluations are used to draw comparisons between a prototype or product and some benchmark (e.g., previous version, competing product). Generally, formative evaluations are small scale and focus on gathering rich qualitative insights that can be used to improve a product. They may be conducted iteratively, with several rounds of evaluation and design changes. Summative evaluations are generally conducted once the design team believes they are done or even after the product has shipped. They are generally larger studies that focus on gathering quantitative data. Summative studies are usually done to validate that the product (hopefully) meets the design requirements or performs better than a previous version or a competing product. Companies may make claims about their product on the basis of these results. For example, a company may be able to show that, as compared with the prior version of their product, in the current version they made it 50 percent faster for users to opt out of behavioral targeting or that they improved the percentage of users who could correctly identify the audience of their social media posts from 40 percent to 80 percent.

5.5.3.2 Small-Scale User Testing

Small-scale testing often provides a large bang for your buck, especially when used as part of a formative evaluation. Relatively inexpensive and quick, a test with even a handful of users can offer important insights about where a UX may need improvement. Such testing may uncover user concerns and confusing interfaces and help to highlight differences in user attitudes and uses of a set of features or products. Small-scale user testing is often done in a lab, where an experimenter is able to directly observe how a user interacts with an interface and ask follow-up questions. Sometimes the experimenter may ask users to think aloud during a test and provide a real time, blow-by-blow commentary of what they are thinking and doing.[102] This can be useful for capturing users' impressions as they try to understand and use a new interface. A related method is participatory design (or codesign), in which UX designers work together with affected stakeholders to identify and develop design solutions that better meet users' needs.[103] Rather than only eliciting needs from users and then evaluating new designs with them, users actively contribute to the design.

Small-scale testing can be quite useful for exploratory studies and for iterative design improvements. However, it generally cannot be used to validate a property of an interface or to make statistical claims.

5.5.3.3 Online Panels and Crowdsourcing Services

Recruiting large numbers of participants for a user study can be time consuming and expensive. Research companies recruit and maintain long-standing online panels of people who agree to receive emails recruiting them to participate in online surveys. Often these panels are comprised of people who have been screened for particular characteristics (e.g., age range, gender, occupation, geographic location). Researchers can pay to gain access to these panels to recruit participants for their studies. Some organizations maintain their own in-house panels, recruited from among users of a particular product or members of a university community, for example.

Crowdsourcing platforms such as Amazon Mechanical Turk allow people to sign up as workers and perform tasks for pay (generally at rates similar to minimum wage). These platforms can be a quick and inexpensive way to recruit participants for online surveys and other online user studies. If you use these services, you need to be especially careful about how you screen participants for any desired attributes and make sure that they don't just randomly click through your survey in order to get a quick payment. Some crowdsourcing services such as Prolific and CrowdFlower have been shown to deliver more diverse participants and achieve higher-quality survey results than the more popular Mechanical Turk.[104]

While people often think about using panels and crowdsourcing services for surveys, they can also be used for online user testing. If you can make relevant parts of a user interface available through a web interface or an instrumented download, you can ask participants to perform tasks using your interface, collect data about how they use it and ask them questions afterwards. Services such as UserTesting.com specialize in online user testing of interfaces and user experiences.

5.5.3.4 A/B testing

A/B testing refers to tests where some users of a product or service see version A and others see version B. Generally, there are small differences between A and B and some metric to compare the two. For example, A and B may differ based on the placement of buttons, wording of menu items, colors, page layout, graphics or presence of a particular feature. This type of testing is not actually limited to two versions—companies may test many versions in a similar fashion. This testing may be done with users of a deployed service, who generally are not aware that they are part of a test.

In the context of privacy, A/B testing might be used to test the impact of interface changes on the time users spend engaging with privacy settings, or the percentage of users who change privacy settings or opt out.

5.5.4 Usability Testing Best Practices

Regardless of what type of usability test you choose, it is important to keep in mind some best practices. Here, we discuss testing metrics related to privacy, ways to provide a realistic context (sometimes using deception) and strategies for recruiting study participants. We conclude with a brief discussion of ethical considerations for usability testing and user studies.

5.5.4.1 Metrics

Before testing begins, it is important to identify a set of metrics that will be used. Metrics might relate to any of the areas discussed in section 5.2, including discoverability, awareness, comprehension, utility, behavior, satisfaction and other factors. For example, when testing a privacy tool or interface, metrics might include speed with which users complete a task, number of errors made while completing a task and accuracy of users when answering questions about the meaning of symbols or other interface components. Other relevant metrics include the percentage of users who exercise privacy choices and the types of choices they make. Choosing the right metrics is important for conducting a meaningful evaluation free of blind spots: You don't want to end up in a situation where you conclude that an interface is wonderful because it performs well on a particular metric you tested, while in reality it would perform horribly on several other highly relevant metrics that you didn't evaluate.

Besides identifying metrics, you should identify desired outcomes or objectives that will allow you to determine whether design requirements have been met. For example, one objective may be to increase speed by a specific percentage over the prior version. Another objective may be to reduce the error rate.

Measuring the effectiveness of consent or opt-out interfaces poses challenges. The percentage of users who consent or opt out is not an indicator of the effectiveness of the interface. What we really want to capture is user awareness and comprehension of choices and whether they successfully select options that align with their personal preferences. One way to measure this would be to survey users shortly after they engage with a consent interface to find out what they believe they consented to (or chose not to consent to) and how this aligns with their preferences.

Metrics regarding user privacy behavior may be difficult to agree upon, especially when they conflict with other values in an organization. For example, a privacy objective may be to increase the number of users who visit the privacy settings page and make changes to their settings. However, when users choose more privacy-protective settings, companies may lose opportunities to profit from user data. Thus, meeting privacy objectives may appear to conflict with meeting revenue objectives. On the other hand,

users who select privacy-protective settings might be less likely to trust the company or may decide to take their business elsewhere if such settings are not available. Companies need to consider these issues and resolve internal tensions between privacy and business goals.

5.5.4.2 Ecological Validity, Context and Deception

Ecological validity refers to the realism of the methods, materials and setting of a user study or usability test. The most reliable results are obtained through field testing. In field tests, designs are evaluated as part of a production system. The advantage of field testing is that new designs are tested under real-world conditions with real users rather than under artificial lab conditions with recruited participants. However, this also limits them to late stages of the design process, when a design has sufficiently matured to be used in production systems without placing users at risk or creating liabilities for the organization.

In order to observe the most natural behavior from non-field-study participants, studies should be designed to mimic the real world as much as possible. Thus, when studies are conducted in a lab, the lab may be set up to look like an office or living room, complete with relevant props. For example, when evaluating privacy settings related to smart home devices, conducting the study in a homey living room setting may elicit from participants the types of privacy concerns and behaviors they would exhibit when using these devices in their own homes.

One component of ecological validity is the context in which user study tasks are embedded. If a study participant sitting in a usability lab is shown a privacy policy or any other privacy interface and asked questions about it without being provided with any context or reason for wanting to read the policy, the resulting usability evaluation will lack ecological validity. In this case, participants may pay more attention to the policy than they would in real life, when privacy is likely not their primary concern. Thus, comprehension testing without a real-world context is likely to result in a best-case scenario. Awareness and behavior testing are likely to be difficult to conduct meaningfully without embedding tasks in context.

To evaluate a privacy interface in context, participants should be given a relevant task. For example, participants might be told that a friend who is concerned about their privacy has asked for help in selecting an email service provider. Kelley et al. evaluated prototype "privacy facts" labels in an app store by asking participants to select apps for a friend. They compared the selections made by participants who either used an app store that had privacy facts labels or did not (as a control condition). They found that the privacy facts labels generally influenced which apps were selected. However, other factors also played a role, and sometimes had more of an influence than privacy.[105]

Often context is provided by embedding tasks in a hypothetical scenario such as those mentioned above. However, participants may not be motivated to behave realistically in a hypothetical scenario, especially one in which they have to imagine a particular privacy or security risk. Sometimes study designs employ *deception* to increase ecological validity and observe how users behave when simulated privacy or security risk is introduced. For example, Bravo-Lillo et al. advertised a study evaluating online video games to crowd workers. Participants played and evaluated two games on real gaming websites not controlled by the experimenters. When they tried to play a third game, a simulated security warning appeared in their browser window. The researchers collected data on how participants responded to the warning, and compared results across several warning designs.[106] In this case, there was deception about the purpose of the study and deception about the warnings, which were designed to appear real even though they were triggered by the experimenters. Deception was needed here to observe participants' natural behavior when confronted with warnings that they believed were real and not a part of the experiment. Section 5.5.4.4 addresses ethical issues to consider when conducting deceptive studies.

5.5.4.3 Testing with Representative and Specific User Populations

The strategy used to recruit study participants can have a large impact on results, especially when conducting small-scale studies. The age, gender, socioeconomic status, digital literacy, technical skills, geographic location, cultural background and other characteristics of study participants influence their behavior and responses to survey questions.

When conducting a user study, it is tempting to reach out to a convenient sample of coworkers, friends or family. This can be useful for pilot testing or to help detect problems with a product or service, and it is better than not doing any user testing. However, for the most useful evaluation, testing should be conducted with a representative sample of current or anticipated users of a product or service, or with specific populations that are of interest. Sometimes, it is helpful to recruit participants who represent particular types of users who may have special issues or represent extreme ends of a user spectrum. For example, aware that older adults are now flocking to social media services, a social media company may conduct a study to assess the privacy needs of adults over age 60 and how they use the service's privacy settings. Likewise, a company developing a privacy tool for extremely privacy-conscious users should test the tool with those types of users.

When recruiting participants for privacy-related user studies, it is advisable to not mention *privacy* in the recruitment materials. Privacy is a contentious topic that can lead to self-selection bias: People who already have an interest or opinion regarding privacy

will respond, but those who prefer to ignore the topic won't. However, the latter group can often provide more useful insights. Good practice is to be vague about the specific focus of the study in recruitment (e.g. by calling it a "study on mobile app use").

Researchers sometimes conduct studies where they aim to recruit a random, representative, census-balanced sample of a country's population. This is important if you want to make claims about the entire population of a country, such as "50 percent of all Americans believe … " However, these studies tend to be expensive to conduct, and this type of sampling is generally not necessary for usability tests.

5.5.4.4 Ethical Considerations

Depending on legal requirements in your country, corporate or institutional policies, and the type of testing you are doing, you may or may not be required to have your study reviewed by an ethics board or institutional review board (IRB) prior to beginning a usability study. In the United States, research conducted at universities that involves identifiable individuals or their information and is designed to contribute to generalizable knowledge must be reviewed by an IRB. Research intended for publication is usually classified as contributing to generalizable knowledge, while research intended to guide development of a specific product or that is being conducted as part of a classroom exercise is usually not.

Even if not required to have a usability study reviewed by an ethics board or IRB, the user study team should still consider ethical issues in their study design. Studies should be designed to minimize the possibility of harm to participants and to treat participants fairly and with respect. If identifiable data is collected only to contact or pay participants, it should be separated from study data. Participants should be able to end their participation at any time and for any reason without penalty. The study team should generally obtain informed consent from participants prior to collecting their data. There are situations where this is not feasible because of concerns about priming participants or because the study uses previously collected data. In cases where consent is not obtained in advance or where deception is used, the study team should brief participants after the conclusion of the study.

While most usability studies related to privacy pose minimal risk to participants, there are some study techniques that have the potential to cause harm and should be considered carefully. For example, a privacy study that collects sensitive data from participants, perhaps to motivate privacy concerns, could increase the risk that participants' information will be exposed. A deceptive study that simulates security or privacy warnings may leave participants uncertain about whether they may have actually been the victim of an attack or breach. These participants may suffer from added anxiety and may unnecessarily take steps to reset passwords, cancel credit cards

or otherwise recover from an attack that did not actually occur. Study teams should take precautions to minimize these and other harms; for example, by establishing data security protocols and debriefing participants at the end of a study session.

5.6 Summary

This chapter provided an overview of people's privacy decision-making processes and common usability issues in privacy interfaces. We presented design principles and a systematic process for creating privacy interfaces that can help avoid usability issues. We further discussed best practices for conducting usability testing and user studies related to privacy.

The design of usable privacy notices and controls is not trivial, but this chapter hopefully explained why it is important to invest the effort in getting the privacy user experience right—making sure that privacy information and controls are not only compliant with regulations but also address and align with users' needs. Careful design of the privacy user experience can support users in developing an accurate and more complete understanding of a system and its data practices. Well-designed and user-tested privacy interfaces provide responsible privacy professionals and technologists with the confidence that an indication of consent was indeed an informed and freely given expression by the user. Highlighting unexpected data practices and considering secondary and incidental users reduces surprise for users and hopefully prevents privacy harms, social media outcries, bad press and fines from regulators. Importantly, a privacy interface is not just a compliance tool but rather an opportunity to engage with users about privacy, to explain the rationale behind practices that may seem invasive without proper context, to make users aware of potential privacy risks and to communicate the measures and effort taken to mitigate those risks and protect users' privacy.

Endnotes

1 Article 29 Data Protection Working Party (WP29), Guidelines on Transparency under regulation 2016/679, adopted on 29 November 2017, as last revised and adopted on 11 April 2018, https://ec.europa.eu/newsroom/article29/item-detail.cfm?item_id=622227.

2 See for example Kirsten Martin, "The penalty for privacy violations: How privacy violations impact trust online," *Journal of Business Research*, vol. 82, (2018), 103–116, https://doi.org/10.1016/j.jbusres.2017.08.034. *See also* Batya Friedman, Peter H. Khan Jr and Daniel C. Howe, "Trust online." Communications of the ACM 43, no. 12 (2000): 34–40, https://doi.org/10.1145/355112.355120.

3 In 2019, the Federal Trade Commission (FTC) settled with Facebook on a $5 billion fine and substantial changes to Facebook's privacy management and compliance structure, https://www.ftc.gov/news-events/blogs/business-blog/2019/07/ftcs-5-billion-facebook-settlement-record-breaking-history.

4 Hana Habib, Yixin Zou, Aditi Jannu, Neha Sridhar, Chelse Swoopes, Alessandro Acquisti, Lorrie Faith Cranor, Norman Sadeh and Florian Schaub, "An Empirical Analysis of Data Deletion and Opt-Out Choices on 150 Websites," *Fifteenth Symposium on Usable Privacy and Security*, (2019), https://www.usenix.org/conference/soups2019/presentation/habib.

5 Y. Wang, S. Komanduri, P.G. Leon, G. Norcie, A. Acquisti and L.F. Cranor, "I regretted the minute I pressed share: A Qualitative Study of Regrets on Facebook," *Seventh Symposium on Usable Privacy and Security* (2011).

6 Y. Wang, P. Leon, A. Acquisti, L.F. Cranor, A. Forget and N. Sadeh. "A Field Trial of Privacy Nudges for Facebook," *ACM SIGCHI Conference on Human Factors in Computing Systems* (2014).

7 Florian Schaub, Rebecca Balebako, Adam L. Durity and Lorrie Faith Cranor," A design space for effective privacy notices," *Proceedings of the Eleventh USENIX Conference on Usable Privacy and Security* (2015).

8 Katie Rogers, "Gmail Formally Adds 'Undo Send' Option," *The New York Times*, June 23, 2015, https://www.nytimes.com/2015/06/24/us/gmails-undo-send-option-can-help-end-email-embarrassment.html.

9 Organization for Economic Co-Operation and Development (OECD), The OECD Privacy Framework, 2013, https://www.oecd.org/internet/ieconomy/privacy-guidelines.htm; FTC, *Privacy online: a report to Congress*, 1998, https://www.ftc.gov/sites/default/files/documents/reports/privacy-online-report-congress/priv-23a.pdf; European Parliament and Council, Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC [General Data Protection Regulation (GDPR)], 2016, http://data.europa.eu/eli/reg/2016/679/2016-05-04.

10 Tamara Dinev and Paul Hart, "An Extended Privacy Calculus Model for E-Commerce Transactions. Information Systems Research," (2006) 17:1, 61–80, https://doi.org/10.1287/isre.1060.0080.

11 Pew Research Center, "Public Perceptions of Privacy and Security in the Post-Snowden Era," November 2014, https://www.pewinternet.org/2014/11/12/public-privacy-perceptions/.

12 Alessandro Acquisti and Ralph Gross, "Imagined communities: Awareness, information sharing, and privacy on the Facebook," International workshop on privacy enhancing technologies, (Springer: Berlin, Heidelberg, 2006), 36-58, https://doi.org/10.1007/11957454_3.

13 M. Sleeper, J. Cranshaw, P.G. Kelley, B. Ur, A. Acquisti, L.F. Cranor and N. Sadeh, "i read my Twitter the next morning and was astonished," a conversational perspective on Twitter regrets, (2013), https://doi.org/10.1145/2470654.2466448; Y. Wang, "I regretted the minute I pressed share."

14 Ashwini Rao, Florian Schaub, Norman Sadeh, Alessandro Acquisti and Ruogu Kang, "Expecting the unexpected: Understanding mismatched privacy expectations online,"*Twelfth Symposium on Usable Privacy and Security*, (2016), 77–96, https://www.usenix.org/conference/soups2016/technical-sessions/presentation/rao; Hazim Almuhimedi, Florian Schaub, Norman Sadeh, Idris Adjerid, Alessandro Acquisti, Joshua Gluck, Lorrie Faith Cranor and and Yuvraj Agarwal, "Your Location has been Shared 5,398 Times! A Field Study on Mobile App Privacy Nudging," *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems,* (2015), https://doi.org/10.1145/2702123.2702210; Joshua Gluck, Florian Schaub, Amy Friedman, Hana Habib, Norman Sadeh, Lorrie Faith Cranor and Yuvraj Agarwal, "How short is too short? implications of length and framing on the effectiveness of privacy notices," *Twelfth Symposium on Usable Privacy and Security,* (2016), 321–340, https://www.usenix.org/conference/soups2016/technical-sessions/presentation/gluck.

15 Josephine Lau, Benjamin Zimmerman, and Florian Schaub," Alexa, Are You Listening? Privacy Perceptions, Concerns and Privacy-seeking Behaviors with Smart Speakers," *Proceedings ACM Human-Computer Interaction* 2, CSCW, Article 102 (November 2018), https://doi.org/10.1145/3274371; Alessandro Acquisti, Idris Adjerid, Rebecca Balebako, Laura Brandimarte, Lorrie Faith Cranor, Saranga Komanduri, Pedro Giovanni Leon, Norman Sadeh, Florian Schaub, Manya Sleeper, Yang Wang and Shomir Wilson, "Nudges for Privacy and Security: Understanding and Assisting Users' Choices Online," *ACM Computing Surveys* 50, 3, Article 44 (August 2017), https://doi.org/10.1145/3054926; Yixin Zou, Abraham H. Mhaidli, Austin McCall and Florian Schaub, "'I've Got Nothing to Lose': Consumers' Risk Perceptions and Protective Actions after the Equifax Data Breach," *Symposium on Usable Privacy and Security*, USENIX Association, (2018), https://www.usenix.org/conference/soups2018/presentation/zou.

16 P.A. Norberg, D.R. Horne, D.A. Horne, "The Privacy Paradox: Personal Information Disclosure Intentions versus Behaviors," *Journal of Consumer Affairs*, (2007), 41(1):100–126, https://dx.doi.org/10.1111/j.1745-6606.2006.00070.x.

17 M. Sleeper, R. Balebako, S. Das, A.L. McConahy, J. Wiese and L.F. Cranor, "The post that wasn't: exploring self-censorship on Facebook," CSCW (2013), https://doi.org/10.1145/2441776.2441865.

18 *See* Jon Penney, "Chilling Effects: Online Surveillance and Wikipedia Use," *Berkeley Technology Law Journal*, vol. 31, no. 1, (2016), 117, https://ssrn.com/abstract=2769645; Alex Marthews and Catherine E. Tucker, "Government Surveillance and Internet Search Behavior," (February 17, 2017), https://ssrn.com/abstract=2412564.

19 Alessandro Acquisti, Laura Brandimarte and George Loewenstein, "Privacy and human behavior in the age of information," *Science,* 347, no. 6221, (2015), 509–514.

20 Joseph Turow, Amy Bleakley, John Bracken, Michael X. Delli Carpini, Nora Draper, Lauren Feldman, Nathaniel Good, Jens Grossklags, Michael Hennessy, Chris Jay Hoofnagle, Rowan Howard-Williams, Jennifer King, Su Li, Kimberly Meltzer, Deirdre Mulligan and Lilach Nir, *Americans, Marketers, and the Internet: 1999-2012,* Annenberg School for Communication, University of Pennsylvania, (April 11, 2014), http://dx.doi.org/10.2139/ssrn.2423753.

21 L. Brandimarte, A. Acquisti and G. Loewenstein, "Misplaced Confidences: Privacy and the Control Paradox," *Social Psychological and Personality Science*, 4(3), (2013), 340–347, https://doi.org/10.1177/1948550612455931.

22 H.A. Simon, *Models of Man, Social and Rational: Mathematical Essays on Rational Human Behavior in a Social Setting*, (New York: Wiley, 1957).

23 Kahneman, Daniel, *Thinking, fast and slow,* (Macmillan, 2011).

24 Acquisti, "Nudges for Privacy and Security: Understanding and Assisting Users' Choices Online"; Acquisti, "Privacy and human behavior in the age of information"; Alessandro Acquisti and Jens Grossklags, "Privacy and rationality in individual decision-making," *IEEE Security & Privacy*, vol. 3, no. 1, (January-February 2005), 26–33, https://doi.org/10.1109/MSP.2005.22.

25 For a comprehensive discussion of privacy-related decision heuristics and biases, *see* Alessandro Acquisti, Idris Adjerid, Rebecca Balebako, Laura Brandimarte, Lorrie Faith Cranor, Saranga Komanduri, Pedro Giovanni Leon, Norman Sadeh, Florian Schaub, Manya Sleeper, Yang Wang and Shomir Wilson, "Nudges for Privacy and Security: Understanding and Assisting Users' Choices Online," *ACM Comput. Surv.* 50, 3, Article 44 (August 2017), https://doi.org/10.1145/3054926.

26 Daphne Chang, Erin L. Krupka, Eytan Adar and Alessandro Acquisti, "Engineering Information Disclosure: Norm Shaping Designs," P*roceedings of the 2016 CHI Conference on Human Factors in Computing Systems,* 587–597, (New York: ACM, 2016), https://doi.org/10.1145/2858036.2858346.

27 Alessandro Acquisti, Leslie K. John and George Loewenstein, "The impact of relative standards on the propensity to disclose," *J. Market. Res.* 49, 2 , (2012), 160–174.

28 Alessandro Acquisti, Leslie K. John, and George Loewenstein, "What is privacy worth?" *The Journal of Legal Studies* 42, 2 (2013), 249–274.

29 N. Jentzsch, S. Preibusch and A. Harasser, *Study on monetising privacy: An economic model for pricing personal information*, ENISA report, (2012), https://www.enisa.europa.eu/publications/monetising-privacy.

30 Yixin Zou, Abraham H. Mhaidli, Austin McCall and Florian Schaub, "I've Got Nothing to Lose: Consumers' Risk Perceptions and Protective Actions after the Equifax Data Breach," *Symposium on Usable Privacy and Security,* (2018), https://www.usenix.org/conference/soups2018/presentation/zou.

31 Acquisti, "Nudges for Privacy and Security."

32 Acquisti, "Privacy and human behavior in the age of information."

33 G.T. Murky Marx, "Conceptual waters: The public and the private," *Ethics and Information Technology* 3: (2001), 157, https://doi.org/10.1023/A:1012456832336.

34 A.E. Marwick and danah boyd, "I tweet honestly, I tweet passionately: Twitter users, context collapse, and the imagined audience," *New Media & Society,* 13(1), (2011) 114–133, https://doi.org/10.1177/1461444810365313; Jenny L Davis and Nathan Jurgenson , "Context collapse: theorizing context collusions and collisions," *Information, Communication & Society,* 17:4 (2014), 476–485.

35 Helen Nissenbaum, *Privacy in Context: Technology, Policy, and the Integrity of Social Life,* (Stanford University Press, 2009).

36 I. Altman, "The Environment and Social Behavior: Privacy, Personal Space, Territory, Crowding," 1975.

37 Leysia Palen and Paul Dourish, "Unpacking 'privacy' for a networked world," *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems,* (ACM, 2003), 129–136, http://dx.doi.org/10.1145/642611.642635.

38 Sandra Petronio, "Communication Privacy Management Theory," *The International Encyclopedia of Interpersonal Communication,*(eds. C. R. Berger, M. E. Roloff, S. R. Wilson, J. P. Dillard, J. Caughlin and D. Solomon), https://onlinelibrary.wiley.com/doi/abs/10.1002/9781118540190.wbeic132.

39 This is a subset of the privacy dark patterns described in Forbrunder Rådet (Norwegian Consumer Council), *Deceived by Design: How tech companies use dark patterns to discourage us from exercising our rights to privacy,* (June 2018), https://www.forbrukerradet.no/undersokelse/no-undersokelsekategori/deceived-by-design/; Acquisti, "Nudges for Privacy and Security"; Christoph Bösch, Benjamin Erb, Frank Kargl, Henning Kopp and Stefan Pfattheicher, "Tales from the dark side: Privacy dark strategies and privacy dark patterns," *Proceedings on Privacy Enhancing Technologies,* no. 4 (2016), 237-254, https://doi.org/10.1515/popets-2016-0038.

40 Sonam Samat and Alessandro Acquisti, "Format vs. content: the impact of risk and presentation on disclosure decisions," *Thirteenth Symposium on Usable Privacy and Security,* (2017), 377–384.

41 *See* example of the LinkedIn profile strength meter in Acquisti, "Nudges for Privacy and Security."

42 Chang, "Engineering Information Disclosure: Norm Shaping Designs."

43 This effect has been shown in multiple experiments, see for example I. Adjerid, A. Acquisti, L. Brandimarte and G. Lowenstein, "Sleights of Privacy: Framing, Disclosure, and the Limits of Transparency," https://doi.org/10.1145/2501604.2501613; S. Patil, R. Hoyle, R. Schlegel, A. Kapadia and A. J. Lee, "Interrupt now or inform later? Comparing immediate and delayed privacy feedback," (2015), https://doi.org/10.1145/2702123.2702165.

44 *See* International Standards Organization, ISO 9241-11: Ergonomics of human-system interaction—Part 11: Usability: Definitions and concepts, (2018), https://www.iso.org/obp/ui/#iso:std:iso:9241:-11:ed-2:v1:en.

45 Quoted from Jakob Nielsen, Usability 101: Introduction to Usability, Nielsen Norman Group, January 2012. https://www.nngroup.com/articles/usability-101-introduction-to-usability/.

46 See Jakob Nielsen, Usability 101.

47 Don Norman and Jakob Nielsen, The Definition of User Experience (UX), Nielsen Norman Group, https://www.nngroup.com/articles/definition-user-experience/.

48 The following book provides a good overview of common design methods and tools: Bruce Hanington and Bella Martin, *Universal Methods of Design: 100 Ways to Research Complex Problems, Develop Innovative Ideas, and Design Effective Solution*, (Rockport Publishers, 2012).

49 *See*, for example, David Wright, "Making Privacy Impact Assessment More Effective," The Information Society, 29:5, (2013), 307–315, "https://doi.org/10.1080/01972243.2013.825687.

50 Initial efforts to compile libraries of privacy design patterns include https://privacypatterns.org/ and https://privacypatterns.eu/.

51 Batya Friedman, Peter H. Kahn, Jr. and Alan Boring, *Value sensitive design and information systems,* The handbook of information and computer ethics, 69–101, (2008), https://onlinelibrary.wiley.com/doi/abs/10.1002/9780470281819.ch4.

52 Batya Friedman, David G. Hendry and Alan Borning, "A Survey of Value Sensitive Design Methods," *Foundations and Trends in Human–Computer Interaction*, vol. 11, no. 2, (2017), 63–125, http://dx.doi.org/10.1561/1100000015.

53 Friedman, *Value sensitive design and information systems*.

54 *See* information life cycle as described in P. Swire and K. Ahmad, *Foundations of Information Privacy and Data Protection,* (Portsmouth: IAPP, 2012).

55 A. M. McDonald and L. F. Cranor, "The cost of reading privacy policies," *I/S: A Journal of Law and Policy for the Information Society,* 4(3): (2008), 540–565.

56 P. M. Schwartz and D. Solove, "Notice and choice," *The Second NPLAN/BMSG Meeting on Digital Media and Marketing to Children*, (2009).

57 This definition is based on the GDPR's definition of consent (GDPR, Art. 7), which reflects the established understanding of what constitutes valid consent.

58 F. Schaub, R. Balebako and L. F. Cranor, "Designing Effective Privacy Notices and Controls" *IEEE Internet Computing*, vol. 21, no. 3, (May-June 2017), 70–77, https://doi.org/10.1109/MIC.2017.75.

59 L. Cranor, "Giving notice: Why privacy policies and security breach notifications aren't enough," *IEEE Communications Magazine,* 43(8) (August 2005):18–19.

60 F. Cate, "The limits of notice and choice," *IEEE Security & Privacy,* 8(2) (March 2010):59–62.

61 Schwartz , "Notice and choice"; McDonald, "The cost of reading privacy policies"; J. R. Reidenberg, T. Breaux, L. F. Cranor, B. French, A. Grannis, J. T. Graves, F. Liu, A. M. McDonald, T. B. Norton, R. Ramanath, N. C. Russell, N. Sadeh and F. Schaub, "Disagreeable privacy policies: Mismatches between meaning and users' understanding," *Berkeley Technology Law Journal*, 30(1): (2015), 39–88.

62 N.A. Draper and J. Turow, "The corporate cultivation of digital resignation," *New Media & Society,* 21(8), (2019), 1824–1839, https://doi.org/10.1177/1461444819833331.

63 N. S. Good, J. Grossklags, D. K. Mulligan and J. A. Konstan, "Noticing notice: A large-scale experiment on the timing of software license agreements," *Proceedings of CHI '07*, 607–616, (New York: ACM, 2007); B. Anderson, A. Vance, B. Kirwan, D. Eargle and S. Howard, "Users aren't (necessarily) lazy: Using NeuroIS to explain habituation to security warnings," *Proceedings of ICIS '14*, (2014).

64 The ability of users to find key user interface components is referred to as *discoverability* by user experience designers. Lawyers use this term to mean something quite different: the fact that one side in a legal case must provide certain documents or information to the other side.

65 Habib, "An Empirical Analysis of Data Deletion and Opt-Out Choices on 150 Websites."

66 See, for example, Rebecca Balebako, Richard Shay and Lorrie Faith Cranor, "Is your inseam a biometric? A case study on the role of usability studies in developing public policy," *Workshop on Usable Security,* USEC, (2014); Pedro Giovanni Leon, Justin Cranshaw, Lorrie Faith Cranor, Jim Graves, Manoj Hastak, Blase Ur and Guzi Xu, "What do online behavioral advertising privacy disclosures communicate to users?" *Proceedings of the 2012 ACM workshop on Privacy in the electronic society,* 19–30, (New York: ACM, 2012), http://dx.doi.org/10.1145/2381966.2381970; Pedro Leon, Blase Ur, Richard Shay, Yang Wang, Rebecca Balebako and Lorrie Cranor. "Why Johnny can't opt out: a usability evaluation of tools to limit online behavioral advertising," *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 589-598, (New York: ACM, 2012), https://doi.org/10.1145/2207676.2207759; Forbrunder Rådet (Norwegian Consumer Council), *Deceived by Design: How tech companies use dark patterns to discourage us from exercising our rights to privacy,* report, (June 2018), https://www.forbrukerradet.no/undersokelse/no-undersokelsekategori/deceived-by-design/.

67 The here-described principles are an expanded description of Schaub, "Designing Effective Privacy Notices and Controls."

68 M. S. Wogalter, ed., "Communication-Human Information Processing (C-HIP) Model," *Handbook of Warnings*, 51–61, (Mahwah, NJ: Lawrence Erlbaum Associates, 2006).

69 L. Cranor, "A Framework for Reasoning About the Human In the Loop," *Usability, Psychology and Security 2008*, http://www.usenix.org/events/upsec08/tech/full_papers/cranor/cranor.pdf.

70 Marian Harbach, Markus Hettig, Susanne Weber and Matthew Smith, "Using personal examples to improve risk communication for security & privacy decisions," P*roceedings of the SIGCHI Conference on Human Factors in Computing Systems,* (New York: ACM, 2014), http://dx.doi.org/10.1145/2556288.2556978.

71 Quoted from GDPR Recital 58, which expands on the child-specific transparency requirements described in GDPR Article 12.1.

72 Kevin Litman-Navarro, "We Read 150 Privacy Policies. They Were an Incomprehensible Disaster," *The New York Times,* June 2019, https://www.nytimes.com/interactive/2019/06/12/opinion/facebook-google-privacy-policies.html.

73 Martin Degeling, Christine Utz, Christopher Lentzsch, Henry Hosseini, Florian Schaub and Thorsten Holz, "We Value Your Privacy ... Now Take Some Cookies: Measuring the GDPR's Impact on Web Privacy," *Network and Distributed Systems Symposium,* (2019), https://www.ndss-symposium.org/wp-content/uploads/2019/02/ndss2019_04B-2_Degeling_paper.pdf.

74 Sameer Badarudeen and Sanjeev Sabharwal, "Assessing readability of patient education materials: current role in orthopaedics," *Clinical Orthopaedics and Related Research* 468, no. 10 (2010): 2572–2580, https://dx.doi.org/10.1007%2Fs11999-010-1380-y.

75 Schaub, "Designing Effective Privacy Notices and Controls."

76 Lau, "Alexa, Are You Listening?"

77 Parts of the process description originate from two prior publications of the authors, which describe the process in more detail in Schaub, "Designing Effective Privacy Notices and Controls"; Schaub, "A design space for effective privacy notices."

78 B. Ur, J. Jung and S. Schechter. "Intruders versus intrusiveness: Teens' and parents' perspectives on home-entryway surveillance," *Proceedings of UbiComp '14*, 129–139, (New York: ACM, 2014).

79 Lau, "Alexa, Are You Listening?"

80 Gluck, "How short is too short?"

81 Microsoft, Privacy Guidelines for Developing Software Products and Services, Technical Report version 3.1, 2008.

82 B. Shneiderman, "The eyes have it: a task by data type taxonomy for information visualizations," *Proceedings 1996 IEEE Symposium on Visual Languages*, (Boulder, CO, 1996), 336–343, https://doi.org/10.1109/VL.1996.545307.

83 A. M. Mcdonald, R. W. Reeder, P. G. Kelley and L. F. Cranor, "A comparative study of online privacy policies and formats," *proceedings of PETS '09*, 37–55, (Berlin: Springer, 2009); H. Nissenbaum, "A Contextual Approach to Privacy Online," *Daedalus*, vol. 140, no. 4, (2011), 32–48.

84 Almuhimedi, "Your Location has been Shared 5,398 Times!"

85 Habib, "An Empirical Analysis of Data Deletion and Opt-Out Choices on 150 Websites."

86 Marc Langheinrich and Florian Schaub, "Privacy in Mobile and Pervasive Computing," *Synthesis Lectures on Mobile and Pervasive Computing 10*, no. 1 (2018): 1–139, https://doi.org/10.2200/S00882ED1V01Y201810MPC013.

87 See the original paper for an expanded discussion of each dimension with extensive examples and pointers to relevant academic research: Schaub, "A design space for effective privacy notices."

88 F. Schaub, B. Könings and M. Weber, "Context-Adaptive Privacy: Leveraging Context Awareness to Support Privacy Decision-making," *IEEE Pervasive Computing*, vol. 14, no. 1, (January-March 2015), 34–43, https://doi.org/10.1109/MPRV.2015.5; Schaub, "Context-Adaptive Privacy Mechanisms," *Handbook of Mobile Data Privacy*, eds. A. Gkoulalas-Divanis and C. Bettini, (Springer, 2018), https://doi.org/10.1007/978-3-319-98161-1_13.

89 R. S. Portnoff, L. N. Lee, S. Egelman, P. Mishra, D. Leung and D. Wagner, "Somebody's watching me? Assessing the effectiveness of webcam indicator lights," *In Proc. CHI '15*, 1649–1658, (New York: ACM, 2015).

90 R. Calo, "Against Notice Skepticism in Privacy (and Elsewhere)," *Notre Dame Law Rev.*, vol. 87, no. 3, (2012) 1027–1072.

91 Lau, "Alexa, Are You Listening?"

92 A standard for machine-readable website privacy policies was published by the World Wide Web Consortium in 2002, but has been retired after failing to receive widespread adoption. See Lorrie Faith Cranor, "Necessary But Not Sufficient: Standardized Mechanisms for Privacy Notice and Choice," *Journal of Telecommunications and High Technology Law*, vol. 10, no. 2, (2012), http://www.jthtl.org/content/articles/V10I2/JTHTLv10i2_Cranor.PDF.

93 Langheinrich, Marc, "A privacy awareness system for ubiquitous computing environments,"" *International conference on Ubiquitous Computing*, 237–245, (Berlin, Heidelberg: Springer, 2002); A. Das, M. Degeling, D. Smullen and N. Sadeh, "Personalized Privacy Assistants for the Internet of Things: Providing Users with Notice and Choice," *IEEE Pervasive Computing*, vol. 17, no. 3, (July-September 2018), 35-46, https://doi.org/10.1109/MPRV.2018.03367733.

94 Maggie Oates, Yama Ahmadullah, Abigail Marsh, Chelse Swoopes, Shikun Zhang, Rebecca Balebako and Lorrie Cranor, "Turtles, Locks, and Bathrooms: Understanding Mental Models of Privacy Through Illustration," *Proceedings on Privacy Enhancing Technologies*, (4): (2018) 5–32; B. Ur, P.G. Leon, L.F. Cranor, R. Shay, and Y. Wang, "Smart, Useful, Scary, Creepy: Perceptions of Online Behavioral Advertising," *Symposium on Usable Privacy and Security,* (2012); Cynthia E Schairer, Cynthia Cheung, Caryn Kseniya Rubanovich, Mildred Cho, Lorrie Faith Cranor and Cinnamon S Bloss, "Disposition toward privacy and information disclosure in the context of emerging health technologies," *Journal of the American Medical Informatics Association*, vol. 26, issue 7, (July 2019), 610–619, https://doi.org/10.1093/jamia/ocz010.; Pardis Emami Naeini, Sruti Bhagavatula, Hana Habib, Martin Degeling, Lujo Bauer, Lorrie Cranor and Norman Sadeh, "Privacy Expectations and Preferences in an IoT World," *Symposium on Usable Privacy and Security*, Santa Clara, CA, (July 12-14, 2017).

95 R. Balebako, F, Schaub, I. Adjerid, A. Acquisti and L. Cranor, "The Impact of Timing on the Salience of Smartphone App Privacy Notices," *5th Annual ACM CCS Workshop on Security and Privacy in Smartphones and Mobile Devices,* (2015).

96 Pedro Giovanni Leon, Justin Cranshaw, Lorrie Faith Cranor, Jim Graves, Manoj Hastak, Blase Ur and Guzi Xu, "What do online behavioral advertising privacy disclosures communicate to users?" *Proceedings of the 2012 ACM workshop on Privacy in the electronic society,* 19–30, (New York: ACM, 2012), DOI=http://dx.doi.org/10.1145/2381966.2381970.

97 Pardis Emami-Naeini, Henry Dixon, Yuvraj Agarwal and Lorrie Faith Cranor, "Exploring How Privacy and Security Factor into IoT Device Purchase Behavior," *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems,* (New York: ACM), https://doi.org/10.1145/3290605.3300764.

98 Kathryn Whitenton, "The Two UX Gulfs: Evaluation and Execution," March 11, 2018, https://www.nngroup.com/articles/two-ux-gulfs-evaluation-execution/.

99 S. Egelman, J. Tsai, L. Cranor and A. Acquisti, "Timing Is Everything? The Effects of Timing and Placement of Online Privacy Indicators," *CHI '09: Proceedings of the SIGCHI conference on Human Factors in Computing Systems* (2009).

100 Leon, "Why Johnny Can't Opt Out."

101 Joyce, "Formative vs. Summative Evaluations," July 28, 2019, https://www.nngroup.com/articles/formative-vs-summative-evaluations/.

102 Jakob Nielsen, "Thinking Aloud: The #1 Usability Tool," January 16, 2012, https://www.nngroup.com/articles/thinking-aloud-the-1-usability-tool/.

103 Douglas Schuler and Aki Namioka, eds., *Participatory design: Principles and practices,* (CRC Press, 1993).

104 Eyal Peer, Laura Brandimarte, Sonam Samat and Alessandro Acquisti, "Beyond the Turk: Alternative platforms for crowdsourcing behavioral research," *Journal of Experimental Social Psychology,* vol. 70, (2017), 153–163, ISSN 0022-1031, https://doi.org/10.1016/j.jesp.2017.01.006.

105 P.G. Kelley, L.F. Cranor and N. Sadeh. "Privacy as Part of the App Decision-Making Process," CHI 2013.

106 C. Bravo-Lillo, L.F. Cranor, J. Downs, S. Komanduri, R.W. Reeder, S. Schechter and M. Sleeper, "Your Attention Please: Designing security-decision UIs to make genuine risks harder to ignore," *Proceedings of the Eighth Symposium On Usable Privacy and Security,* Newcastle, United Kingdom, (2013).

Index

B

C

D

E

F

G

H

I

J

K

L

M

N

O

P

Q

R

S

T

U

V

W

X

Y

Z