Chapter 16

User-Tailored Privacy

**Bart P. Knijnenburg, Reza Ghaiumy Anaraky, Daricia Wilkinson,
Moses Namara, Yangyang He, David Cherry, and Erin Ash**

Abstract Modern information systems require their users to make a myriad of privacy decisions, but users are often neither motivated nor capable of managing this deluge of decisions. This chapter covers the concept of tailoring the privacy of an information system to each individual user. It discusses practical problems that may arise when collecting data to determine a user's privacy preferences, techniques to model these preferences, and a number of adaptation strategies that can be used to tailor the system's privacy practices, settings, or interfaces to the user's modeled preferences. Throughout the chapter, we provide recommendations on how to develop user-tailored privacy solutions, depending on the requirements and characteristics of the system and its users.

16.1 Introduction

As our digital and personal lives become increasingly intertwined, the frequency with which we encounter privacy decisions is on the rise. Moreover, given the complexity of modern information systems, users often report feeling helplessly overwhelmed by the privacy decision-making required to effectively manage the boundaries around the collection and use of their personal information.

In this chapter we present User-Tailored Privacy (UTP) as a means to reduce the burden of privacy decision-making. Combining the best aspects (and avoiding

B. P. Knijnenburg (✉) · R. G. Anaraky · D. Wilkinson · M. Namara · Y. He
School of Computing, Clemson University, Clemson, SC, USA
e-mail: bartk@clemson.edu; rghaium@clemson.edu; dariciw@clemson.edu;
mosesn@clemson.edu; yyhe@clemson.edu

D. Cherry
Computer Science Department, Morehouse College, Atlanta, GA, USA
e-mail: David.cherry@morehouse.edu

E. Ash
Department of Communication, Clemson University, Clemson, SC, USA
e-mail: ash3@clemson.edu

the downsides) of the existing privacy management paradigms of “notice and choice” and “privacy nudging” (see Chap. 2 for an overview of existing paradigms), the concept of UTP can be implemented alongside technical privacy-preserving solutions (see Chap. 8) and Privacy by Design (see Chap. 2) to provide a system whose privacy settings are tailored to the level of privacy each individual user is most comfortable with.

In the remainder of this section we first critically appraise the shortcomings of existing privacy management paradigms. In Sect. 16.2 we outline the UTP framework, which consists of three phases that will be discussed in subsequent sections: Measuring users’ privacy preferences (Sect. 16.3), modeling these preferences and the appropriate decision context (Sect. 16.4), and adapting the system according to these models (Sect. 16.5). We conclude with an overview of the various goals UTP can accomplish if implemented correctly (Sect. 16.6).

16.1.1 The Limitations of Technical Solutions to Privacy

Engineers tend to consider technical solutions to privacy problems, which broadly fall into two categories: Architectures, platforms, and standards designed specifically to minimize data leakage (including distributed architectures, portable user profiles, and client-side personalization techniques that provide limited access to and “linkability” of user data [1, 2]), and algorithmic techniques for data protection (including anonymization, obfuscation, differential privacy, and homomorphic encryption [3, 4]). These solutions have numerous well-documented limitations: Distributed architectures [1] and encryption [3] are notoriously slow, client-side personalization leaves users vulnerable to data loss and theft [2], and providing full anonymity or even pseudonymity is often not feasible in modern information systems [5].

A notorious example of the latter is the case where Netflix released anonymized user data as part of a contest to improve its recommendation algorithm. Soon after the release of the data, researchers were able to de-anonymize the data by cross-referencing anonymous user ratings with IMDb profiles [5]. In response, a closeted lesbian mother sued Netflix, alleging that the de-anonymization procedure could “out” her based on her viewing behavior [6].

Importantly, technical privacy-preserving solutions do not apply to social networking applications, where disclosure is at the heart of the functionality of the application. In such applications, it is not the disclosure of the information per se that users are worried about, but rather the determination of *who* has access to the disclosed information [7].

Moreover, while technical solutions provide some protection against privacy violations, the existence of such protection does in itself not necessarily mean that users will disclose more information. For example, work on client-side personalization shows that users’ perception of the privacy afforded by this technology is modest and has only a very slight impact on their subsequent sharing decisions [2]. In other

cases, the use of the technical solution itself is cumbersome, and many users avoid it (see Chap. 8). Hence technical solutions to privacy must be supplemented with socio-technical solutions.

16.1.2 The Limitations of Privacy by Design

As discussed in Chap. 2, Privacy by Design is a design philosophy in which privacy aspects are addressed early in the system design and development process [8]. Privacy by Design is often regarded as an alternative to outfitting a system with a vast array of privacy settings. Unfortunately, the existence of privacy settings is inevitable in many modern information systems, for two important reasons. First of all, the functionalities embedded in many modern information systems, such as social networking and personalization, are not feasible without data collection [9, 10]. In these cases, privacy and functionality are in direct opposition, and a decision must be made regarding how much data collection is justified to provide a certain level of functionality. Second, users tend to differ extensively regarding this decision. The goal of the system is to ensure that all users' privacy preferences are adhered to without limiting certain users' ability to use the system to the fullest extent [11].

16.1.3 The Limitations of Notice and Choice

Where privacy is in direct opposition with system functionality and users differ in the amount of privacy they prefer to trade off for functionality, privacy experts argue that users must be given controls (e.g., “privacy settings”) as a means to effect this trade-off, as well as a certain amount of information that will help them operate these controls [12]. The idea of “notice and choice” is also at the heart of existing or planned regulatory schemes (see Chap. 18).

Unfortunately, for most modern information systems, privacy notices fall prey to the “transparency paradox” [13]: notices that are sufficiently detailed to have an impact on people's privacy decisions are often too long and complex for people to read. Moreover, notices may actually decrease disclosure, even if they are supposed to indicate positive privacy protection practices [14, 15]. For example, e-commerce practitioners have documented several cases where privacy seals *decrease* conversion rates rather than increasing them [16, 17]. Finally, the effect of notice may be very fleeting: even the slightest distraction can easily nullify any effect of privacy notices [18]. Indeed, while many people claim to read online privacy policies [19], many do not actually read them [20] or do not read closely enough to understand them [21].

Likewise, users of modern information systems tend to fall prey to the “control paradox,” which states that users claim they want full control over their privacy but

often do not actually take control, even when it is offered [22]. Indeed, many users tend to pay little attention to privacy seals [23], social navigation cues [24], privacy assurances [25], and permission requests [26].

Finally, many researchers no longer believe that users always make “calculated” privacy decisions, but often employ heuristic decision strategies instead (see Chap. 4). As such, even if users do take control over their privacy, it is not certain that they will do so effectively.

16.1.4 The Limitations of Privacy Nudging

Privacy nudges attempt to relieve some of the burden of privacy decision-making by making it easier for users to make the “right” decisions regarding their privacy [18, 27, 28]. Traditionally, nudges have been defined as covert changes to the structure, framing, and defaults of a decision environment, but in the field of privacy they have also been used to describe designs that steer users in a desirable direction in a more overt manner [27].

Ample research has demonstrated that the more overt nudges like justifications [24, 29], privacy seals [30], and audience/sentiment feedback [28, 31] fail to have a consistent effect on users’ disclosure and privacy concerns. Moreover, while the traditional nudges have been found to be more effective [32], they have typically only been tested for behavioral impact, disregarding the question of whether they reduced users’ privacy concerns or their privacy decision burden [33]. Indeed, researchers worry that defaults may threaten consumer autonomy, especially when they work outside of users’ explicit awareness [34].

16.1.5 A Case for User-Tailored Privacy

We summarize the problems with existing privacy solutions as follows:

- **Technical solutions cannot always effectively be implemented**, hence they must be complemented with user-centric solutions.
- **Privacy by Design is not universally applicable**; it fails when privacy and functionality are not in conflict, and when users differ in their inherent trade-off between privacy and functionality. In these cases, privacy settings are inevitable.
- **Notice and choice assume that users will take control over their privacy**. In reality, they often fail to take control effectively, either due to a lack of motivation or ability.
- **Privacy nudging takes a one-size-fits-all approach to privacy**, and by making universal assumptions regarding the “best” privacy decisions, they threaten user autonomy.

User-Tailored Privacy (UTP) acknowledges the need for privacy settings, but rather than putting the full burden of managing these settings on the user, it uses *personalized* nudges to simplify and/or automate part of these privacy management responsibilities. As such, it takes into account both the wide variety in users' privacy preferences and their inability to effectively implement these preferences themselves in the context of complex modern information systems.

The concept of tailoring privacy to users' needs was first explicitly discussed by Kobsa in his keynote to the 2001 *User Modeling* conference [35]. Since then, the idea has been explored in several areas of computing.

One existing application of UTP makes user-tailored suggestions for privacy settings or permissions. In this light, Liu et al. [36] developed a profile-based personalized privacy assistant for smartphone app permissions. The app groups users into different profiles based on their privacy preferences. Based on these profiles, the assistant recommends permission settings that the user could change. User study results show that the recommendations were adopted by the majority of users. And that the recommendations led to more restrictive permission settings without compromising on user comfort with these configurations.

UTP can also recommend Web sites or applications to users, based on their adherence to the user's privacy preferences, for example, in the context of an app store. The idea of automated evaluation of the privacy practices of Web sites or applications has a long history that started with the invention and eventual demise of P3P [37]. Several studies have looked into evaluating Web sites or applications from a privacy perspective [38], and some have found that users who have access to such evaluations may end up paying a premium for privacy [39]. While these studies provided users with privacy information about the Web site or app, none of these studies took an active approach to provide privacy-based *recommendations*. UTP could leverage the generated privacy descriptions (cf. [38]) and provide automatic recommendations.

In a social networking setting, UTP has been used to help users decide what information should be shared with whom. For example, Fang and LeFevre use hierarchical clustering on social network structures to predict the most suitable audience for users' personal information [40], Ravichandran et al. [41] demonstrated how a small number of default policies can accurately capture most users' location-sharing preferences, and Knijnenburg and Jin recommend audience-related sharing settings for a location-sharing service [42].

Finally, research has shown that social network users employ privacy management strategies that go beyond selective information sharing (see Chap. 7). Recent work shows that social network users can be classified into six profiles when it comes to these privacy boundary management strategies [43]—a few studies have demonstrated how UTP can leverage these profiles to adapt the privacy-setting interface of the social network to highlight the user's most-used privacy functionalities [44, 45].

Fig. 16.1 Roots of the UTP approach

Beyond these existing applications of UTP, we note that the idea is rooted in decades of user-centered privacy research as it builds upon the strengths of existing privacy paradigms while avoiding the aforementioned limitations (see Fig. 16.1):

- Like other technical solutions, it leverages computational power to help users manage their privacy, thereby shifting some of the burden from the user to the system.
- Like Privacy by Design, it aims to make privacy “effortless.”
- Like notice and choice, it allows users to take control over their privacy.
- Like nudging, it acknowledges (or even leverages) users’ heuristic decision-making practices.

In the following sections, we describe UTP in more detail and give advice regarding its implementation.

Fig. 16.2 Overview of UTP framework

16.2 The UTP Framework

UTP is an approach to privacy that models users' privacy concerns and provides them with adaptive privacy decision support. With UTP, a system measures user privacy-related characteristics and behaviors, uses this as input to model their privacy preferences, and then adapts the system's privacy settings to these preferences (see Fig. 16.2). UTP solutions attempt to provide the right amount of privacy-related information while allowing users to maintain control, without being misleading or overwhelming.

The following three steps are common among the many variations of UTP; in the following sections we describe each of these steps in more detail:

- **Measure user characteristics and behavior** in order to tailor the provided support to the user and the context of the decision.
- **Model users' privacy decisions using machine learning algorithms** and then plan adaptations based on this model. Models can be built based on direct observation of users' behaviors or via inference from their attitudes.
- **Adapt the system to the users' privacy concerns.** UTP can adapt the system's privacy settings, justifications, privacy-setting user interface, and/or personalization procedure.

16.3 Measuring the User

The first step of UTP is to model users' privacy decisions. Users vary extensively in their privacy decisions; no two people make exactly the same privacy decisions, and even for the same person the decision tends to depend on the context.

The dynamics of contextually determined privacy norms are the central topic of Nissenbaum's popular *Contextual Integrity* framework (see [13], covered in more detail in Chap. 2). The key contextual and personal variables that have been shown to influence users' privacy decisions are:

- **The data requested (What)**—Users' privacy behaviors and decisions vary by the type of data that is being collected [46, 47].
- **The user him/herself (Who)**—There exist distinct profiles of privacy behaviors among users [47, 48].
- **The recipient of the information (To whom)**—The recipient of the information plays an important role in users' disclosure decisions as well, both in “commercial” and “social” settings [49–51].
- **Other factors, usually system- or purpose-specific**—In certain types of systems, privacy preferences depend on other contextual factors [52–54].

The system should identify as much of these contextual and personal variables as possible to be able to accurately model users' privacy decision.

16.3.1 *The Data (What)*

Each data category has a different degree of sensitivity. Research shows that the least sensitive types of information are first name, email address, physical characteristics (age, gender, height), and interests and preferences, while the most sensitive types of information are contact information (aside from email address), financial information, information regarding sex and birth control, and social security number [33, 55]. These attitudinal results have been confirmed behaviorally [2, 29, 43].

Beyond these general rankings, research has found that privacy behaviors are multidimensional, meaning that disclosure is more complex than a single tendency, but also not completely unstructured [47, 48]. The differences between the uncovered dimensional structures also show that the underlying dimensionality of disclosure behaviors varies by context.

16.3.2 *The User (Who)*

Users differ in their level of privacy concern and behavior. Most privacy surveys report sizable standard deviations in their estimates of disclosure tendencies and privacy behaviors [47, 56]. Moreover, research shows that using machine learning techniques, users can be sorted into groups that demonstrate similar behaviors. These “privacy profiles” can be built on top of the dimensions of “what” (see above) or based directly on their privacy-setting behaviors [48]. For example, Fig. 16.3

Fig. 16.3 Facebook users’ privacy management strategies

shows the profiles related to a number of personal privacy management strategies uncovered by Wisniewski et al. [43].

As explained in Chap. 12, there also exist differences in privacy behavior based on universal cultural dimensions. Notably, Li et al. demonstrate that if UTP takes culture into account, this will likely result in a significant improvement in the accuracy of privacy predictions [57]. Finally, at the individual level, demographic differences and differences in personality have been used to predict users’ privacy concerns in prior research. The most prominent demographic characteristics that affect privacy behavior are age, gender, education level, rural/urban background, and income level [58]. Results regarding personality seem rather inconsistent, though [59].

16.3.3 The Recipient (To Whom)

The recipient of the information is another important variable affecting users’ privacy decision-making. Whereas decisions regarding people as recipients are usually governed by social conventions, decisions about applications as recipients are governed by information privacy concerns [7, 47]. Similar to the “what” and “who” aspects, recipients can be clustered into groups or “circles” [49, 50]. This approach imposes some structure on the effect of recipient, and these decisions are mainly moderated by the degree the user trusts those recipients [60].

16.3.4 Other Factors

There are other factors such as the location, time of day, the usefulness of the information, and users' emotion [52–54]. The relevance of these factors differs per system, so researchers and practitioners should study them to further improve their predictions of users' privacy decisions.

Many contextual influences on privacy concerns and behavior are due to purpose specificity. This is related to the concept of contextual integrity, which is defined in terms of informational norms which render certain attributes appropriate or inappropriate in certain contexts, under certain conditions [13]. For example, people will be less likely to share their location with their employer at night, when they are not on the job. As such, UTP can leverage the logic of purpose specificity to improve privacy predictions.

16.4 Modeling Privacy

Researchers and practitioners interested in implementing UTP can model users' privacy preferences using off-the-shelf machine learning algorithms. In doing so, they should consider the potentially dynamic aspect of users' privacy decisions (i.e., their privacy preferences may change over time), the way the algorithm may balance the cost of over- versus under-disclosure, the potential trade-off between privacy and other user and/or system goals, and the impact of traditional machine learning problems like overfitting and the cold start problem.

16.4.1 Types of Input

There are two ways to model users' privacy preferences: direct observation of user behaviors and inference from users' attitudes. The most rudimentary privacy decision behavior is users' decision to disclose or withhold information (or, analogously, to allow or reject a certain collection or sharing of data to occur automatically). Such behaviors have been successfully used in user privacy modeling [36, 40, 41, 53, 54, 61]. Some platforms (especially social networks) may offer users a plethora of features to control various aspects of their privacy. Users' use of such controls has been modeled successfully as well [43]. Moreover, when UTP gives users explicit privacy recommendations, users' acceptance or rejection of these recommendations can be another important user modeling input, especially when it is combined with other data [62].

Users' privacy preferences can also be measured directly [42, 48, 61]. While attitudinal data are generally more difficult to collect than behaviors, they are often more stable and precise, because behavior tends to fall prey to heuristic influences.

Likewise, users' privacy preferences can be derived from their traits, such as their culture, demographics, job title, or personality, although these relationships are often weaker [48, 53, 57, 63].

Aside from users' preferences, context is an essential user-modeling input for UTP, as privacy decisions are likely to be heavily context-dependent. It is important to automatically include context in UTP's user model, as the absence of contextual information would likely increase the number of times users have to interact with their privacy settings [64].

Contextual variables can introduce a large amount of sparsity to a user model or result in overfitting. Luckily, various mechanisms exist that allow a recommender system to integrate contextual variables without significantly reducing the amount of data available for each prediction context [54]. Another way to prevent overfitting and sparsity is to have a psychological theory behind the measurement of certain contextual variables [53]. An example of this is Toch et al.'s [65] realization that entropy is the most important aspect of location sensitivity and Li et al.'s [57] finding that country-level cultural variables are a better predictor than country itself.

In sum, we recommend the following practices for gathering input data:

- **Model users' decisions** to disclose/withhold information, to allow/reject tracking, and to use/avoid privacy features.
- **Collect implicit feedback** on privacy recommendations provided by UTP.
- **Elicit privacy attitudes/preferences** to factor out the influence of heuristic decision-making.
- **Collect user traits** to kick start UTP's user model.
- **Take context into account** when modeling users' privacy preferences.
- **Respect the user's privacy** in gathering all this input data.

16.4.2 Algorithms

Algorithms can calculate privacy recommendations using one of two methods: "collaborative filtering" methods, which rely on other users' behaviors, and "case-based reasoning" methods, which rely on the target user's behavior only.

Collaborative filtering leverages other users' privacy behaviors to help predict the current user's privacy preferences [36, 54]. An example is the nearest neighbor approach, where the target user's behaviors are matched with other users' behaviors in an attempt to find the users who are most similar to the target user. Once a set of nearest neighbor users has been found, any unknown preferences of the target user can be predicted using the preferences of these neighbors. This approach is called "user-based collaborative filtering," as opposed to "item-based collaborative filtering," which applies the same approach to the items instead of the users.

Privacy recommendations that are based on collaborative filtering can "leak" information about users' privacy preferences, and thus create security violations. For example, if a hacker has access to users' privacy settings, they may be able to

derive from these settings what kind of information users find most sensitive. Zhao et al. propose a system that treats users' privacy recommendations themselves as sensitive information, and in response, they build a differentially private privacy recommender using standard data obfuscation techniques [66].

Case-based reasoning applies contextualized rules to decide on the best outcome in a given situation [57, 61, 64]. The rules can be based on common sense (e.g., recommendations could be predefined for various types of applications) or based on (past) data of other users (e.g., past user data could be used to establish a "privacy score" for each type of app, which then informs future recommendations). Regardless, one benefit of case-based recommendation is that the system does not require "live" user data (which reduces the chance of "privacy leaking") and can easily be implemented on the client side (which voids the need for user data to be shared with the recommender) [67].

Personalized case-based reasoning systems are usually profile-based, where the specific set of rules to apply to the target user's behavior depends on the profile that was assigned to them. Profiles turn the user modeling from a multidimensional tracking problem into a simpler classification problem [48] and offer personalization without requiring a central server to calculate the recommendations. A downside of case-based privacy recommendations is that its rules are static: unless the algorithm behind the rules is re-trained, the rules will not change even if users' behavior evolves.

Finally, in predicting users' privacy preferences, it is important to keep the user "in the loop" regarding the origin of and reasoning behind these predictions. UTP should therefore adopt explainable and user-controllable algorithms, even if this is at the cost of prediction accuracy. The following practices are recommended in this regard:

- **Make it easy for users to give preference input** or explicit feedback, so that they can control and, where needed, correct the recommendations [62].
- **Build conversational recommenders** that engage in a dialogue with users to uncover their privacy preferences [68].
- **Implement explanations** as a means to increase users' understanding of the recommendation process, their trust in the quality of the recommendations, and their perception of competence and benevolence of the system [69].

16.4.3 The Adaptation Target: What Should UTP Try to Accomplish?

Once the user's privacy behavior or attitude is known, the question remains how UTP should adapt to this behavior/attitude. This can be done in several ways:

- **Match the user's current behaviors**—this alleviates their decision-making burden, especially when done through automation.

- **Recommend practices that dovetail with users' current behavioral patterns**—this solidifies their behavioral practices. It is best to use suggestions or highlights to make such recommendations.
- **Move beyond current behavioral patterns**—this encourages exploration and self-actualization. Research shows that users appreciate such a personalized educational approach, especially when it is done through active, well-explained recommendations [45].

Note that these suggestions pair each recommendation target (match, dovetail, move beyond) with a particular recommendation method (automate, highlight, explain). Other pairings may not work well. For example, actively recommending users privacy behaviors that they already engage in can be regarded as redundant or a nuisance [45]. Likewise, pushing users beyond their current behavioral patterns without careful explanation can cause reactance (see Sect. 16.7.2).

16.5 Adapting the System

While most existing work on UTP covers the modeling aspect, it is of utmost importance to also cover various ways in which systems can leverage these user privacy models to provide user-tailored privacy decision support. Particularly, this section covers the following adaptations:

- **Intelligent privacy settings**—These adaptations alleviate the burden of privacy decision-making, either through fully automated adaptive defaults or adaptive nudges in the form of highlights or suggestions.
- **Augmented privacy notices**—These adaptations inform users about the reasons behind a recommendation or act as a nudge that gives users a rationale for engaging in a privacy-related behavior. More complex justifications can educate users about the risks and benefits involved in a privacy decision.
- **Adaptive privacy-setting interfaces**—These adaptations restructure the user interface of the system to make certain privacy actions easier to accomplish.
- **Privacy-aware personalization**—These adaptations influence the types of personalization a system can engage in based on the collected user data, thereby preventing potential unwanted inferences to be made.

16.5.1 *Intelligent Privacy Settings*

The most commonly studied application of UTP is “adaptive privacy settings.” These are essentially adaptive versions of nudges or defaults [34]. Unlike traditional nudges and defaults, they consider the crucial role of the user and decision context, thereby limiting threats to user autonomy.

Adaptive default settings make it easier for users to choose the right settings, since most settings will already be aligned with their preferences [70]. A large number of existing works on privacy prediction assume that users will benefit from this relief in privacy-setting burden, but very few works test whether users indeed appreciate—and go along with—such adaptive defaults. Early work in this regard by Namara et al. suggests that UTP should implement a “hybrid” adaptive privacy-setting procedure, along the following guidelines [45]:

- **Automatically apply settings to alleviate users** from frequent privacy behaviors but avoid automating decisions with far-reaching consequences.
- **Highlight suggested settings to reduce users’ cognitive burden** in a subtle but useful manner.
- **Suggest privacy settings to keep users involved in their privacy decisions** but avoid making suggestions with awkward social consequences.

Finally, in cases where settings and/or disclosure requests are presented in a sequential manner, the order of sensitive versus less sensitive requests has an impact on disclosure [29]. For example, disclosure rates are lower when asking less sensitive questions first, and sharing rates in social networks are higher when users are asked to share with weaker ties first [51]. Generally speaking, disclosure is higher for information that is requested first, so UTP can adapt the order of sequentially presented settings and information requests to prioritize the disclosure of certain types of information.

16.5.2 *Augmented Privacy Notices*

A justification is usually accompanied by a recommended setting or action and provides a succinct reason to engage or not engage in the recommended privacy-related behavior. In general, justifications seem to have no positive effect on users’ privacy decision-making [29, 71], but justifications that are adapted to the user seem to have a positive effect [63]. Moreover, context-relevant justifications are likely more effective if shown when users are actually in the process of making a decision, rather than “post-hoc” [26]. And finally, justifications can be used to frame a privacy decision, which can significantly influence the level of disclosure, with negative framing leading to significantly lower levels of disclosure than positive framing [71]. As such, we make the following recommendations:

- **Make justifications context-relevant**—This helps users understand the purpose of the request or decision.
- **Time justifications carefully**—Justifications should only happen in situations where they may have a short-term (e.g., impact the user’s decision) or long-term (e.g., increase the user’s privacy knowledge) impact.
- **Tailor justification types** (explanation, usefulness, or social norm) to the user’s personal characteristics.
- **Leverage the framing of justifications** to adaptively influence disclosure.

Fig. 16.4 Bringing certain privacy features to the forefront: (a) Users can directly change the audience of a post with toggle buttons, without having to use the standard drop-down list. (b) A more prominent design for blocking apps, events, and people that is displayed directly in the notifications

16.5.3 Adaptive Privacy-Setting Interfaces

Privacy features are often difficult to access, and they create an unwieldy “labyrinth” of privacy functionality that users find difficult to use [72]. UTP can be used to tailor the design of the interface itself to the user. This approach, labeled User-Tailored Privacy by Design (UTPbD) [43], emphasizes features the user is expected to use most often and de-emphasizes features they only seldom use (see Fig. 16.4). UTPbD can be implemented in two ways:

- **A direct application of UTPbD** involves profiling the users of a system based on their privacy behaviors, and then tailoring the privacy controls of the system in a way that changes their salience depending on the profile of the current user.
- **An extrapolated application of UTPbD** involves turning the profiles uncovered in a baseline system (e.g., Facebook) into “personas” to develop privacy design guidelines for a yet-to-be-implemented system that is envisioned to have similar privacy features (e.g., a new social network).

The latter approach may not result in a user-tailored solution per say, but it uses the “measure” and “model” aspects of the UTP framework for persona development.

16.5.4 Privacy-Aware Personalization

UTP models should not only acknowledge and account for potential differences in users’ attitudes toward data *collection*, but also data *use*. The latter is particularly important in systems that use users’ data to personalize their content or its presentation. In providing a personalized experience, such systems may make unexpected (or unwanted) inferences about the user. As such, users may prefer

that the system *not* use (a subset of) their data for certain personalization purposes. To address users' preferences regarding the use of their data for personalization purposes, UTP can be used to adapt a system's personalization procedures to the users' data use preferences.

Particularly, UTP can be integrated into a dynamic privacy-enabling user modeling framework. Wang and Kobsa developed such a framework to consider prevailing privacy laws and regulations based on the country of residence of the user [73]. Our suggestion is to support a much more granular level of personalized constraints based on UTP: Users' privacy preferences as modeled by UTP can serve as a basis to determine whether certain types of inferences or data use should be allowed and disallowed for each particular user.

16.6 What Should Be the End-Goal of UTP?

UTP can serve multiple goals. While UTP has traditionally been envisioned to alleviate the user's privacy decision-making burden, UTP can also support users by taking on a "teaching role" and giving them the tools they need to decide for themselves how to meet their privacy goals. To further complicate things, UTP can take the privacy requirements of the recipient of the information, other users, and the community or organization in which it operates into account as well. Below, we discuss these goals in more detail and outline strategies for reconciling conflicting goals among multiple stakeholders.

16.6.1 *Support the User*

The main goal of UTP is to support the user. However, in providing such support we should note that users have conflicting motivations and goals in making such privacy decisions. For example, users of recommender systems, the goals of privacy, and recommendation quality are in conflict, and the balance depends on what level of recommendation quality is considered "acceptable" and what kind of tracking users consider to be "comfortable." It will be up to system designers to decide which of these goals to prioritize and what policies to adopt. In this regard, collecting more data (see Sect. 16.3) will help the system to identify the conflicts and the contexts of each decision.

Finally, a conflict exists in that users usually claim to want control over their privacy, without wanting to spend a significant amount of time on setting privacy settings [22]. UTP can help users by automating their privacy decisions, but this could eventually result in users who are disinterested and ill-equipped to set any settings manually. This could in turn result in an erosion of user autonomy [74]

which can cause severe privacy violations. Hence, we make the following recommendations:

- **Make most of the individual privacy decisions automatically**—Given the complexity of the privacy trade-offs in most systems, this is the only way to avoid overwhelming users.
- **Allow the user to control the general rules by which these decisions are made**—Arguably, users will be much better equipped at making higher-level privacy decisions, rather than lower-level individual settings.
- **Where appropriate, raise exceptions to ask for user feedback**, and incorporate such feedback into the user privacy model maintained by UTP.
- **Allow users to share their “stewardship” over their data** with a trusted third party.

16.6.2 *Teach the User*

In addition to alleviating the user’s decision burden, UTP adaptations can also be used to teach the user about privacy based on their behavior and current level of knowledge. This way, UTP allows users to decide for themselves how to meet their privacy goals. The field of recommender systems has shown that explanations increase users’ understanding of the recommendation process [75]. Hence, they ascertain that users are actively involved in the decision process, which is instrumental in increasing their understanding and teaching them about their privacy. We recommend two useful venues for such explanations:

- **Provide personalized privacy “tips”** that highlight users’ inconsistent behaviors and bring them to their attention.
- **Use tailored privacy education** to give users more confidence in their overall privacy strategy and to support the evolution of this strategy.

Ghaiumy Anaraky et al. [71] introduced tips in a privacy decision-making scenario and demonstrated that combined with other nudges they can affect disclosure behavior. Privacy tips are envisioned to work particularly well when users exhibit inconsistent privacy-related behaviors (e.g., withholding a piece of information from a colleague, while at the same time posting it on a public site), or when users fail to engage in synergistic privacy behaviors (e.g., grouping Facebook friends into circles, but not using these circles to selectively target posts to specific audiences) [43]. UTP can personalize the tips to these occasions, which should increase the effectiveness of privacy tips and minimize the chance that they will be ignored.

In contrast, privacy education works best when recommending actions outside the user’s current purview. This helps users to adopt new privacy protection strategies. Note, however, that education is not effective for the practices the user already engages in, nor for practices the user has decided they do not want to engage

in. UTP can actively avoid these situations as a means to optimize the effectiveness of privacy education.

16.6.3 Help the Recipient

UTP can also advocate for the recipient of the user's information, or even society at large. For example, an app that tracks epidemics will function better if users are willing to share their health status and location, and this would allow authorities to appropriately allocate resources to mitigate or prevent outbreaks. In this case, the privacy of each individual user is in conflict with the success of the application and its benefit for society. As such, governments (or, in a corporate context, the user's employer) may mandate certain data collection practices (e.g., reporting job training results) or even prohibit certain disclosures (e.g., for mission-critical data to be shared on public servers).

In case the user's privacy is in conflict with organizational or societal goals or regulations, it is important to increase user trust in the recipients of data, if their long-term data disclosure is to be maintained. Therefore, we make the following recommendations:

- **Make sure that users never feel pressured to provide sensitive data**—In these cases, they are likely to provide fake data instead.
- **Use justifications** to explain why certain organizational or legal constraints are in place.
- **Utilize collected data in a way that aligns with users' expectations** with regard to the purposes of disclosure.

16.6.4 Reconciling the Differences

In some cases, UTP must make trade-offs to reconcile the preferences of stakeholders with conflicting goals (e.g., users, recipients of data, other users, the organization). The field of group recommender systems has studied various ways to integrate the preferences of multiple stakeholders in the recommendation process [72]. A useful ethical principle for reconciling the goals of multiple stakeholders is to always put the end-user first, using the principles of fairness [76] and reciprocity [77]. However, these principles do not give sufficient recourse in cases of organizational constraints (which often cannot be traded off against users' preferences) or conflicts between users (who may have equal stake in the management of a piece of co-created data). We recommend the following practices in reconciling the conflicting goals of multiple stakeholders:

- **Develop UTP as a multi-stakeholder recommender system** if it is likely that there exist multiple conflicting goals around privacy.

- **Put the end-user first** and apply the principles of reciprocity and fairness to reconcile conflicting goals among multiple stakeholders.
- **Honestly inform users about the optimization strategy** of UTP, that is, to explain to them how different conflicting internal and external goals are taken into consideration.
- **Allow users to reflect upon UTP’s privacy policies at a meta level**—UTP can manage the user’s privacy settings, and users can manage the principles by which UTP operates to reconcile multi-stakeholder differences.

16.7 Problems That May Arise When Implementing UTP

16.7.1 Problems Related to Privacy Modeling

Potential problems related to modeling users’ privacy decisions should be considered when implementing UTP. For instance, due to the privacy paradox [78], there will likely be a difference between UTP user models that are based on users’ and models that are based on their attitudes or preferences of input is more suitable for privacy modeling purposes.

Likewise, typical metrics of prediction accuracy such as F1 and AUC treat “false positives” and “false negatives” as equally bad, but this may not be desirable in privacy prediction settings. To complicate matters, privacy decisions are rarely made in isolation but usually considered as a trade-off with other goals. In making this trade-off, one would have to decide whether the trade-off should be a compensatory (linear) trade-off or a non-compensatory (threshold-based) trade-off. Moreover, it requires estimates of the decision options in light of these various goals, which can be either derived from objective system parameters or subjective user experiences [79].

Most user-modeling systems have a problem of overfitting: the more granular a user’s contextual preferences get, the less data the predictions will be based on [54]. Privacy prediction needs a certain amount of input before it can determine the user’s privacy preferences. Without input, it is impossible to create an accurate user model. In the user modeling community, this problem is known as the “cold start” problem.

Finally, we must acknowledge the fact that privacy recommendations can have a persuasive effect on users [42], which may result in “positive feedback loops” of users accepting a suboptimal privacy recommendation and thereby reinforcing the user model in the wrong direction.

To overcome these problems, we recommend that UTP should be implemented using a layered and gracefully degrading approach:

- **Start with simple “smart profile”-based approaches** [61] when implementing UTP in new systems, then move to more complex user privacy modeling solutions once more user modeling data is available.

- **Study the relative cost of over-disclosure versus under-disclosure** and build this cost into the UTP algorithms.
- **Trade-off privacy with other user goals**, such as the goals of the system, the institution, and other users.
- **Couch the recommendation logic in psychological principles** to prevent overfitting [53].

Combining simple and complex user modeling within the same system allows for “graceful degradation” of the user modeling approach. For example, a collaborative filtering recommender will not work when too little user data is available or when the user is offline. In such cases, the system can fall back on a profile-based approach or even just recommend the settings of the average user.

Finally, to ascertain the quality and validity of the UTP module, it should be evaluated using the “layered evaluation” approach [80], which allows input, processing, and output procedures to be evaluated separately.

16.7.2 Problems Related to Adapting the System

Adaptations span a variety of “degrees of automation,” which present problems for implementing UTP: fully automated UTP can be overly persuasive and difficult to control, while low-automation UTP can be a burden on the user.

Users may show reactance toward adaptations, especially in domains where adaptations are not expected—privacy is one such domain. In this case, it is better to start off with a less automated approach. Once users have gained enough trust in the privacy adaptation procedure, they may choose to accept subsequent adaptations automatically, thereby reducing their decision-making burden [45].

Adaptations that users may find unexpected should be explained to the users to allow them to judge the integrity, benevolence, competence, accuracy, and overall usability of the adaptations [69]. However, it is important to avoid pressuring users into accepting adaptations that they might not want to accept [81]. Users should be given various options to choose from to possibly empower them to make better decisions on their own.

Given these potential problems, we make the following recommendations:

- **Find the optimal adaptation method** with regard to the goals of UTP, be it automation, awareness, guidance, or education (see Sect. 16.6).
- **Give users explicit suggestions if they are unfamiliar with a privacy feature**—This allows for the adaptive behavior to be explained, which can increase trust.
- **When users use a feature frequently, use the fully automated approach**—Users in this situation are willing to give up some control in return for the significant reduction in the burden that this approach offers them.

16.8 Conclusion

In this chapter we introduced the concept of User-Tailored Privacy (UTP) as a means to support the privacy management practices of the users of modern information systems. We hope that our recommendations will help researchers and practitioners to implement UTP in their systems.

We made a case for UTP by highlighting the shortcomings of technical solutions, Privacy by Design, notice and choice, and privacy nudging. We recommend that researchers and practitioners should not avoid these existing practices, but rather complement them with the personalized approach afforded by UTP.

We then defined UTP as a “measure, model, adapt” framework. And covered each of these steps separately:

- We argued that in measuring privacy, researchers and practitioners should acknowledge the plurality and multidimensionality of users’ privacy decision-making practices. They should also note that the variability of users’ practices can often be captured by a concise set of “privacy profiles” and that data recipients can often similarly be grouped into a number of groups or “circles.”
- In modeling privacy, we particularly noted that matching the users’ current privacy practices may not always be the best modeling strategy; in certain cases, UTP should recommend complementary practices, while in other cases UTP can completely move beyond users’ current practices. Researchers and practitioners should carefully balance these various approaches. Moreover, since privacy modeling may not always be successful, UTP should be implemented as a layered and gracefully degrading modeling component.
- In adapting privacy, we argued that UTP can personalize the privacy settings of an application, the justification it gives for requesting certain information, its privacy-setting interface, and its personalization practices. Researchers and practitioners should carefully balance proactive and conservative adaptation strategies in order to reduce users’ burden but at the same time give them sufficient control and reduce undue persuasion.

Finally, we argued that researchers and practitioners should carefully consider the various goals that UTP can support. They should acknowledge that UTP must reconcile users’ potentially conflicting goals, and they should balance the goal of replacing the users’ privacy decision-making practices with the longer-term goal of teaching them about privacy. Moreover, researchers and practitioners should consider that UTP’s support can help other stakeholders in the privacy decision-making process as well. Regarding this, they should carefully consider how to reconcile the potentially conflicting goals of these various stakeholders.

While we were able to leverage existing work to make extensive recommendations regarding the implementation of UTP, we must at the same time acknowledge that UTP is still a relatively novel and underexplored solution to users’ privacy problems. As such, we encourage privacy researchers to investigate UTP in its various incarnations and to contribute to the growing body of literature around this

topic. We hope that this chapter provides a starting point for them as well, as we have highlighted gaps in existing research throughout the chapter.

References

1. Berkovsky, S., Y. Eytani, T. Kuflik, and F. Ricci. 2006. Hierarchical neighborhood topology for privacy enhanced collaborative filtering. In *Proceedings of PEP06, CHI 2006 Workshop on Privacy-Enhanced Personalization*, Montreal, Canada, ed. Kobsa, A., R. Chellappa, and S. Spiekermann, 6–13. http://www.isr.uci.edu/pep06/papers/PEP06_BerkovskyEtAl.pdf.
2. Kobsa, A., H. Cho, B.P. Knijnenburg. 2016. The effect of personalization provider characteristics on privacy attitudes and behaviors: An elaboration likelihood model approach. *Journal of the Association for Information Science and Technology*. <https://doi.org/10.1002/asi.23629>, <http://onlinelibrary.wiley.com/doi/10.1002/asi.23629/abstract>.
3. Nikolaenko, V., S. Ioannidis, U. Weinsberg, M. Joye, N. Taft, and D. Boneh. 2013. Privacy-preserving matrix factorization. In *Proceedings of the 2013 ACM SIGSAC Conference on Computer & Communications Security, CCS '13*, 801–812. New York, NY: ACM. <http://doi.org/10.1145/2508859.2516751>
4. Dwork, C., and M. Naor. 2008. On the difficulties of disclosure prevention in statistical databases or the case for differential privacy. *Journal of Privacy and Confidentiality* 2 (1). <http://repository.cmu.edu/jpc/vol2/iss1/8>.
5. Narayanan, A., and V. Shmatikov. 2008. Robust de-anonymization of large sparse datasets. In *2008 IEEE Symposium on Security and Privacy*, 111–125. IEEE. <https://doi.org/10.1109/SP.2008.33>.
6. Singel, R. 2009. Netflix spilled your brokeback mountain secret, lawsuit claims, Dec 2009. <http://www.wired.com/2009/12/netflix-privacy-lawsuit/?+wired27b+%2528Blog++27B+Stroke+6+%2528Threat+Level%2529%2529/>.
7. Lederer, S., J.I. Hong, A.K. Dey, and J.A. Landay. 2004. Personal privacy through understanding and action: Five pitfalls for designers. *Personal and Ubiquitous Computing* 8 (6): 440–454 (2004). <https://doi.org/10.1007/s00779-004-0304-9>, <http://www.springerlink.com/index/10.1007/s00779-004-0304-9>.
8. Cavoukian, A. 2010. Privacy by Design. Information and Privacy Commissioner of Ontario, Canada. <http://www.privacybydesign.ca/content/uploads/2010/03/PrivacybyDesignBook.pdf>.
9. Awad, N.F., and M.S. Krishnan. 2006. The personalization privacy paradox: An empirical evaluation of information transparency and the willingness to be profiled online for personalization. *MIS Quarterly* 30 (1): 13–28.
10. Strater, K., and H.R. Lipford. 2008. Strategies and struggles with privacy in an online social networking community. In *Proceedings of the 22nd British HCI Group Annual Conference on People and Computers*, 111–119. Swinton: British Computer Society.
11. Hann, I.H., K.L. Hui, S.Y. Lee, and I. Png. 2007. Overcoming online information privacy concerns: An information-processing theory approach. *Journal of Management Information Systems* 24 (2): 13–42. <https://doi.org/10.2753/MIS0742-1222240202>, <http://mesharpe.metapress.com/openurl.asp?genre=article&id=doi:10.2753/MIS0742-1222240202>.
12. Taylor, D., D. Davis, and R. Jillapalli. 2009. Privacy concern and online personalization: The moderating effects of information control and compensation. *Electronic Commerce Research* 9 (3): 203–223. <https://doi.org/10.1007/s10660-009-9036-2>.
13. Nissenbaum, H. 2011. A contextual approach to privacy online. *Daedalus* 140 (4): 32–48. https://doi.org/10.1162/DAED_a_00113.
14. Gardner, J. 2012. 12 Surprising A/B test results to stop you making assumptions. <http://unbounce.com/a-b-testing/shocking-results/>.
15. Pollach, I. 2007. What's wrong with online privacy policies? *Communications of the ACM* 50 (9): 103–108. <https://doi.org/10.1145/1284621.1284627>.

16. Aagaard, M. 2013. How privacy policy affects sign-ups surprising data from 4 A/B tests. <http://contentverve.com/sign-up-privacy-policy-tests/>.
17. Bustos, L. 2012. Best practice gone bad: 4 shocking A/B tests. <http://www.getelastic.com/best-practice-gone-bad-4-shocking-ab-tests/>.
18. Adjerd, I., A. Acquisti, L. Brandimarte, and G. Loewenstein. 2013. Sleights of privacy: Framing, disclosures, and the limits of transparency. In *Proceedings of the Ninth Symposium on Usable Privacy and Security*, SOUPS '13, 9:1–9:11. New York, NY: ACM. <https://doi.org/10.1145/2501604.2501613>.
19. Milne, G.R., and M.J. Culnan. 2004. Strategies for reducing online privacy risks: Why consumers read (or don't read) online privacy notices. *Journal of Interactive Marketing* 18 (3): 15–29. <https://doi.org/10.1002/dir.20009>.
20. Jensen, C., C. Potts, and C. Jensen. 2005. Privacy practices of internet users: Self-reports versus observed behavior. *International Journal of Human-Computer Studies* 63 (1–2): 203–227. <https://doi.org/10.1016/j.ijhcs.2005.04.019>, <http://www.sciencedirect.com/science/article/pii/S1071581905000650>.
21. Pan, Y., and G.M. Zinkhan. 2006. Exploring the impact of online privacy disclosures on consumer trust. *Journal of Retailing* 82 (4): 331–338. <https://doi.org/10.1016/j.jretai.2006.08.006>, <http://www.sciencedirect.com/science/article/pii/S0022435906000558>.
22. Compagno, R., and W. Lusoli. 2010. The policy maker's anguish: Regulating personal data behavior between paradoxes and dilemmas. In *Economics of Information Security and Privacy*, ed. Moore, T., D. Pym, and C. Ioannidis, 169–185. New York, NY: Springer. https://doi.org/10.1007/978-1-4419-6967-5_9.
23. Larose, R., and N.J. Rifon. 2007. Promoting I safety: Effects of privacy warnings and privacy seals on risk assessment and online privacy behavior. *Journal of Consumer Affairs* 41 (1): 127–149. <https://doi.org/10.1111/j.1745-6606.2006.00071.x>, <http://onlinelibrary.wiley.com/doi/10.1111/j.1745-6606.2006.00071.x/abstract>.
24. Besmer, A., J. Watson, and H.R. Lipford. 2010. The impact of social navigation on privacy policy configuration. In *Proceedings of the Sixth Symposium on Usable Privacy and Security*, Redmond, Washington, 7:1–7:10. <https://doi.org/10.1145/1837110.1837120>.
25. Metzger, M.J. 2006. Effects of site, vendor, and consumer characteristics on web site trust and disclosure. *Communication Research* 33 (3): 155–179. <https://doi.org/10.1177/0093650206287076>.
26. Felt, A.P., E. Ha, S. Egelman, A. Haney, E. Chin, and D. Wagner. 2012. Android permissions: User attention, comprehension, and behavior. In *Proceedings of the Eighth Symposium on Usable Privacy and Security*, SOUPS '12, 3:1–3:14. New York, NY: ACM. <http://doi.org/10.1145/2335356.2335360>.
27. Acquisti, A. 2009. Nudging privacy: The behavioral economics of personal information. *IEEE Security and Privacy* 7:82–85. <https://doi.org/10.1109/MSP.2009.163>, ACM ID: 1685896.
28. Wang, Y., P.G. Leon, A. Acquisti, L.F. Cranor, A. Forget, and N. Sadeh. 2014. A field trial of privacy nudges for facebook. In *Proceedings of the 32nd Annual ACM Conference on Human Factors in Computing Systems*, CHI '14, 2367–2376. Toronto: ACM. <https://doi.org/10.1145/2556288.2557413>.
29. Knijnenburg, B.P., and A. Kobsa. 2013. Making decisions about privacy: Information disclosure in context-aware recommender systems. *ACM Transactions on Interactive Intelligent Systems* 3 (3): 20:1–20:23. <https://doi.org/10.1145/2499670>, <http://bit.ly/tiis2013>.
30. Rifon, N.J., R. LaRose, and S.M. Choi. 2005. Your privacy is sealed: Effects of web privacy seals on trust and personal disclosures. *Journal of Consumer Affairs* 39 (2): 339–360 (2005). <https://doi.org/10.1111/j.1745-6606.2005.00018.x>.
31. Jedrzejczyk, L., B.A. Price, A.K. Bandara, and B. Nuseibeh. 2010. On the impact of real-time feedback on users' behaviour in mobile location-sharing applications. In *Proceedings of the Sixth Symposium on Usable Privacy and Security*, Redmond, Washington, 14:1–14:12. <https://doi.org/10.1145/1837110.1837129>, <http://portal.acm.org/citation.cfm?doid=1837110.1837129>.

32. Johnson, E.J., S. Bellman, and G.L. Lohse. 2002. Defaults, framing and privacy: Why opting in opting out. *Marketing Letters* 13 (1): 5–15. <https://doi.org/10.1023/A:1015044207315>, <http://www.springerlink.com/content/vrf4lmw6jgnq3pvd/abstract/>.
33. Knijnenburg, B.P. 2015. A user-tailored approach to privacy decision support. Ph.D. thesis, UC Irvine.
34. Smith, N.C., D.G. Goldstein, and E.J. Johnson. 2013. Choice without awareness: Ethical and policy implications of defaults. *Journal of Public Policy & Marketing* 32 (2): 159–172. <https://doi.org/10.1509/jppm.10.114>, <http://search.ebscohost.com/login.aspx?direct=true&db=bth&AN=91886736&site=ehost-live>.
35. Kobsa, A. 2001. Tailoring privacy to users' needs (invited keynote). In *User Modeling 2001*, ed. Bauer, M., P.J. Gmytrasiewicz, and J. Vassileva, 303–313. No. 2109 in Lecture Notes in Computer Science, Springer. https://doi.org/10.1007/3-540-44566-8_52.
36. Liu, B., J. Lin, and N. Sadeh. 2014. Reconciling mobile app privacy and usability on smartphones: Could user privacy profiles help? In *Proceedings of the 23rd International Conference on World Wide Web*, 201–212. WWW '14, International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, Switzerland. <https://doi.org/10.1145/2566486.2568035>.
37. Cranor, L.F. 2002. *Web Privacy with P3P*. Sebastopol, CA: O'Reilly & Associates, Inc.
38. Harkous, H., R. Rahman, and K. Aberer. 2016. Data-driven privacy indicators. In *Twelfth Symposium on Usable Privacy and Security (SOUPS 2016)*. Denver, CO: USENIX Association. <https://www.usenix.org/conference/soups2016/workshop-program/wpi/presentation/harkous>.
39. Egelman, S., J. Tsai, L.F. Cranor, and A. Acquisti. 2009. Timing is everything? The effects of timing and placement of online privacy indicators. In *Proceedings of the 27th International Conference on Human Factors in Computing Systems*, 319–328.
40. Fang, L., and K. LeFevre. 2010. Privacy wizards for social networking sites. In *Proceedings of the 19th International Conference on World Wide Web*, WWW '10, 351–360. New York, NY: ACM. <http://doi.org/10.1145/1772690.1772727>.
41. Ravichandran, R., M. Benisch, P. Kelley, and N. Sadeh. 2009. Capturing social networking privacy preferences. In *Privacy Enhancing Technologies*, ed. Goldberg, I., and M. Atallah. Lecture Notes in Computer Science, 1–18. Vol. 5672. Berlin: Springer. https://doi.org/10.1007/978-3-642-03168-7_1.
42. Knijnenburg, B.P., and H. Jin. 2013. The persuasive effect of privacy recommendations. In *Twelfth Annual Workshop on HCI Research in MIS* (2013)
43. Wisniewski, P.J., B.P. Knijnenburg, and H.R. Lipford. 2017. Making privacy personal: Profiling social network users to inform privacy education and nudging. *International Journal of Human-Computer Studies* 98: 95–108. <https://doi.org/10.1016/j.ijhcs.2016.09.006>, <http://www.sciencedirect.com/science/article/pii/S1071581916301185>.
44. Wilkinson, D., S. Sivakumar, D. Cherry, B.P. Knijnenburg, E.M. Raybourn, P. Wisniewski, and H. Sloan. 2017. User-tailored privacy by design. In *Proceedings of the Usable Security Mini Conference*, USEC '17. San Diego, CA: Internet Society. <https://doi.org/10.14722/usec.2017.23007>.
45. Namara, M., H. Sloan, P. Jaiswal, and B.P. Knijnenburg. 2018. The potential for user-tailored privacy on facebook. In *2018 IEEE Symposium on Privacy-Aware Computing (PAC)*, 31–42. IEEE.
46. Spiekermann, S., J. Grossklags, and B. Berendt. 2001. E-privacy in 2nd generation e-commerce: privacy preferences versus actual behavior. In *Proceedings of the 3rd ACM conference on Electronic Commerce*, 38–47. ACM.
47. Olson, J.S., J. Grudin, and E. Horvitz. 2005. A study of preferences for sharing and privacy. In *CHI'05 Extended Abstracts on Human Factors in Computing Systems*, 1985–1988. ACM (2005)
48. Knijnenburg, B.P., A. Kobsa, and H. Jin. 2013. Dimensionality of information disclosure behavior. *International Journal of Human-Computer Studies* 71 (12): 1144–1162.

49. Kairam, S., M. Brzozowski, D. Huffaker, and E. Chi. 2012. Talking in circles: Selective sharing in google+. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 1065–1074. ACM (2012)
50. Watson, J., A. Besmer, and H.R. Lipford. 2012. + Your circles: Sharing behavior on google+. In *Proceedings of the Eighth Symposium on Usable Privacy and Security*, 12. ACM.
51. Knijnenburg, B.P., and A. Kobsa. 2014. Increasing sharing tendency without reducing satisfaction: Finding the best privacy-settings user interface for social networks. In *International Conference on Information Systems (ICIS)*
52. Benisch, M., P.G. Kelley, N. Sadeh, and L.F. Cranor. 2011. Capturing location-privacy preferences: quantifying accuracy and user-burden tradeoffs. *Personal and Ubiquitous Computing* 15 (7): 679–694.
53. Dong, C., H. Jin, and B.P. Knijnenburg. 2016. PPM: A privacy prediction model for online social networks. In *International Conference on Social Informatics*, 400–420. Springer
54. Xie, J., B.P. Knijnenburg, and H. Jin. 2014. Location sharing privacy preference: analysis and personalized recommendation. In *Proceedings of the 19th International Conference on Intelligent User Interfaces*, 189–198. ACM.
55. Wang, Y., G. Norcie, and L.F. Cranor. 2011. Who is concerned about what? A study of American, Chinese and Indian users privacy concerns on social network sites. In *International Conference on Trust and Trustworthy Computing*, 146–153. Springer. https://doi.org/10.1007/978-3-642-21599-5_11.
56. Ackerman, M.S., L.F. Cranor, and J. Reagle. 1999. Privacy in e-commerce: Examining user scenarios and privacy preferences. In *Proceedings of the 1st ACM Conference on Electronic Commerce*, EC '99, 1–8. Denver, CO: ACM Press. <https://doi.org/10.1145/336992.336995>, ACM ID: 336995.
57. Li, Y., A. Kobsa, B.P. Knijnenburg, and M.C. Nguyen. 2017. Cross-cultural privacy prediction. *Proceedings on Privacy Enhancing Technologies* 2017 (2): 113–132.
58. Hoy, M.G., and G. Milne. 2010. Gender differences in privacy-related measures for young adult facebook users. *Journal of Interactive Advertising* 10 (2): 28–45. <https://doi.org/10.1080/15252019.2010.10722168>.
59. Page, X., B.P. Knijnenburg, and A. Kobsa. 2013. FYI: Communication style preferences underlie differences in location-sharing adoption and usage. In *Proceedings of the 2013 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, 153–162. ACM (2013). <http://dl.acm.org/citation.cfm?id=2493487>
60. Mesch, G.S. 2012. Is online trust and trust in social institutions associated with online disclosure of identifiable information online? *Computers in Human Behavior* 28 (4): 1471–1477. <https://doi.org/10.1016/j.chb.2012.03.010>, <http://www.sciencedirect.com/science/article/pii/S0747563212000763>.
61. Bahirat, P., Y. He, A. Menon, and B. Knijnenburg. 2018. A data-driven approach to developing IoT privacy-setting interfaces. In *23rd International Conference on Intelligent User Interfaces*, IUI '18, 165–176. Tokyo: ACM. <http://doi.org/10.1145/3172944.3172982>.
62. Kelley, P.G., P. Hanks, Drielsma, N. Sadeh, and L.F. Cranor. 2008. User-controllable learning of security and privacy policies. In *Proceedings of the 1st ACM Workshop on Workshop on AI/Sec*, AI/Sec '08, 11–18. New York, NY: ACM. <http://doi.org/10.1145/1456377.1456380>.
63. Knijnenburg, B.P., and A. Kobsa. 2013. Helping users with information disclosure decisions: potential for adaptation. In *Proceedings of the 2013 International Conference on Intelligent User Interfaces*, 407–416. ACM.
64. Pallapa, G., S.K. Das, Di Francesco, M., and T. Aura. 2014. Adaptive and context-aware privacy preservation exploiting user interactions in smart environments. *Pervasive and Mobile Computing* 12: 232–243.
65. Toch, E., J. Cranshaw, P.H. Drielsma, J.Y. Tsai, P.G. Kelley, Springfield, J., L. Cranor, J. Hong, and N. Sadeh. 2010. Empirical models of privacy in location sharing. In *Proceedings of the 12th ACM International Conference on Ubiquitous Computing*, 129–138. ACM (2010).

66. Zhao, Y., J. Ye, and T. Henderson. 2014. Privacy-aware location privacy preference recommendations. In *Proceedings of the 11th International Conference on Mobile and Ubiquitous Systems: Computing, Networking and Services*, 120–129. ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering).
67. Bradley, K., R. Rafter, and B. Smyth. 2000. Case-based user profiling for content personalisation. In *International Conference on Adaptive Hypermedia and Adaptive Web-Based Systems*, 62–72. Berlin: Springer.
68. Cranshaw, J., J. Muga, and N. Sadeh. 2011. User-controllable learning of location privacy policies with Gaussian mixture models. In *Proceedings of the Twenty-Fifth AAAI Conference on Artificial Intelligence*, AI '11, San Francisco, CA, 1146–1152. <http://www.aaai.org/ocs/index.php/AAAI/AAAI11/paper/viewPDFInterstitial/3785/4052>.
69. Wang, W., and I. Benbasat. 2007. Recommendation agents for electronic commerce: Effects of explanation facilities on trusting beliefs. *Journal of Management Information Systems* 23 (4): 217–246. <https://doi.org/10.2753/MIS0742-1222230410>.
70. Dinner, I., E.J. Johnson, D.G. Goldstein, and K. Liu. 2011. Partitioning default effects: Why people choose not to choose. *Journal of Experimental Psychology: Applied* 17 (4): 332–341. <https://www.learnlib.org/p/64714/>.
71. Anaraky, R.G., T. Nabizadeh, B.P. Knijnenburg, and M. Risius. 2018. Reducing default and framing effects in privacy decision-making. In *SIGHCI 2018 Proceedings*.
72. Madejski, M., M. Johnson, and S. Bellovin. 2012. A study of privacy settings errors in an online social network. In *Fourth International Workshop on Security and Social Networking*, SEC-SOC '12, Lugano, Switzerland, 340–345. <https://doi.org/10.1109/PerComW.2012.6197507>
73. Wang, Y., and A. Kobsa. 2013. A PLA-based privacy-enhancing user modeling framework and its evaluation. *User Modeling and User-Adapted Interaction* 23 (1): 41–82. <https://link.springer.com/article/10.1007/s11257-011-9114-8>.
74. Knijnenburg, B.P., S. Sivakumar, and D. Wilkinson. 2016. Recommender systems for self-actualization. In *Proceedings of the 10th ACM Conference on Recommender Systems*, 11–14. ACM.
75. Gedikli, F., D. Jannach, and M. Ge. 2014. How should I explain? A comparison of different explanation types for recommender systems. *International Journal of Human-Computer Studies* 72 (4): 367–382.
76. Ekstrand, M.D., R. Joshaghani, and H. Mehrpouyan. 2018. Privacy for all: Ensuring fair and equitable privacy protections. In *Conference on Fairness, Accountability and Transparency*, 35–47. <http://proceedings.mlr.press/v81/ekstrand18a.html>
77. Pai, P., and H.T. Tsai. 2016. Reciprocity norms and information-sharing behavior in online consumption communities: An empirical investigation of antecedents and moderators. *Information & Management* 53 (1): 38–52. <https://doi.org/10.1016/j.im.2015.08.002>, <http://www.sciencedirect.com/science/article/pii/S0378720615000865>.
78. Norberg, P.A., D.R. Horne, and D.A. Horne. 2007. The privacy paradox: Personal information disclosure intentions versus behaviors. *Journal of Consumer Affairs* 41 (1): 100–126 (2007). <https://doi.org/10.1111/j.1745-6606.2006.00070.x>, <http://onlinelibrary.wiley.com/doi/10.1111/j.1745-6606.2006.00070.x/abstract>.
79. Knijnenburg, B.P., E.M. Raybourn, D. Cherry, D. Wilkinson, S. Sivakumar, and H. Sloan. 2017. Death to the privacy calculus? In *Proceedings of the 2017 Networked Privacy Workshop at CSCW*. Social Science Research Network, Portland, OR, Feb 2017. <https://papers.ssrn.com/abstract=2923806>.
80. Paramythi, A., S. Weibelzahl, and J. Masthoff. 2010. Layered evaluation of interactive adaptive systems: Framework and formative methods. *User Modeling and User-Adapted Interaction* 20 (5): 383–453. <https://doi.org/10.1007/s11257-010-9082-4>.
81. Bösch, C., B. Erb, F. Kargl, H. Kopp, and S. Pfattheicher. 2016. Tales from the dark side: Privacy dark strategies and privacy dark patterns. *Proceedings on Privacy Enhancing Technologies* 2016 (4): 237–254. <https://doi.org/10.1515/popets-2016-0038>. <https://www.degruyter.com/view/j/popets.2016.2016.issue-4/popets-2016-0038/popets-2016-0038.xml?format=INT>.

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.