From procedures to peril: Towards risk transparency in information privacy for users

Telecommunications Policy 50 (2026) 103195

Contents lists available at ScienceDirect

Telecommunications Policy

journal homepage: www.elsevier.com/locate/telpol

From procedures to peril: Towards risk transparency in information privacy for users

Nico Ebert [a,*], Simone Fischer-Hübner [b,l,m], Soheil Human [c], Agnieszka Kitkowska [d], Konrad Kollnig [e], Jelena Mitrović [f,k], Shidong Pan [g], Thierry Schaltegger [a,h], Florian Schaub [i], Daniel Smullen [j], Lu Xian [i]

[a] Zurich University of Applied Sciences (ZHAW), Switzerland
[b] Karlstad University, Sweden
[c] Vienna University of Economics and Business, Austria
[d] Jönköping University, Sweden
[e] Maastricht University, Netherlands
[f] University of Passau, Germany
[g] New York University, USA and Columbia University, USA
[h] University of Zurich, Switzerland
[i] University of Michigan, USA
[j] Independent, USA
[k] Institute for AI Research and Development of Serbia, Serbia
[l] Chalmers University of Technology, Sweden
[m] Gothenburg University, Sweden

ARTICLE INFO

Keywords:
Information privacy
Privacy policies
Privacy risk
Risk transparency

ABSTRACT

Information privacy is an integral part of users' lives, as many digital services and their business models heavily rely on personal data. For example, conversational agents will use massive amounts of user conversations to hyper-personalize ads. Although privacy information is provided through policies and app notifications, and regulation increasingly adopts risk-based approaches, users remain largely uncertain about the risks they face. Design tweaks such as privacy icons or nutrition labels have yielded little improvement, as the central issue lies not in how privacy information is presented, but in what is omitted: the emphasis on disclosing data practices alone does not sufficiently reduce users' uncertainty about potential harms. This paper develops an argument for complementing the current paradigm of "procedural transparency" with "risk transparency." Risk transparency prioritizes the clear communication of privacy risks to individuals using digital services, similar to established practices in domains such as drug safety, public health, or consumer protection, where explicitly informing users about risks is considered the main priority. In this article, we discuss risk transparency terminology, illustrate how risk can be communicated, and review the evidence on the effectiveness of risk communication as well as its associated challenges. A shift towards privacy risk transparency aims to provide consumers and data subjects with more meaningful information that supports their informed decision-making in the data economy.

* Corresponding author. ZHAW, Theaterstrasse 17, 8401, Winterthur, Switzerland.
E-mail address: nico.ebert@zhaw.ch (N. Ebert).

https://doi.org/10.1016/j.telpol.2026.103195
Received 5 December 2025; Received in revised form 23 February 2026; Accepted 12 March 2026
Available online 31 March 2026

1. Introduction

Every time we open an app, scroll through social media, use a smart TV, or use virtually any digital product, we encounter communications ostensibly intended to inform us about personal data collection, use, and sharing - often with limited effectiveness (McDonald & Cranor, 2008; Obar & Oeldorf-Hirsch, 2020). These communications can take the form of lengthy terms and conditions, cookie banners, or privacy notices proclaiming, e.g., "Our conversational agent uses your conversations for training our models." Others appear mid-use as prompts such as, "Do you want to allow this app to access your location?" While they all may create the appearance of empowering users to make informed privacy choices, their underlying function is often more pragmatic: to satisfy legal requirements or to limit liability in case of disputes or misuse. For example, the GDPR (EU, 2016) requires detailed information about processing purposes, recipients, and legal bases (Art. 12–14), and the CCPA (California State Legislature, 2018) mandates disclosure of categories of personal information collected, used, and shared (Section 1798.100). Other privacy frameworks have similar provisions.

Not only laypersons, but even privacy experts often struggle to interpret the information presented in privacy notices (Reidenberg et al., 2015). This difficulty persists despite numerous proposals to make such information more accessible: by simplifying language (Schaub et al., 2017), using icons (Holtz et al., 2011), introducing standardized "nutrition labels" (Cranor et al., 2024; Emami-Naeini et al., 2022; Kelley et al., 2009), using standardized short-form notice templates (Xian et al., 2025), or, more recently, leveraging large language models (LLMs) for the interpretation of policy documents (Freiberger et al., 2025). Yet, even when users do understand the content of a privacy notice, much of it is often irrelevant from their perspective (McDonald & Cranor, 2008). In practice, these communications contribute little to meaningful engagement with privacy risks and instead function primarily as compliance mechanisms (Amos et al., 2021) that are deliberately generic, lengthy, vague, and ambiguous (Reidenberg et al., 2015).

We argue that a central issue is not just *how* privacy communication is designed, but *what* is not being communicated. Privacy communication should inform privacy-related decisions and reduce uncertainty for the person who is deciding whether to use a digital service (Acquisti et al., 2015, 2020). In today's data economy, informed decisions cannot be made effectively based on generic information but rather requires awareness of specific privacy risks in a user's context (Acquisti et al., 2015). Despite omnipresent data breaches, unauthorized sales of personal data, large-scale scraping of data for AI training, emotion detection with AI and now hyper-personalized ads based on LLM inputs (Hammond & Criddle, 2026; Tang et al., 2025), users are rarely made aware of potential harms they face. For example, many users were likely unaware that highly sensitive dialogues shared with conversational AI agents could end up being publicly accessible on the Internet (Morris, 2025). Although many users are sensitive to privacy risks (Pew Research, 2023), they struggle to understand how technologies create such risks and are easily nudged into accepting more than they are comfortable with (Acquisti et al., 2020). Yet, current privacy communication fails to highlight these risks.

The ineffectiveness of privacy communication is rooted in its focus on creating *procedural transparency*: describing data processing procedures and user rights. This conceptualization is firmly entrenched in both historical and contemporary frameworks, including regulatory guidelines (FTC, 2000; OECD, 2001), as well as legislation (California State Legislature, 2018; EU, 2016). For example, the FTC's fair information practice principles require organizations to "be transparent about information policies and practices regarding PII" and demand "clear and accessible notice regarding creation, collection, use, processing, storage, maintenance, dissemination, and disclosure of PII." (FTC, 2000). While procedural transparency is essential for system design and governance, it presupposes a level of privacy expertise that laypeople typically lack to recognize potential risks. Despite the global shift toward risk-based approaches in privacy and AI regulation including mandatory risk assessments (e.g., EU AI Act (EU, 2024)), communication still centers on disclosing data practices towards users rather than making privacy risks salient, thereby inhibiting effective privacy self-management (Acquisti et al., 2015).

In this paper, we advocate for complementing procedural transparency with *risk transparency* to reduce uncertainty for users in privacy decisions. Risk transparency is well established in other domains: for example, in drug safety, the goal is not to comprehensively explain detailed interaction mechanisms of a drug and the human body to ordinary consumers and then let them come to their own conclusion whether health risks are associated with the drug. Instead, major risks such as adverse reactions, drug interactions, and contraindications must be communicated clearly (Pelizzari, 2025). In the remainder of the paper we dissect the fallacy of procedural transparency, motivate the need for risk transparency in information privacy, and outline a path for shifting from procedural to risk transparency.

2. The fallacy of procedural transparency

2.1. Procedural transparency as the prevailing paradigm

Information privacy refers to the collection, use, and sharing of personal data, as well as individuals' control over these processes (Awad & Krishnan, 2006; Smith et al., 2011) (cf. Table 1). Regarding organizations' use of personal data, the current approach to enabling user control over data processing relies on a common assumption: "telling users how their data is processed will help them make informed data sharing decisions." This idea of procedural transparency is reflected in many established regulatory guidelines (FTC, 2000; OECD, 2001) as well as in legislation worldwide, including in the EU 1995 Data Protection Directive (European Parliament and Council, 1995, pp. 31–50), the EU GDPR (EU, 2016), the CCPA (California State Legislature, 2018), the South African POPIA (South African Government, 2020), and the Australian Privacy Act (Australian Government, 2019), as well as many others. These frameworks emphasize principles such as transparency, informed, and openness about data processing practices. For instance, in the EU, the GDPR requires data controllers to provide information about why and how personally identifiable data is collected, processed,

Table 1
Central concepts in this article.

Concept	Definition	Example
Information privacy	The collection, use, and sharing of personal data, as well as individuals' control over these processes (Awad & Krishnan, 2006; Smith et al., 2011). Here, it refers to institutional data processing.	A user allows an app to collect and share location data for personalized ads.
Procedural transparency	Transparency achieved by disclosing personal data processing practices to support informed decision-making (adapted from "Information Transparency" (Turilli & Floridi, 2009)). Negative outcomes are not described.	The app discloses that it collects location data with third-party advertising partners, as required under GDPR Art. 13.
Risk communication	The process of conveying information about risks, uncertainties, and harms in order to reduce uncertainty and support decision-making (Balog-Way et al., 2020; Morgan et al., 2001)	The app explains that third-party location sharing increases the likelihood of misuse or data breaches.
Risk transparency	Transparency achieved by communicating the privacy risks that may result from data practices to support informed decision-making (adapted from"Information Transparency" (Turilli & Floridi, 2009) and "Risk Awareness" (Aven et al., 2018).	A smartphone operating system warns that location sharing may enable tracking or profiling.
Privacy harm	Adverse consequences resulting from privacy violations, such as economic, social, psychological, or reputational harm (Citron & Solove, 2022).	Users may suffer stalking, discrimination, reputational damage, or distress.
Privacy risk	The potential for privacy harm due to a privacy violation characterized by uncertainty (adapted from "Risk" (Aven et al., 2018; Citron & Solove, 2022)).	Location sharing bears the risk of leakage or misuse, revealing sensitive routines.
Privacy violation	A privacy-related event in which established norms or legal obligations of personal data processing are breached (adapted from "Event" (Aven et al., 2018; Nissenbaum, 2011)).	An advertising partner breaches or unlawfully sells users' location data.
Privacy threat	A risk source such as a condition, actor, or data practice that has the potential to lead to privacy harm (adapted from "Threat" (Aven et al., 2018)).	Continuous collection of precise location data and sharing it with multiple advertising partners.

stored, and shared by an organization and its partners (Art. 13–15 GDPR (EU, 2016)). The more recent AI Act, which emphasizes risk-based regulation of AI technologies, also focuses on procedural transparency towards the end user (Art. 50 EU AI Act (EU, 2024)).

Beyond regulatory requirements, procedural transparency is often also implemented voluntarily, for example, by browser and operating systems (OS) vendors. For example, a mobile OS may notify users that an app requests access to their contacts, location, or seeks permission to share data with third parties (Almuhimedi et al., 2015). Users are then expected to evaluate these requests in the context of their app use and decide whether to allow or deny access, hypothesizing what risks can be associated with their decisions.

Since their introduction, documents that describe data practices for users (i.e., procedural transparency) have been criticized (e.g., overly legalistic language, difficult to comprehend (McDonald & Cranor, 2008)). Various attempts have been made to improve procedural transparency (cf. Table 2). For example, some approaches focus on making privacy notices more comprehensible and better structured (Schaub & Cranor, 2024; Schaub et al., 2015), while others use visual tools like privacy icons and privacy nutrition labels, now commonly used in app stores, to represent data practices (Holtz et al., 2011). Some approaches aim to better connect relevant data practices with the user's context, such as location access information (Almuhimedi et al., 2015), or provide technological support for users (Harkous et al., 2018; Pan et al., 2024).

Despite these efforts, privacy notices keep increasing in length, with more and more data practices being described by service providers, and require increasingly advanced reading skills (Amos et al., 2021; Degeling et al., 2019). Still today, data practices stated within privacy notices are often contradictory and important information is absent (Xian et al., 2025). With a growing number of described data practices it has become even harder for users to see what is relevant or what is missing (Adhikari et al., 2023; Amos et al., 2021; Wagner, 2023). Greater descriptive detail also increases cognitive load and privacy fatigue (Agozie & Kaya, 2021). Indeed, deciding whether to use a digital service based on privacy information resembles many everyday decisions bound by cognitive constraints, leading users to make spontaneous decisions rather than optimal ones (Acquisti et al., 2015, 2020; Solove, 2013). As a consequence, privacy policies and other presentations of data practices are often ignored (Obar & Oeldorf-Hirsch, 2020).

More importantly, research indicates that users often struggle to infer downstream, long-term risks even when they engage with the description of data practices (Acquisti et al., 2020; Korunovska et al., 2020). For example, in one study, half of the participants misinterpreted data practices, and many did not understand the risks of secondary data use (Korunovska et al., 2020). Even when data practices were presented with visual support, such as privacy icons, users often misinterpreted them (Habib et al., 2021; R. Sun et al., 2024), leading some scholars to describe these visual simplifications as “privacy empowerment illusions” (R. Sun et al., 2024).

Psychology offers several possible explanations. First, at a basic cognitive level, recognizing risky situations and anticipating negative outcomes depends on prior experience (Klein, 2017). In the absence of such experience, for instance, when users are unfamiliar with data practices and their potential adverse consequences, individuals may be unable to adequately recognize privacy risks. Second, at a higher cognitive level, laypeople's mental models of “how things work” often differ substantially from those of experts (Morgan et al., 2001). For example, a study compared experts' and laypeople's mental models of chemical hazards in the workplace (Cox et al., 2003). Experts reasoned about long-term risks based on abstract, causal representations of handling processes, whereas laypeople relied primarily on their limited personal experience and salient attributes of chemicals (e.g., visible vapors). As a result, long-term risks were weakly represented in laypeople's mental models and resulting handling practices (Cox et al., 2003). In privacy, several studies indicate that ordinary users have more simplistic and sometimes inaccurate mental models of data practices compared to experts (Ebert et al., 2024; Frik et al., 2019; Mayer et al., 2021; K. Sun et al., 2021; Zhang et al., 2024). Additionally, descriptions in privacy policies tend to align only with expert mental models (Bashir et al., 2015). Accordingly, it can be assumed that laypersons are unlikely to accurately infer privacy risks from procedural descriptions, due to their underlying mental models.

Despite the difficulties users face in deriving privacy risks from descriptions of data practices, the law does not mandate disclosure of privacy risks to ordinary users, unlike in other regulatory domains such as drug safety (e.g., Food, Drug, and Cosmetic Act (United States Congress, 1938)). While data breach notifications have been widely implemented (Romanosky et al., 2011), these inform customers only after privacy violations have occurred. When it comes to disclosing risk ex ante, risk communication is limited to narrow stakeholder groups. For example, the GDPR requires organizations to conduct Data Protection Impact Assessments (DPIAs) to

Table 2
Selected efforts to improve procedural transparency for end users.

Approach	Description
Linguistic assistance	
Simplified language	Translate legalese into comprehensible plain language (Jensen & Potts, 2004).
Layered notices	Use visual layers to balance brevity and detail (e.g., App store disclosures)
Visual assistance	
Privacy icons	Use of symbols to represent key privacy aspects (e.g., App store icons (Holtz et al., 2011)).
Privacy nutrition labels	Present data practices in a standardized, tabular form (e.g., iOS app labels (Kelley et al., 2009)).
Privacy seals	Visual cues to indicate privacy protection level (e.g., TRUSTe label (Rifon et al., 2005)).
Privacy ratings	Rate services based on privacy friendliness (e.g., tools using privacy scores (Tsai et al., 2011)).
Contextual & system assistance	
Just-in-time notices	Display context-sensitive notices during usage (e.g., location access (Almuhimedi et al., 2015)).
Contextual privacy policies	Align contexts with privacy policy segments (e.g., SeePrivacy (Pan et al., 2024)).
Machine-readable notices	Provide information in a machine-readable format (e.g., P3P (Reagle and Cranor, 1999)).

Table 3
Illustrative example of procedural transparency vs. risk transparency for privacy.

Scenario	Procedural transparency	Risk transparency
Website informs about third party data sharing	Your personal data may be shared with selected third party partners and service providers to help us deliver, maintain, and improve our services, in accordance with our privacy policy.	Your data is shared with 173 vendors that help us deliver, maintain, and improve our services. While we carefully vet our vendors, there is a risk that vendors may accidentally leak information or suffer a data breach, which could lead to spam, fraud, or identity misuse (informative example).
Mobile OS informs that a dating app wants to access the user's location	This app requests access to your precise location.	**Warning:** This app may expose you to a serious privacy risk! The app has sold precise location data in the past, making it likely that your location data is sold to data brokers. Once disclosed, this data cannot be controlled and may enable stalking, discrimination, or surveillance (fear appeal example).

identify and mitigate risks in projects that are likely to result in a high risk to users. However, the identified risks are not communicated to users (Goncalves, 2020); they are addressed only internally or, if high risks remain, reported to supervisory authorities (Art. 35 GDPR (EU, 2016)). The GDPR mandates transparency regarding the "significance and envisaged consequences" of processing only in the limited context of automated decision-making and profiling (Art. 13–15 GDPR (EU, 2016)).

2.2. The need for risk transparency

Scholars have argued that it is crucial for individuals to understand risks in decision-making due to (1) normative, (2) instrumental, and (3) substantive reasons (Balog-Way et al., 2020). First, it is ethically essential to communicate known or anticipated risks to allow a person to act autonomously and informed instead of being confronted with diffuse concerns (Covello et al., 1989). This idea is reflected in concepts such as informed consent, freedom of information, or public-right-to-know. The concepts can be found in privacy laws, too, which demand protection, for example, from "the risk to the rights and freedoms of natural persons" (Recital 72ff. of the GDPR (EU, 2016)). Second, for individuals to achieve privacy protection goals, they need to first understand the potential harms (Citron & Solove, 2022). Finally, understanding privacy risks and resulting harms can substantively improve the quality of privacy-related decisions. We illustrate these three points using a weather app: users should be informed that the app sells location data and understand the associated risks (1). They should recognize possible harms, such as profiling based on location data, so they can pursue their individual privacy preferences on their smartphone (2). Finally, they should be able to decide whether to use the app, switch to another app, or restrict location sharing in their OS in line with their privacy preferences (3).

3. From procedural to risk transparency for privacy

Highlighting privacy risks ("What could happen?") instead of merely describing data practices ("What is done?") requires a reconceptualization of what transparency means in the privacy context. In the following, we draw on established risk research and privacy literature to conceptualize key notions of privacy risk transparency (cf. Table 1), examine how privacy risks can be communicated, and discuss challenges of privacy risk communication.

3.1. Defining privacy risks

The Society for Risk Analysis characterizes risk as the potential for adverse consequences of an event under uncertainty (Aven et al., 2018). In the context of privacy, risks can therefore be understood as the potential for privacy harm to occur due to a privacy violation (Aven et al., 2018; Citron & Solove, 2022). Sources of risk are referred to as privacy threats, such as conditions, actors, or data practices (e.g., the massive collection of personal data) (Aven et al., 2018).

A privacy violation can be defined as an event in which established norms or legal obligations are breached (e.g., the event of a data breach) (Aven et al., 2018; Nissenbaum, 2011). Privacy harms range from physical, economic, and psychological damage to manipulation and discrimination (Citron & Solove, 2022). Some forms are subjective (e.g., anxiety), while others are more objective and measurable (e.g., job loss due to information disclosure) (Citron & Solove, 2022).

The privacy harms resulting from a privacy violation may not manifest immediately or be fully understood initially (Mayer et al., 2021; Solove, 2013), which can lead to the perception of privacy risks as abstract or diffuse (Acquisti et al., 2015). For instance, when personal data is shared or leaked, willingly or not, by organizations, this may remain unnoticed for a long time while the data is actively being used or misused. Data breaches, for example, have often been detected only after several months or longer (Mayer et al., 2021; Roumani, 2022). Unnoticed privacy violation may also create new privacy risks potentially leading to new violations. For instance, data breaches may increase the likelihood that affected individuals will become victims of further cybercrime, e.g., when leaked passwords are used to compromise additional accounts. Subsequent identity theft as well as the misuse of stolen credit card details, social security numbers, or other sensitive information repeatedly result in tangible financial harm (Ruohonen et al., 2024). Also, a seemingly identical privacy violation may create very different risks based on the context (Nissenbaum, 2011): A leaked GPS coordinate, for instance, poses vastly different risks depending on whether it identifies a tourist at a café, a priest in a gay bar, or a soldier in a conflict zone.

3.2. *Communicating privacy risks*

Risk communication is a well-established approach in domains such as public health, environmental science, consumer safety, and disaster management. Its central purpose is to reduce uncertainty in ways that support informed decision-making by highlighting concrete risks (e.g., diseases associated with smoking) (Balog-Way et al., 2020). Depending on context, the goals of risk communication range from raising awareness about potential threats to influencing attitudes or behavior (Balog-Way et al., 2020). To achieve these goals, risk communication employs various approaches. These include the presentation of statistical information in an understandable form, as well as the use of personal experiences, narratives, and anecdotes that help individuals imagine potential consequences (Balog-Way et al., 2020).

The communication of risks can focus on an individual's analytical decision-making processes that assess risks and benefits (e.g., a privacy calculus (Dinev & Hart, 2006)) by employing education and factual information (Morgan et al., 2001). As individuals often evaluate risks and benefits based on affective responses rather than on objective facts alone (Lerner et al., 2015; Slovic et al., 2007), risk communication can also deliberately target automatic, affect-driven mental processes, for instance, through intentionally fear-inducing messages (i.e., fear appeals) (Bigsby & Albarracín, 2022).

Numerous meta-analyses have examined the effectiveness of risk communication, especially in the fear appeal and warning literature. The current consensus is that risk communication can be effective in raising risk awareness, however, it must be complemented by additional mechanisms, for example, those that strengthen the users' ability to cope with threats (i.e., increase efficacy) (Albarracín et al., 2024; Bigsby & Albarracín, 2022; Hancock et al., 2020; Kok et al., 2018). For example, meta-analyses in health communication research indicate that risk perceptions can be increased by providing risk-related information in an understandable form (Bakhit et al., 2024; Zipkin et al., 2014). Regarding fear appeals, several meta-analyses have found overall positive effects on attitudes, intentions, and behaviors, with effect sizes ranging from small to moderate (Boster & Mongeau, 1984; Earl & Albarracín, 2007; Floyd et al., 2000; Sheeran et al., 2014; Tannenbaum et al., 2015; Witte & Allen, 2000). At the same time, one meta-analysis reported null effects of fear appeals (Peters et al., 2013). Criticism of fear appeals has focused on potential boomerang effects (Kok et al., 2018), ethical concerns (Hastings et al., 2004; Kok et al., 2018), and a limited likelihood of producing long-term behavioral change (Ruiter et al., 2014). Comparable meta-analyses in safety science have examined the effectiveness of warning labels on products such as cigarette packs or chemicals. They conclude that warnings can increase safety attitudes and behavioral compliance, provided they are concise, salient, and embedded in the user's task context (Argo & Main, 2004; Cox III et al., 1997; Hancock et al., 2020; Noar et al., 2016). For example, a meta-analysis related to the effectiveness of cigarette pack warnings suggests that pictorial warnings are more effective than textual warnings (Noar et al., 2016).

In the context of privacy, comparatively fewer empirical studies have investigated risk communication, and meta-analyses dedicated to risk communication have not yet emerged. However, a meta-analysis that covered various forms of privacy nudges found that eleven out of thirteen studies on risk communication reported positive effects (Ioannou et al., 2021), while two did not (Junger et al., 2017; Meier et al., 2020). Also, more recent studies have reported positive effects of risk communication on risk perception, protective intentions, and protective behavior (Ebert et al., 2021, 2023; Franzen et al., 2022; Massara et al., 2021; Momenzadeh et al., 2020; Saeidi et al., 2022; Shulman et al., 2023; Zou et al., 2024). Two factors are particularly noteworthy for risk communication in privacy: personal relevance and habituation. First, prior research shows that emphasizing concrete personal consequences of data practices can improve risk perception and protective behavior (Almuhimedi et al., 2015; Ebert et al., 2024; Emami-Naeini et al., 2021; Zou et al., 2024). As discussed earlier, this directly addresses users' difficulties in inferring downstream risks from otherwise abstract descriptions. Second, prior work has indicated that users quickly habituate when repeatedly confronted with excessive or irrelevant security and privacy warnings, thereby reducing their effectiveness (Karegar et al., 2020; Vance et al., 2018). A substantial proportion of the aforementioned privacy studies have been conducted in laboratory settings (e.g., Ebert et al., 2021, 2023), which limits their generalizability. This underscores the need to better understand the determinants of effective risk communication in privacy.

In summary, the evidence suggests that risk communication in the domain of privacy can indeed be effective; however, to increase its impact and to prevent boomerang effects (e.g., resignation, resistance) it should be complemented by additional mechanisms, for example by providing readily available and actionable privacy-friendly options to cope with risks (Albarracín et al., 2024; Bigsby & Albarracín, 2022). Evidence also suggests that risk communication in privacy requires some form of prioritization to increase personal relevance and avoid habituation. Rather than communicating all possible risks, communication should either focus on those privacy risks that are most severe and relevant for groups of individuals, or be strongly personalized to reflect individual risk preferences and risk levels (e.g., using personal privacy risk profiles (Novikova et al., 2025)).

Illustrative example. The two scenarios shown in Table 3 better illustrate risk communication in privacy. For illustrative purposes, the examples are deliberately more elaborate than what would typically be appropriate in practice. They also do not describe coping mechanisms (e.g., access to a more privacy-friendly alternative), which are crucial in a practical implementation (Bigsby & Albarracín, 2022). The procedural transparency examples in the second column reflect current practice. In the first scenario, the risk transparency approach informs users about a privacy risk, for instance, because the provider is legally required to do so. In contrast, the second scenario uses a fear appeal to heighten emotional arousal: rather than merely stating that data is accessed, the OS warns the user. In theory, risks in both scenarios could also be quantified, for example, in the same way as health risks ("There is a 1 in 10,000 chance your location data is made public within the next 12 months due to a data breach"). However, in practice, the necessary data is hardly available. Yet, communicating what is not known (i.e., uncertainty) can also be part of risk communication (Balog-Way et al., 2020). For example, even the statement that it is "uncertain what happens if a dating app shares location data with 200 partners" provides more information than merely stating that the app shares data with 200 partners.

4. Reconceptualizing transparency in information privacy

Our article has made two contributions: first, it challenges the prevailing focus on procedural transparency in privacy notice requirements and introduces privacy risk transparency as a necessary complement. Although privacy scholarship and regulation increasingly promote structural safeguards, including risk-based regulatory approaches (California State Legislature, 2018; EU, 2016; Nissenbaum, 2011; Solove, 2013), we argue that privacy self-management remains indispensable to preserve individual autonomy. However, privacy self-management requires meaningful direct communication with users to be effective. To date, research has primarily sought to improve how information about data practices is presented and understood, focusing on accessibility and design (Cranor et al., 2024; Schaub et al., 2017). We argue that equal attention must be paid to whether users receive conceptually adequate information about privacy risks. Second, the article proposes and illustrates a vocabulary for privacy risk communication grounded in established risk research and presents evidence on the effectiveness of such communication in other domains, as well as in the privacy context.

However, risk communication to users is not a panacea. Ill-designed risk disclosures may exacerbate privacy fatigue, the belief that privacy protection is meaningless (Choi et al., 2018), and provoke reactance, which manifests as backlash against perceived threats to freedom (Bigsby & Albarracín, 2022). Prior research also documents instances in which risk communication failed to produce sustained effects (Li et al., 2016). Research on behavior change further suggests that the strongest determinant of privacy-protective behavior is not knowledge about risks, but easy access to privacy-friendly alternatives (Albarracín et al., 2024). Still, we argue that providing concise and accurate privacy risk information or warnings constitutes an important transparency cornerstone for laypeople. From a normative perspective, such transparency is required, and empirical evidence suggests that it can produce positive behavioral effects (Balog-Way et al., 2020; Bigsby & Albarracín, 2022). Ideally, risk communication may better address diffuse privacy concerns than traditional privacy policies and therefore reduce privacy fatigue.

Important questions remain, however, regarding how risk communication can be implemented in privacy (cf. Table 4). Both risk analysis and risk communication warrant further investigation and active research. For example, would it be possible to estimate probabilistic privacy risks, and can such risks be meaningfully communicated to laypeople (Gigerenzer et al., 2007)? As we have discussed, current evidence remains limited to isolated studies. The long-term effectiveness of different risk communication strategies, as well as the specific determinants shaping their impact, remains insufficiently understood (Ioannou et al., 2021). Insights from related domains, such as information security, may provide a useful starting point (e.g., habituation effects (Vance et al., 2018), visual saliency (Ebert et al., 2023)). At the same time, risk research and safety science offer a rich and well-developed body of theoretical and empirical knowledge that can inform and guide future work in privacy risk communication.

Even more importantly, effective mechanisms to institutionalize risk disclosure to users are still lacking. Each potential mechanism entails distinct advantages and limitations that require careful examination (e.g., practicality, liability). To encourage voluntary disclosure, existing recommendations such as the FIPPs (FTC, 2000) could be updated to explicitly incorporate risk transparency. Market-based mechanisms, such as reduced cyber insurance premiums, could also serve as incentives for voluntary privacy risk disclosure (Khalili et al., 2017).

Table 4
Open questions for advancing privacy risk transparency.

Domain	Questions
Communicating privacy risks	What are effective risk communication strategies that produce sustained effects on risk perception and protective behavior? How can privacy risks (including probabilistic risks) be translated into formats understandable to lay users and other non-privacy professionals such as employees? What data can support the assessment of privacy risks (e.g., probabilities for privacy violation)? How can privacy risks of emerging technologies such as LLMs be assessed and communicated?
Institutionalizing privacy risk transparency	What are the advantages and disadvantages of different forms of institutionalization (e.g., voluntary disclosure, mandatory disclosure)? What incentives could encourage voluntary disclosure of privacy risks? What liability implications arise from privacy risk communication? How could privacy risk communication be embedded in existing regulatory guidelines and legal frameworks? How could privacy risks be assessed and verified through independent audits?

Alternatively, mandatory disclosure, which is common in other regulatory domains (e.g., for food and drug safety (United States Congress, 1938)), could be considered. If mandatory privacy risk disclosure were enforced through legislation, similar to existing data breach notification requirements (e.g., Art. 33 GDPR (EU, 2016)), privacy law and regulators' guidance on how to implement it would need to be amended accordingly. Specifically, these frameworks would need to require risk transparency not only toward internal specialists or supervisory authorities but also toward users of digital services (Gonçalves, 2020). Since many laws and regulatory guidelines already require privacy risk assessments (e.g., Art. 34 DSA (EU, 2022), Art. 35 GDPR (EU, 2016)) or encourage them (e.g., the NIST Privacy Framework), scholars have proposed making these assessments publicly accessible (Iwaya et al., 2024; Nas, 2019). Privacy risk disclosure could also be carried out by independent third parties, such as NGOs or supervisory authorities, that analyze services and certify or audit organizations. For example, NGOs have developed privacy risk assessments for the public that evaluate selected widely-used services for privacy risks (Mozilla, 2023). Similar mechanisms could be implemented by supervisory authorities. For instance, New York City publicly grades restaurant hygiene using a standardized A/B/C system (NYC DOHMH, 2010). A comparable model based on mandatory independent audits or certifications by public authorities could likewise be envisioned in the privacy domain, particularly given that voluntary certification mechanisms are already provided for in some legislation (e.g., Art. 42–43 GDPR (EU, 2016)).

While challenges and open questions remain, risk transparency in privacy is a necessary evolution of privacy transparency requirements towards a digital world where privacy notices can actually fulfill their intended purpose of informing individuals' decisions about their privacy rather than being a nuisance to be clicked away without consideration.

CRediT authorship contribution statement

Nico Ebert: Writing – original draft, Conceptualization. **Simone Fischer-Hübner:** Writing – review & editing. **Soheil Human:** Writing – review & editing. **Agnieszka Kitkowska:** Writing – review & editing. **Konrad Kollnig:** Writing – review & editing. **Jelena Mitrović:** Writing – review & editing. **Shidong Pan:** Writing – review & editing. **Thierry Schaltegger:** Writing – review & editing. **Florian Schaub:** Writing – review & editing. **Daniel Smullen:** Writing – review & editing. **Lu Xian:** Writing – review & editing.

Declaration of competing interest

This research was funded by the Digitalization Initiative of the Zurich Higher Education Institutions (DIZH), the Swiss National Sciences Foundation (Grant 207550) and by the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation.

Acknowledgements

The idea for this article originated at the Dagstuhl Seminar 25021 (Grand Challenges for Research on Privacy Documents) in early 2025. We would like to thank all participants for the stimulating discussions. Furthermore, we would like to thank Angela Bearth for sparking our interest in risk communication and sharing her extensive knowledge with us over the past few years. We would also like to thank the reviewers for their feedback, which has been instrumental in improving the article.

Data availability

No data was used for the research described in the article.

References

Acquisti, A., Brandimarte, L., & Loewenstein, G. (2015). Privacy and human behavior in the age of information. *Science, 347*(6221), 509–514.

Acquisti, A., Brandimarte, L., & Loewenstein, G. (2020). Secrets and likes: The drive for privacy and the difficulty of achieving it in the digital age. *Journal of Consumer Psychology, 30*(4), 736–758. https://doi.org/10.1002/jcpy.1191

Adhikari, A., Das, S., & Dewri, R. (2023). Evolution of composition, readability, and structure of privacy policies over two decades. *Proceedings on Privacy Enhancing Technologies (PoPETs), 2023*(3), 138–153. https://doi.org/10.56553/popets-2023-0074

Agozie, D. Q., & Kaya, T. (2021). Discerning the effect of privacy information transparency on privacy fatigue in e-government. *Government Information Quarterly, 38*(4), Article 101601. https://doi.org/10.1016/j.giq.2021.101601

Albarracín, D., Fayaz-Farkhad, B., & Samayoa, J. A. G. (2024). Determinants of behaviour and their efficacy as targets of behavioural change interventions. *Nature Reviews Psychology, 3*(6), 377–392. https://doi.org/10.1038/s44159-024-00305-0

Almuhimedi, H., Schaub, F., Sadeh, N., Adjerid, I., Acquisti, A., Gluck, J., Cranor, L. F., & Agarwal, Y. (2015). Your location has been shared 5,398 times! a field study on mobile app privacy nudging. In *Proceedings of the 33rd annual ACM conference on human factors in computing systems* (pp. 787–796). https://doi.org/10.1145/2702123.2702210

Amos, R., Acar, G., Lucherini, E., Kshirsagar, M., Narayanan, A., & Mayer, J. (2021). Privacy policies over time: Curation and analysis of a million-document dataset. In *Proceedings of the web conference 2021* (pp. 2165–2176). https://doi.org/10.1145/3442381.3450048

Argo, J. J., & Main, K. J. (2004). Meta-analyses of the effectiveness of warning labels. *Journal of Public Policy and Marketing, 23*(2), 193–208.

Australian Government. (2019). *Chapter 1: App 1 open and transparent management of personal information*.

Aven, T., Ben-Haim, Y., Andersen, H. B., Cox, T., Droguett, E. L., Greenberg, M., Guikema, S. D., Kröger, W., Renn, O., Thompson, K. M., et al. (2018). Society for risk analysis glossary. https://www.sra.org/risk-analysis-introduction/risk-analysis-glossary/.

Awad, N. F., & Krishnan, M. S. (2006). The Personalization Privacy Paradox: An Empirical Evaluation of Information Transparency and the Willingness to be Profiled Online for Personalization1 [Publisher: MIS Quarterly]. *MIS Quarterly, 30*(1), 13–28. https://doi.org/10.2307/25148715

Bakhit, M., Fien, S., Abukmail, E., Jones, M., Clark, J., Scott, A. M., Glasziou, P., & Cardona, M. (2024). Cardiovascular disease risk communication and prevention: A meta- analysis. *European Heart Journal, 45*(12), 998–1013. https://doi.org/10.1093/eurheartj/ehae002

Balog-Way, D., McComas, K., & Besley, J. (2020). The evolving field of risk communication. *Risk Analysis, 40*(S1), 2240–2262.
Bashir, M., Hayes, C., Lambert, A. D., & Kesan, J. P. (2015). Online privacy and informed consent: The dilemma of information asymmetry. *Proceedings of the Association for Information Science and Technology, 52*(1), 1–10.
Bigsby, E., & Albarracín, D. (2022). Self- and response efficacy information in fear appeals: A meta-analysis. *Journal of Communication, 72*(2), 241–263. https://doi.org/10.1093/joc/jqab048
Boster, F. J., & Mongeau, P. (1984). Fear-arousing persuasive messages. *Communication Yearbook, 8*(1), 330–375. https://doi.org/10.1080/23808985.1984.11678581
California State Legislature. (2018). *California consumer Privacy act (CCPA)*. oag.ca.gov. . (Accessed 8 April 2025).
Choi, H., Park, J., & Jung, Y. (2018). The role of privacy fatigue in online privacy behavior. *Computers in Human Behavior, 81*, 42–51.
Citron, D. K., & Solove, D. J. (2022). Privacy harms. *BUL Rev, 102*, 793.
Covello, V. T., McCallum, D. B., & Pavlova, M. T. (Eds.). (1989). *Effective risk communication: The role and responsibility of government and nongovernment organizations*. Springer US. https://doi.org/10.1007/978-1-4613-1569-8.
Cox III, E. P., Wogalter, M. S., Stokes, S. L., & Murff, E. J. T. (1997). Do product warnings increase safe behavior? A meta-analysis. *Journal of Public Policy and Marketing, 16*(2), 195–204.
Cox, P., Niewöhner, J., Pidgeon, N., Gerrard, S., Fischhoff, B., & Riley, D. (2003). The use of mental models in chemical risk protection: Developing a generic workplace methodol-ogy. *Risk Analysis: International Journal, 23*(2), 311–324.
Cranor, L. F., Agarwal, Y., & Emami-Naeini, P. (2024). Internet of things security and privacy labels should empower consumers. *Communications of the ACM, 67*(3), 29–31. https://doi.org/10.1145/3637630
Degeling, M., Utz, C., Lentzsch, C., Hosseini, H., Schaub, F., & Holz, T. (2019). We value your privacy. now take some cookies: Measuring the gdpr's impact on web privacy. In *Network and distributed systems security (NDSS) symposium 2019*. https://doi.org/10.14722/ndss.2019.23378
Dinev, T., & Hart, P. (2006). An extended privacy calculus model for e-commerce transactions. *Information Systems Research, 17*(1), 61–80.
Earl, A., & Albarracín, D. (2007). Nature, decay, and spiraling of the effects of fear-inducing arguments and HIV counseling and testing: A meta-analysis of the short- and long-term outcomes of HIV-prevention interventions [Place: US Publisher: American Psychological Association]. *Health Psychology, 26*(4), 496–506. https://doi.org/10.1037/0278-6133.26.4.496
Ebert, N., Ackermann, K. A., & Bearth, A. (2023). When information security depends on font size: How the saliency of warnings affects protection behavior. *Journal of Risk Research, 26*(3), 233–255. https://doi.org/10.1080/13669877.2022.2142952
Ebert, N., Alexander Ackermann, K., & Scheppler, B. (2021). Bolder is better: Raising user awareness through salient and concise privacy notices. In *Proceedings of the 2021 CHI conference on human factors in computing systems*. https://doi.org/10.1145/3411764.3445516
Ebert, N., Geppert, T., Knieps, M., Zarouali, B., Schaltegger, T., Wiedemann, A., & Ambuehl, B. (2024). Reflective data sharing on tiktok: Encouraging adolescents to engage with privacy settings [paper 5]. In *Proceedings of the 32nd European conference on information systems*. ECIS 2024). https://aisel.aisnet.org/ecis2024/track24socialmedia/track24socialmedia/5.
Emami-Naeini, P., Dheenadhayalan, J., Agarwal, Y., & Cranor, L. F. (2021). Which privacy and security attributes most impact consumers' risk perception and willingness to purchase iot devices?. In *2021 IEEE symposium on security and privacy (SP)* (pp. 519–536). https://doi.org/10.1109/SP40001.2021.00112
Emami-Naeini, P., Dheenadhayalan, J., Agarwal, Y., & Cranor, L. F. (2022). An informative security and privacy "nutrition" label for internet of things devices. *IEEE Security & Privacy, 20*(2), 31–39. https://doi.org/10.1109/MSEC.2021.3132398
EU. (2016). *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)]. from https://data.europa.eu/eli/reg/2016/679/oj. (Accessed 13 April 2023).
EU. (2022). Regulation (EU) 2022/2065 of the European parliament and of the council of 19 October 2022 on a single market for digital services and amending directive 2000/31/EC (digital services act). https://eur-lex.europa.eu/eli/reg/2022/2065/oj.
EU. (2024). *Regulation (Eu) 2024/1689 of the European parliament and of the council of 13 June 2024 laying down harmonised rules on artificial intelligence and amending certain union legislative acts (artificial intelligence act)* [See Article 3, point 8].
European Parliament and Council. (1995). *Directive 95/46/ec of the European parliament and of the council of 24 October 1995 on the protection of individuals with regard to the processing of personal data and on the free movement of such data* [Official Journal L 281, 23/11/1995 . (Accessed 4 June 2025).
Floyd, D. L., Prentice-Dunn, S., & Rogers, R. W. (2000). A meta-analysis of research on protection motivation theory. *Journal of Applied Social Psychology, 30*(2), 407–429. https://doi.org/10.1111/j.1559-1816.2000.tb02323.x
Franzen, D., Nuñez von Voigt, S., Sörries, P., Tschorsch, F., & Müller-Birn, C. (2022). Am i private and if so, how many? Communicating privacy guarantees of differential privacy with risk communication formats. In *Proceedings of the 2022 ACM SIGSAC conference on computer and communications security* (pp. 1125–1139).
Freiberger, V., Fleig, A., & Buchmann, E. (2025). You don't need a university degree to comprehend data protection this way": Llm-powered interactive privacy policy assessment. In *Proceedings of the extended abstracts of the CHI conference on human factors in computing systems*. https://doi.org/10.1145/3706599.3719816
Frik, A., Nurgalieva, L., Bernd, J., Lee, J., Schaub, F., & Egelman, S. (2019). Privacy and security threat models and mitigation strategies of older adults. In *Fifteenth symposium on Usable privacy and security (SOUPS 2019)* (pp. 21–40). https://www.usenix.org/conference/soups2019/presentation/frik.
FTC. (2000). *Privacy online: Fair information practices in the electronic marketplace (report to congress) (contains discussion of the federal trade Commission's fair information practice principles (FIPPs))*. Federal Trade Commission. https://www.ftc.gov/sites/default/files/documents/reports/privacy-online-fair-information-practices-electronic-marketplace-federal-trade-commission-report/privacy2000.pdf.
Gigerenzer, G., Gaissmaier, W., Kurz-Milcke, E., Schwartz, L. M., & Woloshin, S. (2007). Helping doctors and patients make sense of health statistics [PMID: 26161749]. *Psychological Science in the Public Interest, 8*(2), 53–96. https://doi.org/10.1111/j.1539-6053.2008.00033.x
Gonçalves, M. E. (2020). The risk-based approach under the new eu data protection regulation: A critical perspective. *Journal of Risk Research, 23*(2), 139–152.
Habib, H., Zou, Y., Yao, Y., Acquisti, A., Cranor, L., Reidenberg, J., Sadeh, N., & Schaub, F. (2021). Toggles, dollar signs, and triangles: How to (in)effectively convey privacy choices with icons and link texts. In *Proceedings of the 2021 CHI conference on human factors in computing systems*. https://doi.org/10.1145/3411764.3445387
Hammond, G., & Criddle, C. (2026). *OpenAI brings advertising to ChatGPT in push for new revenue*. Financial Times. Retrieved February 5, 2026, from https://www.ft.com/content/ec1656cd-e07b-48ed-92a8-26c7fe517899.
Hancock, P., Kaplan, A., MacArthur, K., & Szalma, J. (2020). How effective are warnings? A meta-analysis. *Safety Science, 130*, Article 104876. https://doi.org/10.1016/j.ssci.2020.104876
Harkous, H., Fawaz, K., Lebret, R., Schaub, F., Shin, K. G., & Aberer, K. (2018). Polisis: Automated analysis and presentation of privacy policies using deep learning. In *27th USENIX security symposium (USENIX security 18)* (pp. 531–548).
Hastings, G., Stead, M., & Webb, J. (2004). Fear appeals in social marketing: Strategic and ethical reasons for concern. *Psychology and Marketing, 21*(11), 961–986. https://doi.org/10.1002/mar.20043
Holtz, L.-E., Nocun, K., & Hansen, M. (2011). Towards displaying privacy information with icons. In S. Fischer-Hübner, P. Duquenoy, M. Hansen, R. Leenes, & G. Zhang (Eds.), *Privacy and identity management for life* (pp. 338–348). Springer Berlin Heidelberg.
Ioannou, A., Tussyadiah, I., Miller, G., Li, S., & Weick, M. (2021). Privacy nudges for disclosure of personal information: A systematic literature review and meta-analysis. *PLoS One*, *16*(8), Article e0256822. 10.1371/journal.pone.0256822.
Iwaya, L.H., Alaqra, A.S., Hansen, M., & Fischer-Hübner, S. (2024). Privacy impact assessments in the wild: A scoping review. *Array*, *23*, Article 100356.
Jensen, C., & Potts, C. (2004). Privacy policies as decision-making tools: An evaluation of online privacy notices. In *Proceedings of the SIGCHI conference on human factors in computing systems* (pp. 471–478).
Junger, M., Montoya, L., & Overink, F. (2017). Priming and warnings are not effective to prevent social engineering attacks. *Computers in Human Behavior, 66*, 75–87. https://doi.org/10.1016/j.chb.2016.09.012
Karegar, F., Pettersson, J. S., & Fischer-Hübner, S. (2020). The dilemma of user engagement in privacy notices: Effects of interaction modes and habituation on user attention. *ACM Trans. Priv. Secur., 23*(1). https://doi.org/10.1145/3372296

Kelley, P. G., Bresee, J., Cranor, L. F., & Reeder, R. W. (2009). A" nutrition label" for privacy. In *Proceedings of the 5th symposium on useable privacy and security* (pp. 1–12).
Khalili, M. M., Naghizadeh, P., & Liu, M. (2017). Designing cyber insurance policies: Mitigating moral hazard through security pre-screening. In L. Duan, A. Sanjab, H. Li, X. Chen, D. Materassi, & R. Elazouzi (Eds.), *Game theory for networks* (pp. 63–73). Springer International Publishing.
Klein, G. A. (2017). *Sources of power: How people make decisions*. MIT press.
Kok, G., Peters, G.-J. Y., Kessels, L. T. E., ten Hoor, G. A., & Ruiter, R. A. C. (2018). Ignoring theory and misinterpreting evidence: The false belief in fear appeals [PMID: 29233060]. *Health Psychology Review, 12*(2), 111–125. https://doi.org/10.1080/17437199.2017.1415767
Korunovska, J., Kamleitner, B., & Spiekermann-Hoff, S. (2020). The challenges and impact of privacy policy comprehension. In *A. for information systems* (pp. 1–17). *Twenty-eighth* european conference on information systems (ecis 2020) https://aisel.aisnet.org/ecis2020rp/51/.
Lerner, J. S., Li, Y., Valdesolo, P., & Kassam, K. S. (2015). Emotion and decision making. *Annual Review of Psychology, 66*(1), 799–823.
Li, S. X., Ye, Z., Whelan, K., & Truby, H. (2016). The effect of communicating the genetic risk of cardiometabolic disorders on motivation and actual engagement in preventative lifestyle modification and clinical outcome: A systematic review and meta-analysis of randomised controlled trials. *The British Journal of Nutrition, 116*(5), 924–934. https://doi.org/10.1017/S0007114516002488
Massara, F., Raggiotto, F., & Voss, W. G. (2021). Unpacking the privacy paradox of consumers: A psychological perspective. *Psychology and Marketing, 38*(10), 1814–1827. https://doi.org/10.1002/mar.21524
Mayer, P., Zou, Y., Schaub, F., & Aviv, A. J. (2021). Now i'm a bit angry:" individuals' awareness, perception, and responses to data breaches that affected them. In *30th USENIX Security symposium (USENIX security 21)* (pp. 393–410). https://www.usenix.org/conference/usenixsecurity21/presentation/mayer.
McDonald, A. M., & Cranor, L. F. (2008). The cost of reading privacy policies. *I/S: A Journal of Law and Policy for the Information Society, 4*, 543.
Meier, Y., Schäwel, J., Kyewski, E., & Krämer, N. C. (2020). Applying protection motivation theory to predict Facebook users' withdrawal and disclosure intentions. In *International conference on social media and society* (pp. 21–29). https://doi.org/10.1145/3400806.3400810
Momenzadeh, B., Gopavaram, S., Das, S., & Camp, L. J. (2020). Bayesian evaluation of user app choices in the presence of risk communication on android devices. In N. Clarke, & S. Furnell (Eds.), *Human aspects of information security and assurance* (pp. 211–223). Springer International Publishing. https://doi.org/10.1007/978-3-030-57404-816.
Morgan, M. G., Fischhoff, B., Bostrom, A., & Atman, C. J. (2001). *Risk communication: A mental models approach* (1st ed.). Cambridge University Press. https://doi.org/10.1017/CBO9780511814679
Morris, C. (2025). OpenAI does away with feature that made ChatGPT conversations discoverable by google. Retrieved February 9, 2026, from https://fortune.com/2025/08/05/openai-google-search-chat-history/.
Mozilla. (2023). *Privacy not included: A buyer's guide for connected products. Retrieved February 9, 2026, from https://www.mozillafoundation.org/en/privacynotincluded/articles/annual-creep-o-meter/.
Nas, S. (2019). Data protection impact assessment: Assessing the risks of using microsoft office proplus. *Eur. Data Prot. L. Rev., 5*, 107.
Nissenbaum, H. (2011). A contextual approach to privacy online. *Dædalus, 140*(4), 32–48.
Noar, S. M., Hall, M. G., Francis, D. B., Ribisl, K. M., Pepper, J. K., & Brewer, N. T. (2016). Pictorial cigarette pack warnings: A meta-analysis of experimental studies [publisher: BMJ publishing group ltd section: Review]. *Tobacco Control, 25*(3), 341–354. https://doi.org/10.1136/tobaccocontrol-2014-051978
Novikova, E., Doynikova, E., & Kotenko, I. (2025). What are your privacy risks? Privacy risk assessment based on privacy policies analysis. *Expert Systems with Applications, 280*, Article 127270. https://doi.org/10.1016/j.eswa.2025.127270
NYC DOHMH. (2010). *Letter grading for restaurants*.
Obar, J. A., & Oeldorf-Hirsch, A. (2020). The biggest lie on the internet: Ignoring the privacy policies and terms of service policies of social networking services. *Information, Communication & Society, 23*(1), 128–147. https://doi.org/10.1080/1369118X.2018.1486870
OECD. (2001). *Oecd guidelines on the protection of privacy and transborder flows of personal data* (originally adopted in 1980, revised in 2001). Paris https://www.oecd.org/en/publications/oecd-guidelines-on-the-protection-of-privacy-and-transborder-flows-of-personal-data9789264196391-en.html.
Pan, S., Tao, Z., Hoang, T., Zhang, D., Li, T., Xing, Z., Xu, X., Staples, M., Rakotoarivelo, T., & Lo, D. (2024). A new hope: Contextual privacy policies for mobile applications and an approach toward automated generation. In *33rd USENIX security symposium (USENIX security 24)* (pp. 5699–5716).
Pelizzari, N. (2025). Changing roles of patient information leaflets in the UK: A corpus-assisted discourse analysis. *Applied Corpus Linguistics, 5*(2), Article 100129. https://doi.org/10.1016/j.acorp.2025.100129
Peters, G.-J. Y., Ruiter, R. A., & Kok, G. (2013). Threatening communication: A critical re-analysis and a revised meta-analytic test of fear appeal theory. *Health Psychology Review, 7*(sup1), S8–S31.
Pew Research. (2023). *How Americans view data privacy*. Pew Research Center. https://www.pewresearch.org/internet/2023/10/18/how-americans-view-data-privacy/. (Accessed 30 May 2025).
Reagle, J., & Cranor, L. F. (1999). The platform for privacy preferences. *Communications of the ACM, 42*(2), 48–55.
Reidenberg, J. R., Breaux, T., Cranor, L. F., French, B., Grannis, A., Graves, J. T., Liu, F., McDonald, A., Norton, T. B., Ramanath, R., et al. (2015). Disagreeable privacy policies: Mismatches between meaning and users' understanding. *Berkeley Tech. LJ, 30*, 39.
Rifon, N. J., LaRose, R., & Choi, S. M. (2005). Your privacy is sealed: Effects of web privacy seals on trust and personal disclosures. *Journal of Consumer Affairs, 39*(2), 339–362.
Romanosky, S., Telang, R., & Acquisti, A. (2011). Do data breach disclosure laws reduce identity theft? [Publisher: John Wiley & Sons, Ltd]. *Journal of Policy Analysis and Management, 30*(2), 256–286. https://doi.org/10.1002/pam.20567
Roumani, Y. (2022). Detection time of data breaches. *Computers & Security, 112*, Article 102508. https://doi.org/10.1016/j.cose.2021.102508
Ruiter, R. A., Kessels, L. T., Peters, G.-J. Y., & Kok, G. (2014). Sixty years of fear appeal research: Current state of the evidence. *International Journal of Psychology, 49*(2), 63–70.
Ruohonen, J., Hjerppe, K., & von Zastrow, M. (2024). An exploratory case study on data breach journalism. In *Proceedings of the 19th international conference on availability, reliability and security*. https://doi.org/10.1145/3664476.3670456
Saeidi, M., Calvert, M., Au, A. W., Sarma, A., & Bobba, R. B. (2022). If this context then that concern: Exploring users' concerns with ifttt applets. *Proceedings on Privacy Enhancing Technologies, 2022*(1), 166–186. https://doi.org/10.2478/popets-2022-0009
Schaub, F., Balebako, R., & Cranor, L. F. (2017). Designing effective privacy notices and controls. *IEEE Internet Computing, 21*(3), 70–77. https://doi.org/10.1109/MIC.2017.75
Schaub, F., Balebako, R., Durity, A. L., & Cranor, L. F. (2015). A design space for effective privacy notices. In *Proceedings of the eleventh USENIX conference on useable privacy and security* (pp. 1–17).
Schaub, F., & Cranor, L. F. (2024). Useable and useful privacy interfaces. In *An introduction to privacy for technology professionals* (2nd ed.). IAPP.
Sheeran, P., Harris, P. R., & Epton, T. (2014). Does heightening risk appraisals change people's intentions and behavior? A meta-analysis of experimental studies [Place: US Publisher: American Psychological Association]. *Psychological Bulletin, 140*(2), 511–543. https://doi.org/10.1037/a0033065
Shulman, Y., Kitkowska, A., & Meyer, J. (2023). Informing users: Effects of notification properties and user characteristics on sharing attitudes. *International Journal of Human-Computer Interaction, 39*(14), 2796–2824. https://doi.org/10.1080/10447318.2022.2086592
Slovic, P., Finucane, M. L., Peters, E., & MacGregor, D. G. (2007). The affect heuristic. *European Journal of Operational Research, 177*(3), 1333–1352. https://doi.org/10.1016/j.ejor.2005.04.006
Smith, H. J., Dinev, T., & Xu, H. (2011). Information Privacy Research: An Interdisciplinary Review [Publisher: Management Information Systems Research Center, University of Minnesota]. *MIS Quarterly, 35*(4), 989–1015. https://doi.org/10.2307/41409970
Solove, D. J. (2013). Privacy self-management and the consent dilemma [GWU legal studies research paper no. 2012-141, GWU law school public law research paper no. 2012-141]. *Harvard Law Review, 126*, 1880–1903. https://ssrn.com/abstract=2171018.
South African Government. (2020). Protection of personal information act (popi act). https://popia.co.za/.

Sun, K., Sugatan, C., Afnan, T., Simon, H., Gelman, S. A., Radesky, J., & Schaub, F. (2021). "they see you're a girl if you pick a pink robot with a skirt": A qualitative study of how children conceptualize data processing and digital privacy risks. In *Proceedings of the 2021 CHI conference on human factors in computing systems*. https://doi.org/10.1145/3411764.3445333

Sun, R., Zhu, Q., Cheng, R. X., Tang, W., Zuo, J., Lv, D., & Qin, S. (2024). Research on the cognitive neural mechanism of privacy empowerment illusion cues regarding comprehensibility and interpretability for privacy disclosures [Publisher: Nature Publishing Group]. *Scientific Reports, 14*(1), 8690. https://doi.org/10.1038/s41598-024-58917-8

Tang, B. J., Sun, K., Curran, N. T., Schaub, F., & Shin, K. G. (2025). Ads that talk back: Implications and perceptions of injecting personalized advertising into llm chatbots. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol., 9*(4). https://doi.org/10.1145/3770640

Tannenbaum, M. B., Hepler, J., Zimmerman, R. S., Saul, L., Jacobs, S., Wilson, K., & Albarracín, D. (2015). Appealing to fear: A meta-analysis of fear appeal effectiveness and theories. *Psychological Bulletin, 141*(6), 1178–1204. https://doi.org/10.1037/a0039729

Tsai, J. Y., Egelman, S., Cranor, L., & Acquisti, A. (2011). The effect of online privacy information on purchasing behavior: An experimental study. *Information Systems Research, 22*(2), 254–268.

Turilli, M., & Floridi, L. (2009). The ethics of information transparency. *Ethics and Information Technology, 11*(2), 105–112. https://doi.org/10.1007/s10676-009-9187-9

United States Congress. (1938). Federal food, drug, and cosmetic act [21 U.S.C. § 352 – Misbranded Drugs and Devices]. https://www.law.cornell.edu/uscode/text/21/352.

Vance, A., Jenkins, J. L., Anderson, B. B., Bjornn, D. K., & Kirwan, C. B. (2018). Tuning out security warnings. *MIS Quarterly, 42*(2), 355–380.

Wagner, I. (2023). Privacy policies across the ages: Content of privacy policies 1996–2021. *ACM Trans. Priv. Secur., 26*(3). https://doi.org/10.1145/3590152

Witte, K., & Allen, M. (2000). A Meta-Analysis of Fear Appeals: Implications for Effective Public Health Campaigns [Publisher: SAGE Publications Inc]. *Health Education & Behavior, 27*(5), 591–615. https://doi.org/10.1177/109019810002700506

Xian, L., Tran, V. H., Lee, L., Kumar, M., Zhang, Y., & Schaub, F. (2025). Layered, overlapping, and inconsistent: A large-scale analysis of the multiple privacy policies and controls of u.s. banks. In *Proceedings of the 2025 ACM SIGSAC conference on computer and communications security* (pp. 3177–3191). https://doi.org/10.1145/3719027.3765072

Zhang, Z., Jia, M., Lee, H.-P., Yao, B., Das, S., Lerner, A., Wang, D., & Li, T. (2024). "it's a fair game", or is it? Examining how users navigate disclosure risks and benefits when using llm-based conversational agents. In *Proceedings of the CHI conference on human factors in computing systems*. https://doi.org/10.1145/3613904.3642385

Zipkin, D. A., Umscheid, C. A., Keating, N. L., Allen, E., Aung, K., Beyth, R., Kaatz, S., Mann, D. M., Sussman, J. B., Korenstein, D., et al. (2014). Evidence-based risk communication: A systematic review [PMID: 25133362]. *Annals of Internal Medicine, 161*(4), 270–280. https://doi.org/10.7326/M14-0295

Zou, Y., Le, K., Mayer, P., Acquisti, A., Aviv, A. J., & Schaub, F. (2024). Encouraging users to change breached passwords using the protection motivation theory. *ACM Transactions on Computer-Human Interaction, 31*(5). https://doi.org/10.1145/3689432