Behaviour & Information Technology

ISSN: 0144-929X (Print) 1362-3001 (Online) Journal homepage: www.tandfonline.com/journals/tbit20

Legally compliant personalised prioritisation of privacy policy information shows no effect on user engagement, comprehension, or workload

Meihe Xu , Jannis Strecker-Bischoff , Clement Guitton , Kenan Bektas , Aurelia Tamo-Larrieux & Simon Mayer

To cite this article: Meihe Xu , Jannis Strecker-Bischoff , Clement Guitton , Kenan Bektas , Aurelia Tamo-Larrieux & Simon Mayer (29 Jun 2026): Legally compliant personalised prioritisation of privacy policy information shows no effect on user engagement, comprehension, or workload, Behaviour & Information Technology, DOI: 10.1080/0144929X.2026.2692098

To link to this article: https://doi.org/10.1080/0144929X.2026.2692098

Published online: 29 Jun 2026.

BEHAVIOUR & INFORMATION TECHNOLOGY
https://doi.org/10.1080/0144929X.2026.2692098

Legally compliant personalised prioritisation of privacy policy information shows no effect on user engagement, comprehension, or workload

Meihe Xu[a], Jannis Strecker-Bischoff[b], Clement Guitton[b], Kenan Bektas[b], Aurelia Tamo-Larrieux[c] and Simon Mayer[b]

[a]Department of Private Law, Maastricht University, Maastricht, the Netherlands; [b]Institute of Computer Science, University of St. Gallen, St. Gallen, Switzerland; [c]Faculty of Law, Criminal Justice and Public Administration, University of Lausanne, Lausanne, Switzerland

ABSTRACT

Privacy policies function as both legal documents and information sources for users, but their length and complexity often discourage engagement. In this paper, we investigate whether a personalised approach can address this issue by prioritising information that concerns individual users most while maintaining a policy's legal compliance on disclosure. We first explored whether personal characteristics can be used to predict a person's most concerned category and, hence, serve as a baseline for personalisation. We then conducted an eye-tracking experiment and interviews (n = 30) to understand the effectiveness of personalised reordering of privacy policies. In the interviews, many participants perceived personalised reordering as helpful, although others raised concerns about the invasion of privacy through this personalisation. The eye-tracking results indicate that personalised reordering leads to higher engagement for the first few sentences of a privacy policy. Based on our findings, we present design recommendations for creating legally compliant forms of privacy disclosures that encourage user engagement as well as discussions and implications on privacy disclosure compliance.

ARTICLE HISTORY
Received 30 October 2025
Accepted 11 June 2026

KEYWORDS
personalisation; privacy policies; personalised law; eye tracking; privacy disclosure; user study

1. Introduction

Privacy policies serve as the primary channel for companies to inform users about their data practices. Regulations like the EU General Data Protection Regulation (GDPR) (European Parliament and Council of the European Union 2016), California's California Privacy Rights Act (CPA) (California Consumer Privacy Act of 2018), and China's Personal Information Protection Law (Standing Committee of the National People's Congress 2021) all mandate the information that privacy disclosures must contain. These regulations also impose transparency requirements on data controllers that are uniform across sectors, technologies involved, and associated risk levels, as demonstrated in GDPR (Ebers and Sein 2024; Xu, Jug, and Tamò-Larrieux 2024). Such *uniform* privacy policies have been found to disadvantage users as they are challenging to understand (Chen et al. 2023; Degeling et al. 2019; Proctor, Athar Ali, and Vu 2008; Reidenberg et al. 2015), too long and time-consuming to read (McDonald and Cranor 2008), and becoming lengthier (Wagner 2023), ultimately causing users to completely ignore them (Obar and Oeldorf-Hirsch 2020). However, with companies' disturbing data practices hidden in the lengthier and more convoluted privacy policies (Wagner 2023), such behaviour leaves users more vulnerable and disempowered.

To address these challenges, transparency-enhancing technologies (e.g. Janic, Wijbenga, and Veugen 2013) and alternative disclosure methods like privacy labels (e.g. Kelley et al. 2009; Kitkowska et al. 2020; Zhang et al. 2024) aim to aid users in better understanding data collection and processing practices. However, these often do not meet the requirements set by privacy regulations and cannot replace privacy policies *as legally binding documents*. Adoption of such alternatives is hence left as optional, and privacy policies remain the fundamental and legally mandated mechanism for informing users about data practices and user rights under current regulations.

The personalisation of privacy policies by adapting their content based on an individual user's personal preferences might present an approach that allows comprehensive yet more engaging policies for users. However, the literature to date has not answered

CONTACT Meihe Xu m.xu@maastrichtuniversity.nl Bouillonstraat 3, 6211 LH, Maastricht, the Netherlands

whether personalised disclosure of privacy policies is effective (e.g. in terms of user engagement or comprehension) nor how to practically *operationalise* personalisation (i.e. a system-initiated adaptation based on personal data (Strecker, Mayer, and Bektaş 2025) in this context. While legal scholars have used general phrases such as 'personalized mandated disclosures' (Ben-Shahar and Porat 2021), 'personalized privacy notices' (Busch 2019), or 'granular consumer information' (Luzak 2021), they did not zoom in on the specific types of legally mandated disclosure documents (e.g. privacy policies).

A straightforward way to personalise a privacy policy for users would be to directly ask them about their privacy concerns and adapt the policy accordingly, as described by Ben-Shahar and Porat's concept of a *descriptive criterion* of need (Ben-Shahar and Porat 2021, 93). However, such an approach can be flawed, as users often choose to skip reading privacy policies entirely (Obar and Oeldorf-Hirsch 2020), and thus users themselves might not know how their privacy concerns relate to or are located in a specific privacy policy.

An alternative is leveraging proxies to predict what a user wants to know. As proxies can come in all shapes and sizes, in this paper, we specifically look at personal characteristics and how they might work as proxies to a user's most concerned category of information in a privacy policy. Our focus on concerns to drive personalisation is motivated by their impact on user decision-making. Privacy concerns are found to be negatively related to a user's willingness to disclose ther personal information to a service provider when he or she reads a privacy policy during sign-ups (Kitkowska et al. 2023). Further thematic analysis from that study shows that a user's agreement or disagreement with a privacy policy is driven less by its overall presentation of information, and more by the specific privacy concerns a user cares about, for example, which data will be collected and how it will be processed (Kitkowska et al. 2023). Consequently, quick access to information that matches each user's primary privacy concern is important if the content and display of a privacy disclosure do not accommodate each user. Earlier studies (Bansal and Gefen 2010; Lee et al. 2019; O'Neil 2001) link personal characteristics to overall privacy-concern levels; here, we test whether those characteristics predict a user's primary concern in a privacy policy. As an initial step, we focus on predicting a user's single most-concerned category.

Considering the context-dependency of users' privacy concerns and preferences (Acquisti, Brandimarte, and Loewenstein 2015; Ebert, Ackermann, and Heinrich 2020), we choose social media as the context due to its embeddedness in many people's lives. Our first research question, therefore, is:

> **RQ1** How well can we predict, based on personal characteristics, a social media user's most concerned category of information in a privacy policy?

If we can successfully predict this category, we can use the result to personalise the display of a privacy policy for specific users. So far, few studies tackle the actual personalisation of *privacy policies* and their *impacts*. For instance, Chang et al. propose automatically providing users with segments of a privacy policy based on their privacy concerns (Chang et al. 2019). However, the issue with *partial* provisioning of information is that this *undermines* full disclosure as granted by law. Since disclosure requirements are set with privacy laws (e.g. the GDPR), providing partial information (including summaries as well as scores or ratings) would be deemed insufficient for legal compliance by a company. Beyond personalisation, researchers have been exploring the design of contextual privacy (Ortloff et al. 2020; Windl et al. 2022), which involves presenting users with relevant information from a privacy policy within the appropriate context. This approach may fall short if different users have varying privacy priorities and concerns within the same context, or if users are indifferent to certain contexts. It may also fail to provide users with a coherent, full picture of a company's personal information processing, as disclosure would be segmented across different areas. To address these challenges, a first step to personalising privacy policies is maintaining the full text in a single document while adjusting the content order based on a user's prioritisation of privacy concerns (e.g. displaying the content most concerning to the user first). This method ensures legal compliance while having the potential to enhance user engagement. We hence address the second research question:

> **RQ2** Does personalizing the order of sections in a privacy policy lead to improved user engagement and comprehension, and does it reduce user cognitive workload compared to a generic privacy policy?

This article is structured into three main sections to tackle the research questions. Section 2 delves into RQ1 with a survey study to explore the possibility of predicting a user's most concerned category of information in a privacy policy based on their personal characteristics. Section 3 addresses RQ2 by exploring personalising privacy policies and comparing their effects to generic privacy policies through an eye-tracking experiment. The eye-tracking experiment is further

complemented by semi-structured interviews to bring further nuances regarding users' perceptions of the design. Finally, Section 4 discusses the implications of our research regarding possible future directions for privacy policies.

2. Study 1: prediction of most concerned category of information in a privacy policy

In our first study, we explore whether and which personal characteristics may be used to predict a user's most concerned category of information in a privacy policy. To do so, we utilise existing research that focuses on personal characteristics and their correlations with an individual's privacy concerns, including sociodemographic factors (age, gender, and education level), the Big Five personality traits (also called the OCEAN model [Costa and McCrae 1999]), online privacy literacy, and motivation for privacy protection behaviour. The rationale behind this is that if these factors are found to be correlated, even with mixed findings, with the degree of privacy concerns in existing research, they might also be indicative of the specific categories of privacy concerns a user has in a privacy policy. An important note needs to be taken into consideration here: these personal characteristics as well as the categories of concerns should be understood within a cultural/regional context (for example, see the differences in shared concerns in different regions in [Bergström 2015; Chen and Zhang 2021; Xu et al. 2012]).

2.1 Related work

2.1.1 Age

While their privacy attitudes are similar (Hoofnagle et al. 2010), older adults are significantly more concerned about online privacy and information privacy than younger adults (Broeck, Poels, and Walrave 2015; Li and Borah 2018; Zukowski and Joseph Brown 2007). Furthermore, adolescents are consistently less concerned than adults about their privacy (Miltgen and Peyrat-Guillard 2014; Steijn and Vedder 2015). We are interested in whether age contributes to the odds of specific categories being a social media user's most concerned category of information in a privacy policy.

2.1.2 Gender

Regarding gender, some studies indicate higher privacy concerns among women than men, e.g. regarding Facebook usage (Lankton, Harrison McKnight, and Tripp 2017), or the activation of privacy settings and untagging of photographs on social media (Tifferet 2019). Other research, however, has not found a difference between genders in the degrees of information privacy concerns (Faja and Trimi 2008; Li and Borah 2018; Zukowski and Joseph Brown 2007). In our study, we are interested in whether gender contributes to a social media user's selection of their most concerned category in a privacy policy.

2.1.3 Education level

Education level has been found to correlate with the level of privacy concerns. However, results are not consistent regarding the directionality of this correlation. Some studies (Lee et al. 2019; Sheehan 2002) find that individuals with higher levels of education are more concerned about their privacy online than their peers with less education, while others (Bergström 2015; Bhatia and Breaux 2018; Zukowski and Joseph Brown 2007) show that higher education level is associated with lower levels of concern about information privacy among Internet users. We are interested in whether an individual's education level influences their most concerned category in a privacy policy.

2.1.4 The big five personality traits

Personality traits, using the Big Five model (extroversion, agreeableness, openness, conscientiousness, and neuroticism) (John, Donahue, and Kentle 1991), are found to significantly impact individuals' degrees of concern about privacy. However, the directionality of personality's impact seems to vary in different contexts. High scores in agreeableness and conscientiousness are associated with increased concern for information privacy (Korzaan and Boswell 2008; Osatuyi 2015). Extroversion is found to negatively impact privacy concerns in less sensitive contexts (such as in e-commerce but not in the health context), and more agreeable people are more concerned about their privacy than those who are less agreeable (Bansal, Zahedi, and Gefen 2016). In the context of location-based services, highly agreeable people have lower concerns for privacy compared to their less agreeable counterparts, while more conscientious and more open people have higher privacy concerns (Junglas, Johnson, and Spitzmüller 2008). In this study, we are interested in exploring whether personality traits can predict a social media user's most concerned category in a privacy policy.

2.1.5 Online privacy literacy

Online privacy literacy is about '*privacy and online privacy-related behaviors*', which serves '*as a principle to support, encourage, and empower users to undertake informed control of their digital identities*' (Park 2013, 217). Online privacy literacy is argued to empower

users to self-protect and self-determine in the online space (Masur 2020). On social networking sites, online privacy literacy contributes to a more cautious activity and higher perceived security of users (Bartsch and Dienlin 2016). Higher literacy often leads to increased concerns about privacy (Prince et al. 2021), and users with higher online privacy literacy are more likely to utilise privacy protective measures (Baruh, Secinti, and Cemalcilar 2017). We are interested in whether the degree of online privacy literacy influences a user's selection of their most concerned category in a privacy policy.

2.1.6 Privacy protection motivation

Protection Motivation Theory posits that individuals' motivation to protect their privacy is influenced by both threat appraisal (perceived susceptibility/vulnerability to and severity of privacy threats) and coping appraisal (self-efficacy and response efficacy of privacy protection behaviours) (Boerman, Kruikemeier, and Zuiderveen Borgesius 2021). Youn has shown that perceived vulnerability to privacy risks among young adolescents significantly heightens their online privacy concerns (Youn 2009). Similarly, perceived severity, self-efficacy, and perceived vulnerability are found to positively correlate with information privacy concerns in the context of social media in Malaysia (Mohamed and Ahmad 2012). Adhikari and Panda found perceived vulnerability, perceived severity, and self-efficacy to be significantly correlated with a user's privacy concerns in the context of social media using a sample of Indian participants (Adhikari and Panda 2018). In healthcare, Zhang et al. found a negative correlation between response efficacy, self-efficacy and privacy concerns, while perceived vulnerability and severity positively influenced such concerns (Zhang et al. 2018). These findings highlight the multifaceted nature of privacy protection motivation and its impacts on the extent of an individual's privacy concerns across different contexts. We are interested in whether motivation for privacy protection behaviour informs a social media user's most concerned category in a privacy policy.

2.1.7 Privacy concerns

Existing measures on privacy concerns mostly focus on the degrees of privacy concerns a person has (Preibusch 2013). The Internet Users' Information Privacy Concerns (IUIPC) framework proposed by Malhotra et al. (Malhotra, Kim, and Agarwal 2004) includes three dimensions: collection, control, and awareness. However, it remains unclear whether such a degree of concern can accurately predict a person's specific concern in a privacy policy. For example, would a high degree of privacy concern measured in IUIPC make users more likely to choose a section in a privacy policy as their most concerned category, such as one that covers a company's information collection practices? In this study, we are interested in exploring degrees of privacy concerns and their influence on a person's selection of their most concerned category of information in a privacy policy.

2.2. Methods

To understand the relationship between individual characteristics (socio-demographics, personality, and privacy-related factors) and the selection of primary privacy concerns in privacy policies, we conducted a survey with participants from Germany. Participants were recruited using Prolific with a prior-approval rate of over 80%. The survey was administered in English in July 2023 and took approximately 15 min. The study received ethical approval from the Ethics Review Committee of the lead author's institution (ERCIC_396_24_11_2022), and all participants were compensated USD 4.60 in local currency for their participation. We recruited 220 participants. After excluding participants who failed one or more of the five attention check questions, we retained 207 valid responses.

2.2.1 Independent variables

In the survey, we first collected socio-demographic information, including age in years, gender (men/women/non-binary/other), and level of education (from early childhood education to Doctorate or equivalent). We aggregated age and level of education into brackets. Age was put into different age groups based on the classification used by UN Trade and Development (U. N. Trade and Development 2023). Level of education was bracketed based loosely on the International Standard Classification of Education developed by UNESCO (International Labour Organization 2024) (see Table A3 in Appendix C). We measured participants' personality traits using the Big Five Inventory (BFI) with 44 items on a 5-point Likert scale (John, Donahue, and Kentle 1991; John, Naumann, and Soto 2008). We measured online privacy literacy using a questionnaire adapted from the Online Privacy Literacy Scale (OPLIS) (Trepte et al. 2015). Since OPLIS was developed prior to the GDPR's entry into force, we updated the questions by adding GDPR-related items and removing those specific to German data protection laws. A law student from the lead author's institution fact-checked these questions, which were further refined through feedback from legal scholars. This

Table 1. Multicollinearity diagnostics for privacy-related predictors.

	Unstandardised B	Coefficients Std. Error	Standardised Coefficient Beta	t	Sig.	Collinearity Statistics	
						Tolerance	VIF*
(Constant)	30.618	0.640		47.858	< .001		
IUIPC-8	1.859	0.717	0.196	2.593	.010	.800	1.250
Privacy motivation	0.133	0.707	0.014	0.188	.851	.824	1.214
Correct rate privacy literacy	1.267	0.668	0.133	1.896	.059	.921	1.086

**Note.* The dependent variable is the age of participants.

section comprised 20 true/false or multiple-choice questions, each with a single correct answer to indicate participants' online privacy literacy. As we did not validate the questions, they are only indicative of a person's online privacy literacy through one's accuracy rate on the questions. Privacy protection motivation was assessed using a scale by Boerman et al. (Boerman, Kruikemeier, and Zuiderveen Borgesius 2021). Privacy concerns were measured using the IUIPC-8 scale (Groß 2021). All scales were adjusted to a 5-point range, aligning with the survey's overall design. All continuous independent variables were standardised.

2.2.2. Dependent variable and statistical analysis

The dependent variable is defined as a participant's most concerned category of information in TikTok's EU privacy policy. We opted for TikTok for its global reach and for future comparison purposes. A one-sentence summary of each section in the privacy policy was manually generated by one author, who is a data protection law scholar. The summarised statements included examples to explain what the sections entail (see Table A1 in Appendix A for details), and we asked users to rank the summarised statements from most to least concerned. A participant's *most concerned category of information* was thus measured as a participant's top choice of the presented summarised statements.

We also asked participants to rank the section headings of TikTok's privacy policy from the most to least concerning. The order of sections in the headings and summarised statements was randomised to minimise the potential order effect. We presented the headings at the very beginning of the survey and the summarised statements at the very end. The tasks and time spent during the survey should minimise the carryover effect for our participants. The goal of ranking twice was to explore further whether summarising the content of each section and adding examples to a one-sentence summary would influence participants' selection of their most concerning category in a privacy policy.

We used multinomial logistic regression for data analysis, with each summarised privacy statement considered a category of the dependent variable. When setting up the multinomial logistic regression model, we chose the statement with the highest frequency as the reference category, i.e. 'What information we collect'. For categorical independent variables, we set the reference categories as 'woman' for gender, '40-64' for age groups, and 'Master's and doctoral' for level of education.

2.3. Results

To evaluate the assumption of multicollinearity among the privacy-related predictors, Variance Inflation Factor (VIF) was examined. In the present model, VIF values ranged from 1.08–1.25, suggesting no significant multicollinearity among the predictors (see Table 1). These results indicate that the predictors are sufficiently independent to ensure stable regression estimates for multinomial logistic regression.

After checking the multicollinearity of privacy-related predictors, we conducted the multinomial logistic regression. Overall, we received an indication that the model is fitting well ($\chi^2 = 161.317$, $p < .05$, $R^2_{McF} = .21$, see Tables 2 and 3). Likelihood ratio tests were used to assess the contribution of each predictor to the model. The results indicated that most predictors did not significantly contribute to the model, except Neuroticism ($\chi^2 = 24.9$, $p < .01$, see Table 4).

To counter Type I errors of multiple comparisons, we adjusted the p value of neuroticism with the Holm procedure (Chen, Feng, and Yi 2017) (Family-wise $\alpha = .05$

Table 2. Model fit for multinomial logistic regression.

Model	Model fitting criteria −2 Log Likelihood	χ^2	df	Sig.
Intercept only	768.122			
Final	606.805	161.317	126	.018

Table 3. Pseudo R-Square statistics for multinomial logistic regression.

Pseudo R-Square	
Cox and Snell	.541
Nagelkerke	.555
McFadden	.210

Table 4. Likelihood ratio tests.

Effect	Model Fitting Criteria −2 Log Likelihood of Reduced Model	Likelihood Ratio Tests		
		Chi-Square	df	Sig.
Intercept	606.805	.000	0	.
Privacy Concerns	618.779	11.974	9	.215
BFI Extroversion	612.731	5.926	9	.747
BFI Agreeableness	616.185	9.380	9	.403
BFI Conscientiousness	618.116	11.311	9	.255
BFI Neuroticism	631.703	24.898	9	.003
BFI Openness	617.797	10.992	9	.276
Privacy Motivation	617.072	10.267	9	.329
Privacy Literacy	623.113	16.308	9	.061
Age	629.070	22.265	18	.220
Education Level	630.736	23.931	27	.634
Gender	616.988	10.183	9	.336

with 81 comparisons), after which the adjusted p value (adjusted $p = .243$) was larger than .05. The nine category-specific p values for neuroticism were further evaluated as one family. After Holm adjustment (Family-wise $\alpha = .05$), only one of the individual p values for neuroticism remained significant for predicting the category of 'Rights and choices' compared to the reference category (adjusted $p < .05$). This indicates that participants with higher neuroticism scores are less likely to choose 'Your rights and choices' than the reference category of 'What information we collect'. For every one-unit increase in neuroticism, the odds of selecting this category decreased by 79.5% (OR = 0.205, 95%$CI = 0.07 - 0.58$, adjusted $p < .05$). For the odds ratios of predictors that are significant prior to Holm adjustment for category-specific comparisons, see Table 5.

To check whether '*the degree to which predicted probabilities agree with actual outcomes*', we used a classification table (Peng et al., 2002, p.6). The classification table (see Table A2 in Appendix B) shows that the multinomial logistic regression model can overall successfully predict the dependent variable based on the predictors 49.8% of the time, with the highest rate being 87.8% for the highest-frequency category 'What information we collect', but performed poorly for categories 'Legal bases and information processing' and 'Young users' (0%).

2.3.1 Headings vs. summarised privacy statements

In addition to the multinomial logistic regression analysis of different predictors on participants' most concerned category of information, we further compared whether the presentation of privacy concerns (headings vs. summarised statements) from a privacy policy to participants would influence the selection of their most concerned category. A Chi-square test for homogeneity was conducted to compare the distribution of participants' most concerned categories when they were presented with headings and with summarised statements. The result showed a significant difference (χ^2 (9, N = 414) = 35.651, $p < .001$, see Table 6, and Figure A1 in Appendix B). The effect size, measured by Cramér's V is 0.29, indicating a moderate effect. The results indicate that simply providing a one-sentence summary with examples for each section can significantly influence participants' most concerned category of information in a privacy policy.

2.4. Discussion

To answer RQ1, our model does not accurately predict a participant's most concerned category of information in a privacy policy, with an overall correct prediction rate of 49.8%. Although the model has an 87.8% accuracy for the section 'What information we collect', the relatively high percentage is likely a result of an imbalanced frequency of categories, which represents a limitation of our study. This imbalance may inflate the prediction accuracy for the most selected category while contributing to poor prediction for categories with low frequencies. For future studies testing personal characteristics on predicting one's most concerned category of information in a privacy policy, we recommend using

Table 5. Odds ratios of predictors that are significant prior to Holm adjustment for category-specific comparisons in the multinomial logistic regression model (Reference category: 'What information we collect').

Category	Predictor	B	OR (95%CI)	*p*	Adjusted *p*
Information usage	Level of education: Below high school	2.776	16.05 (1.484–173.578)	.022	.60
	Level of education: High school and post-high school	1.570	4.806 (1.089–21.221)	.038	.96
	Gender: Men	1.122	3.070 (1.024-9.206)	.045	1
Rights and choices	Neuroticism	−1.585	0.205 (0.072–0.579)	.003	.027
Security and retention	Conscientiousness	−1.074	0.341 (0.134-0.873)	.025	.20
	Neuroticism	−1.577	0.206 (0.056–0.768)	.019	.152
Global operations and transfers	Conscientiousness	−0.822	0.440 (0.230–0.842)	.013	.117
Young users	Age class: 18–24	−3.308	0.037 (0.002–0.720)	.030	1
Updates	Agreeableness	1.931	6.898 (1.253–37.981)	.027	.243
	Privacy literacy	−1.925	0.146 (0.035–0.613)	.009	.081

Note: Adjusted *p* values are corrected with the Holm-Bonferroni method within each predictor (Family-wise $\alpha = .05$). Unadjusted odds ratios and 95% CIs are shown.

Table 6. Chi-Square result and its effect size for participants' most concerned category of information under headings and summarised statements.

	Value	df	Asymptotic Significance (2-sided)
Pearson Chi-Square	35.651	9	<.001
Likelihood Ratio	43.208	9	<.001
Linear-by-Linear Assoc.	.003	1	.957
N of Valid Cases	414		
Nominal by Nominal	Value		Approximate Significance
Cramer's V	.293		< .001

purposive sampling so that there will be an equal distribution of samples for each category in a privacy policy.

Although neuroticism displayed significance before correction ($p < .01$) in the likelihood ratio tests, it did not pass the Holm adjustment. The result of this multinomial logistic regression should therefore be interpreted cautiously. Barlow et al. connect neuroticism to the feeling of uncontrollability when facing stress (Barlow et al. 2014), which is shown to be linked to higher intolerance of uncertainty (Yang et al. 2015; Zhang et al. 2022). If individuals with higher neuroticism are more likely to be intolerant of uncertainty, then in principle, knowing specific categories of information in a privacy policy might help them feel in control of their data and potentially mitigate feelings of uncertainty. Prior to the adjustment, while neuroticism is significant for the model, it is not significant across all the categories in the dependent variable. Other predictors significantly influence the odds of specific categories (see Table 5). Interestingly, the predictor measuring privacy concerns did not significantly contribute to the general model or the odds of specific categories. This indicates that an individual's overall level of privacy concern does not necessarily affect their selection of the most concerned category of information when reading a privacy policy. In other words, if a user scores high on privacy concerns measured by IUIPC-8, it does not mean they would pick the measured concern(s) as their most concerned category of information when it appears in a privacy policy.

Overall, this study provides valuable exploratory insights that demonstrate personal characteristics alone do not allow accurate predictions of a person's most concerned category of information in a privacy policy. Alternatively, future research could explore additional predictors, consider alternative modelling approaches, and investigate the potential for developing more nuanced models incorporating interaction effects. Additionally, further analysis could benefit from a larger dataset and/or more balanced data to improve the model.

In the study, we also compared participants' most concerned categories under section headings and summarised statements. The significant result (Table 6) indicates that participants made different choices on their most concerned category of information depending on whether they were given headings or summarised privacy statements. Considering people generally skip privacy policies (Obar and Oeldorf-Hirsch 2020), they likely are unaware of what is actually included in each section. Thus, in our study, when participants were better informed about the content of each section with summarised statements, they were more likely to switch their choices to align better with what they actually were concerned about. Xu et al. demonstrate that information organisation in a privacy policy does not necessarily follow the structure of legal requirements for disclosure (Xu, Jug, and Tamò-Larrieux 2024). The discrepancy between law and compliance by companies could add additional hurdles for users to pinpoint their most concerned categories of information, considering the information could be dispersed in multiple sections or mixed with other information in one section. This finding suggests a possible direction for service providers when presenting privacy policies to users. Rather than having a table of contents consisting only of section headings, service providers could consider providing a one-sentence summary of the key privacy implications instead. The one-sentence summary could further help users to better pinpoint the sections containing their most concerned information.

3. Study 2: effects of privacy policy personalisation on user engagement, comprehension, and workload

Despite the clear limitations of using personal characteristics as proxies for an individual's most concerned category of information, reordering privacy policy sections for a more personalised experience could still significantly enhance user engagement compared to a generic policy. We use the term personalised to describe a system-initiated adaptation of an interface based on (personal) data (Fan and Poole 2006; Salonen and Karjaluoto 2016; Strecker, Mayer, and Bektaş 2025). In our second study, we explore whether a personalised reordering of privacy policies leads to improved user engagement and comprehension, and whether it reduces user workload compared to a generic policy.

3.1. Related work and hypotheses

Research on user engagement with privacy policies consistently shows that people often choose not to engage with them, frequently skipping them altogether (Obar

and Oeldorf-Hirsch 2020). This finding is supported by eye-tracking experiments that are able to show how users actually behave when reading texts compared to self-reporting methods (Steinfeld 2016; Vu et al. 2007). Factors influencing engagement with privacy disclosures often include a person's degrees of privacy concerns (Milne and Culnan 2004), the way disclosures are presented (Ebert, Ackermann, and Scheppler 2021), and disclosure comprehensibility (Ibdah et al. 2021; Milne and Culnan 2004). Additionally, there is often a disconnection between what users prioritise and what companies emphasise in these policies (Kununka et al. 2017), with current privacy policies being more company-centred than user-centred (Ding and Huang 2024). This misalignment suggests that uniform privacy policies in their current form do not motivate users to engage and read. Alternative forms of disclosures address these shortcomings of privacy policies, such as privacy labels (Kelley et al. 2009; Zhang et al. 2022), privacy notices (Ebert, Ackermann, and Scheppler 2021), user-tailored privacy (Knijnenburg et al. 2022), or summarised privacy policies (Woodring, Perez, and Ali-Gombe 2024). While these approaches present privacy-related information in a more engaging form than traditional privacy policies, they do not comply with privacy regulations due to their simplification of mandated disclosure content; thus, they cannot replace privacy policies as legally binding documents.

In Study 2, we therefore investigate whether a personalised privacy policy design that remains legally binding increases user engagement (i.e. reading time) and comprehension by prioritising a user's most concerned category of information at the beginning of the policy. Study 1 shows that predicting a user's most concerned category of information through proxies is challenging; Study 2 asks users directly about their preferences, allowing for more effective tailoring to test the design. We thus hypothesise that:

H2.1 Participants will spend more time reading the personalized privacy policies than generic ones.

H2.2 Participants' comprehension of a privacy policy's content will be higher in the personalized condition than in the generic condition.

Reading a privacy policy can be hindered by a user's information overload, which occurs '*when an individual's efficiency and effectiveness in using information are hampered by the amount of relevant, and potentially useful, information available to them*' (Bawden and Robinson 2020, 12). Obar and Oeldorf-Hirsch (Obar and Oeldorf-Hirsch 2020) found that this overload negatively impacts users' willingness to read terms of service agreements, especially during sign-ups, service changes, or privacy policy updates. When cognitive workload is reduced by making the same content more digestible, people's comprehension and verification accuracy increase (Moody 2004). While our design does not alter the content or overall structure of a privacy policy, the approach (i.e. rearranging the order of content) could minimise users' need to filter out irrelevant details and reduce their need to navigate through less concerned information compared to a generic privacy policy, hence, reducing their cognitive workload. Therefore, we further hypothesise that:

H2.3 Participants will have a lower perceived workload when reading a personalized privacy policy than a generic one.

3.2. Methods

In an eye-tracking experiment, we studied H2.1 through eye-tracking measurements, H2.2 with the help of Cloze tests, and H2.3 with the NASA-TLX questionnaire. Semi-structured interviews were conducted, where we asked participants for their opinions on our personalisation design and on privacy policies in general. To elicit participants' natural reactions towards privacy policies, we framed the experiment as a usability study on a novel app marketplace interface where users experienced the interface during two app installations. In fact, it was a within-subject design in a randomised order. Each installation presented a condition, a personalised or generic privacy policy, and we compared participants' engagement, comprehension, and perceived workload in each condition. We opted for a within-subject design to minimise the influence of inter-individual variability in factors such as usual reading behaviour, prior experience, and general willingness to engage with privacy policies. In the personalised condition, the content order in the policy was rearranged based on a user's most concerned category of information. We used TikTok's EU privacy policy in both conditions to ensure consistency with Study 1. After the experiments, participants took semi-structured interviews. Then, they were debriefed and compensated according to the guidelines of the host institute.

3.2.1 Participants

We randomly recruited 31 individuals ('women': 9, 'men': 20, 'Prefer not to say': 1) from a Swiss university using an internal recruitment platform. We excluded the data from one participant due to the poor quality of their recorded eye-tracking data. The average age of

the participants was 25.5 years (SD = 3.6), with the majority (N = 24) holding a Bachelor's degree or higher (see Table A4 in Appendix C). Six participants reported wearing vision correction. The data collection sessions took approximately 45–60 min. The study received approval from the ethical review board of the lead author's institution (ERCIC_425_02_03_2023), and participants were monetarily compensated (≈ USD 30 in local currency) for their time.

3.2.2 Pre-experiment survey

A pre-experiment survey was sent to participants to gather their demographic data and to identify their most concerned category of information, which served to personalise the privacy policy for the experiment. Participants were not informed that their answers would be used later in the main study. Their lab sessions were scheduled to ensure a minimum of 24 h between the survey and the experiment. For the ranking of categories, we used summarised privacy statements from the first study (see Appendix A). Most participants (N = 11) picked 'What information we collect' as their most concerned category, followed by 'How we use your information' (N = 6; see Table A4 in Appendix C). Since we framed Study 2 as understanding the usability of an app marketplace interface, the pre-study survey also included questions about people's experiences with existing marketplaces, such as Apple's App Store, to align the pre-experiment survey with our cover story.

3.2.3 Prototypes

We used Figma to create a mock marketplace for two app installation processes. Both apps were designed to be visually distinct and were framed as apps for specific aspects of TikTok: TrendTrails (see Figure 1(b)), which enables the discovery of current trends on TikTok for the generic condition, and FilterFusion (see Figure 1(a)), which enables the creation of Augmented Reality filters for TikTok videos for the personalised condition. The marketplace and the app details mimic existing marketplaces and apps so that they appear familiar to the participants. We created the displayed description texts for the two apps with the help of ChatGPT (excluding the privacy policies).

For each app, participants first saw a marketplace overview with the icon of the respective app. Upon clicking the icon, a page with information about the app was visible with an 'Install' button which led to a page that showed a description of the app's functionality. Then, two pages with app preferences followed, to make the installation process more credible. Afterwards, the privacy policy was shown on eleven pages (see Figure 1 and Figure 2) that the participants could navigate with a 'Continue' and a 'Back' button. We opted for clickable buttons instead of scrolling to navigate the privacy policy because this format allows more precise analysis of the eye-tracking data. The participants had to click a checkbox to accept the privacy policy on the last page of the privacy policy. The participants clicked the 'Finish' button to end the installation process, thereby completing one trial.

3.2.4 Apparatus

The app prototypes (see Figure 1) were presented on a 1920 × 1080px monitor (HP E24 G4) with the Tobii Pro Fusion (120 Hz) eye tracker attached to it. The research software iMotions was used to record participants' eye movements while they interacted with the prototypes. Participants were seated around 60 cm away from the monitor.

3.2.5 Measuring text comprehension

To measure participants' comprehension of the privacy policies, we employed Cloze tests (Bormuth 1968), in which participants had to fill in every fifth word in excerpts sampled from three different sections in the two versions of the privacy policies. Cloze tests are shown to correlate with established measures of reading comprehension (Gellert and Elbro 2013; Greene 2001; Williams, Ari, and Santamaria 2011), indicating that they are likely to measure the same or similar constructs as comprehension tests. Furthermore, they were not intended to measure participants' general cognitive abilities but comprehension. This approach is particularly appropriate due to the within-subject design of the study, since comparisons were made only between how the same participant comprehended texts without the confounding factors such as the different cognitive abilities or the different reading comprehension levels of participants. Thus, Cloze tests can serve as a reliable measurement for assessing comprehension. In addition, alternatives such as multiple-choice questions were considered to measure comprehension during the brainstorming phase. However, such alternatives were dropped due to concerns that multiple-choice questions could pose a risk of false positives, as participants might guess the correct answers. As a result, Cloze tests were deemed more appropriate to accurately measure comprehension.

For the personalised condition, we created different versions of the Cloze tests corresponding to the three sections a participant would see during the experiment: the first section, one in the middle, and one at the end. Each participant only received one version of the Cloze test based on the selection of their most concerned

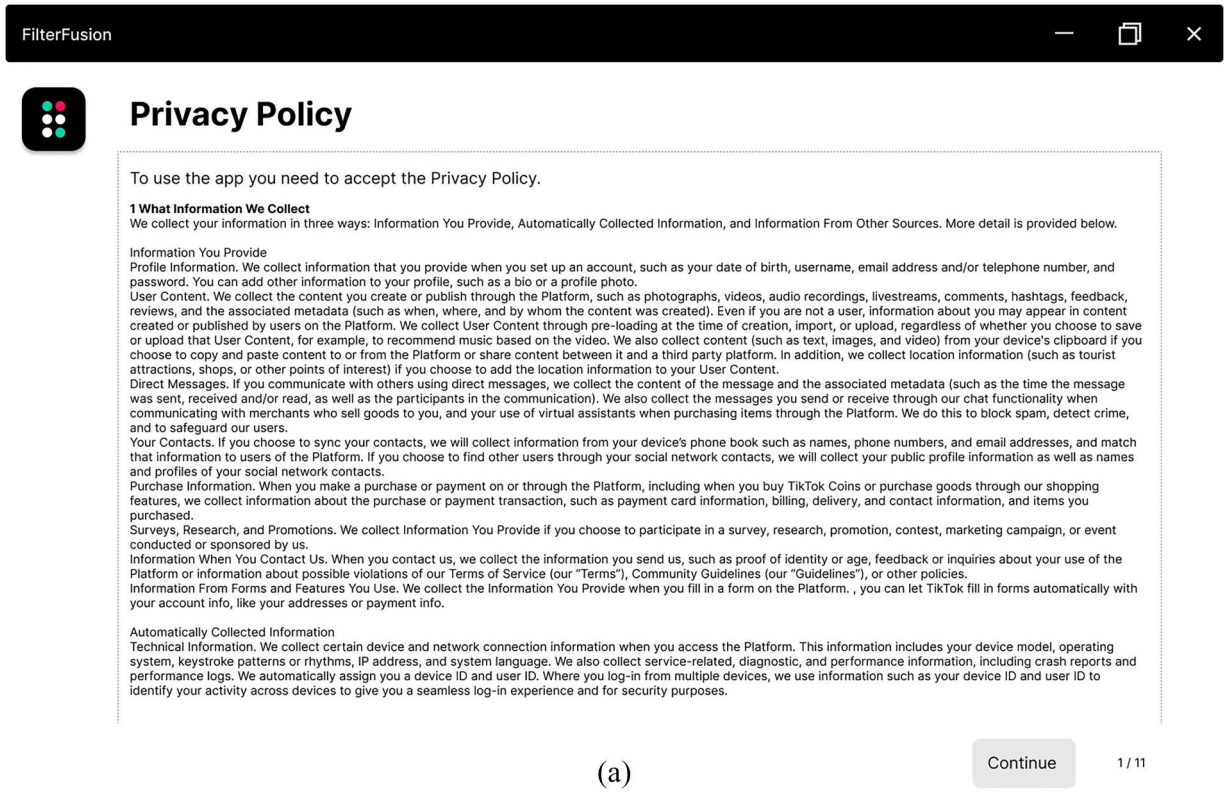

Figure 1. (a) The prototype of the app as seen by the participants in the browser during the study in the personalised condition. (b) A screenshot from the analysis of a trial recording from the generic condition, showing an AOI (here: the most concerned section at its original position in a generic privacy policy; in blue) on the privacy policy with the participant's current gaze overlayed.

category of information. For the generic condition, the Cloze tests were the same for all participants. Participants had ten blanks to fill out in the first section, five in the middle section, and five in the last.

3.2.1.1 Procedure. Participants were told to install two apps in a randomised order, during which they encountered the two versions of privacy policies. The participants were not instructed to specifically read the

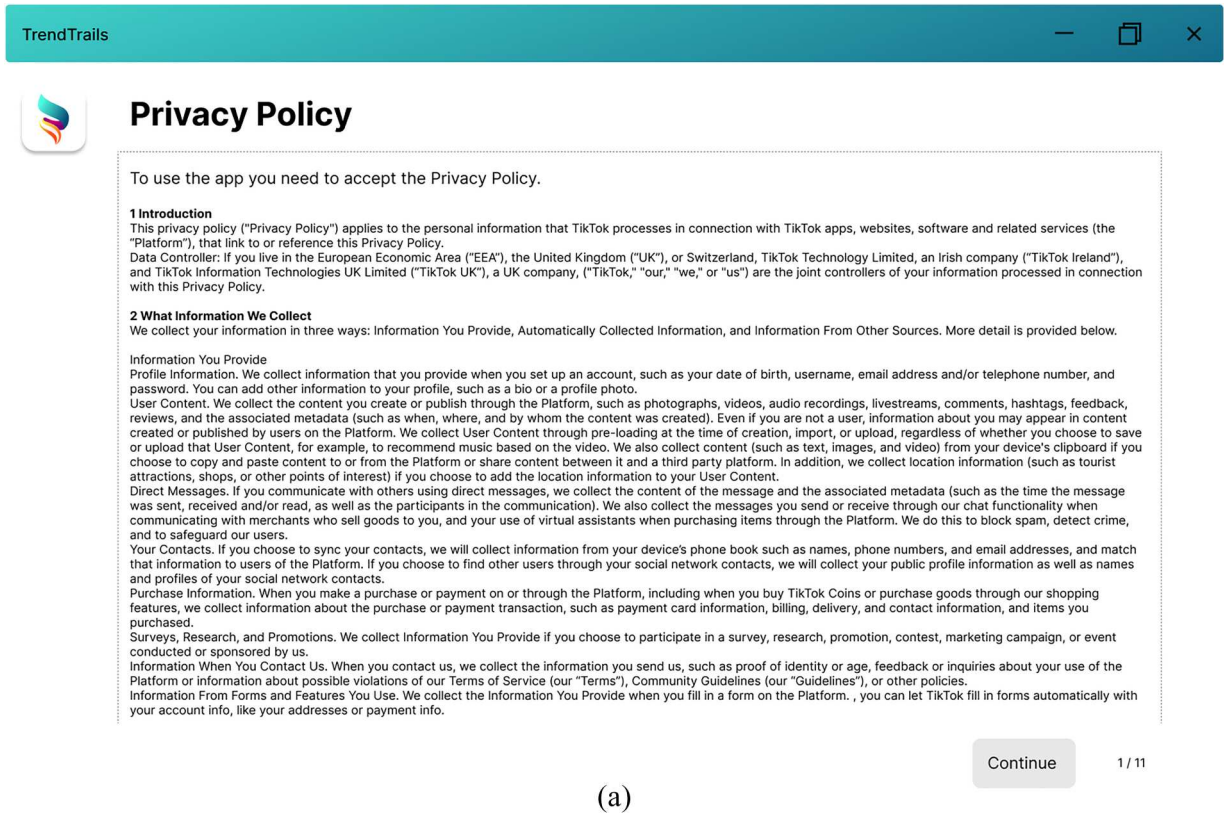

(a)

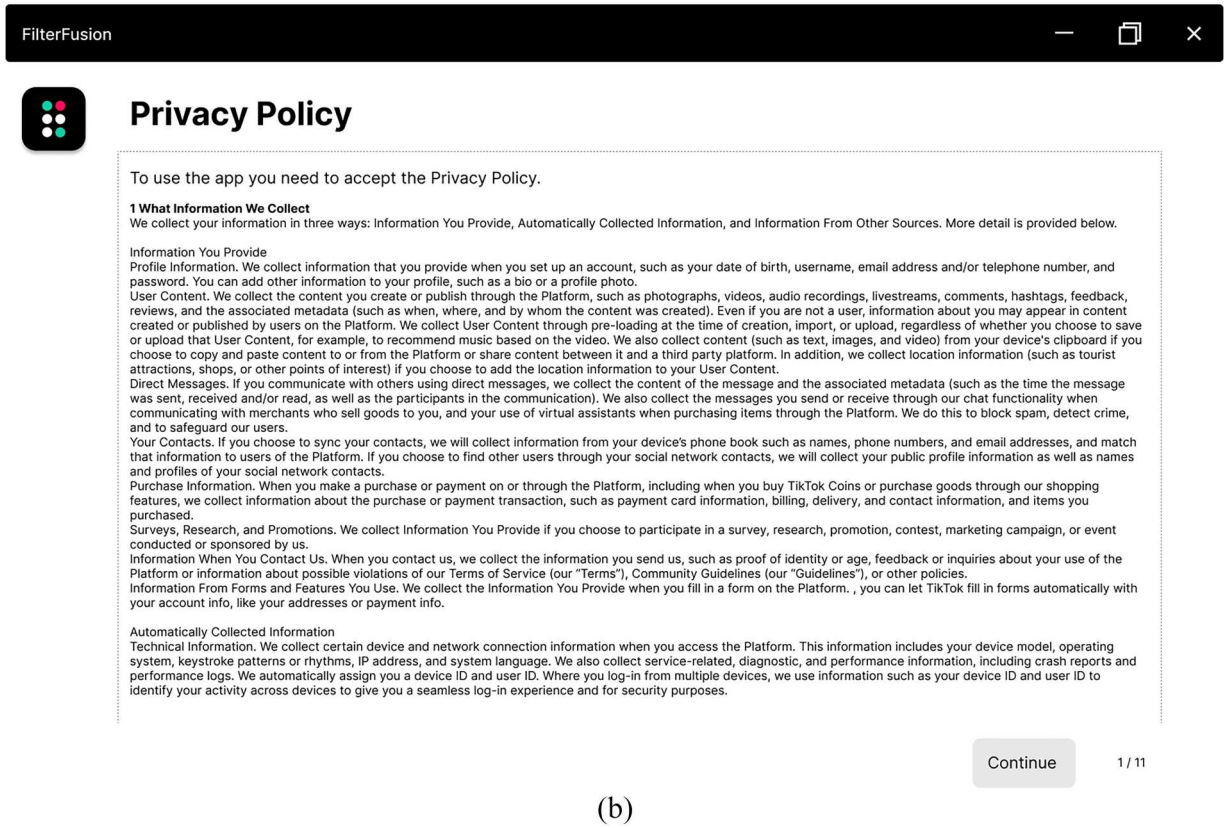

(b)

Figure 2. Screenshots from the prototype showing (a) the generic privacy policy as displayed in the generic condition, and (b) the personalised one as displayed in the personalised condition (here with 'What information we collect' as the participant's most concerned and thus first displayed section).

privacy policies, nor that their reading behaviour would be measured, but rather to go through the installation process 'as they normally would'. In the personalised condition, the first section of the privacy policy corresponded to a participant's most concerned category of information indicated in the pre-experiment survey, instead of the 'Introduction' section in the generic policy. Following the first installation, participants completed the NASA-TLX questionnaire (Hart and Staveland 1988) to assess their perceived workload. Then, they repeated the process with the second app. After completing both conditions, participants completed a post-experiment survey to maintain the study's cover story and to minimise the carryover effect on their next task. They then completed two Cloze tests corresponding to the order of the conditions. Finally, participants participated in a semi-structured interview (see Appendix D for the questions) to provide feedback on the perceived helpfulness of the personalised privacy policy, suggestions to make privacy policies more engaging, and potential disclosure alternatives to privacy policies. The study concluded with debriefing and compensation.

3.2.2.2 Data analysis. To analyse the eye tracking data, in iMotions, we defined important regions (e.g. the first section, the first page of a privacy policy, the overall privacy policy) of each prototype as the Areas of Interest (AOI) (see, e.g. Figure 1). Similar to previous research (see Section 3.1), we measured the average dwell times on these AOIs. We then used iMotions to calculate fixation-based dwell times with an I-VT-based (Salvucci and Goldberg 2000) fixation threshold of 30°/s and a dwell threshold of 100 ms. We use these fixation-based dwell times as an approximation for participants actually reading the text. The NASA TLX responses and Cloze tests were analysed via Python scripts.

The interview data were transcribed, anonymised and analysed using ATLAS.ti. Inductive qualitative content analysis was applied to these transcriptions to provide '*a description of patterns or regularities found in the data*' (Drisko and Maschi 2016, 86), where '*[b]oth manifest and latent content are examined, as are meanings in context*' (Drisko and Maschi 2016, 87–88). The code generation followed the steps by Mayring for inductive category development using qualitative content analysis (Mayring 2014, Figure 14). The first author and the third author initially coded three randomly selected interview transcripts independently. They then met to discuss and refine the codes and their definitions, agreeing on an initial codebook for further analysis. The first author, who designed the interview guide and conducted half of the interviews, proceeded to code the remaining transcripts, leveraging their familiarity with the interviews and their context. Code generation was iterative throughout the process. As new codes emerged, the first author added them to the shared code lists. All authors could examine the code lists and the first author's coding decisions at any time, hence reducing potential bias. Previously coded transcripts were revisited by the first author to maintain consistency in applying codes. During this process, the first author followed Braun and Clarke's (Braun and Clarke 2019, 592) emphasis on reflexivity, acknowledging "*the centrality of researcher subjectivity*" by documenting their own reflections to critically reflect on potential biases and assumptions.

3.3. Results

H2.1 (*Participants will spend more time reading personalised privacy policies than generic ones.*) To investigate H2.1, we compared the reading times of the full privacy policy, the section corresponding to their most concerned category of information, and the first page between the two conditions. Overall, participants did not spend more time reading the privacy policies in the personalised condition (M = 247.4s, SD = 216.2s) than in the generic condition (M = 235.4s, SD = 207.1s; see Figure 3(c)). Given that the assumption of normality was violated, we used a Wilcoxon signed-rank test which did not show a significant difference (W = 208, $p = .62$), with a rank-biserial correlation r = 0.44, suggesting a medium effect size. Since the main aspect of our personalisation was that the category a participant was most concerned about was presented as the first section, we compared the reading time on this section in the personalised condition with the same section in its original position in the generic condition. We found again no statistical significance using a paired t-test, given that the data were normally distributed ($t = -1.1$, $p = .28$, $d\ f = 29$, *Cohen's* $= -0.201$; Personalised: $M = 45.5s$, $SD = 49.9s$; Generic: $M = 35s$, $SD = 50.6s$).

We further compared the reading time on the first page between the two conditions to examine whether personalised privacy policies increase participants' initial interest in reading. When comparing the reading time of the Introduction section in the generic version (see Figure 2(a)) with that of an AOI of the same size in the personalised version, the results show that participants read this initial paragraph significantly longer ($W = 103$, $p = .0066$, $r = -0.83$) in the personalised version ($M = 12.6s$, $SD = 10.8s$) than in the generic one ($M = 7s$, $SD = 6.4s$; see Figure 3(a)). Yet, this effect does not persist for the reading time of the whole first page.

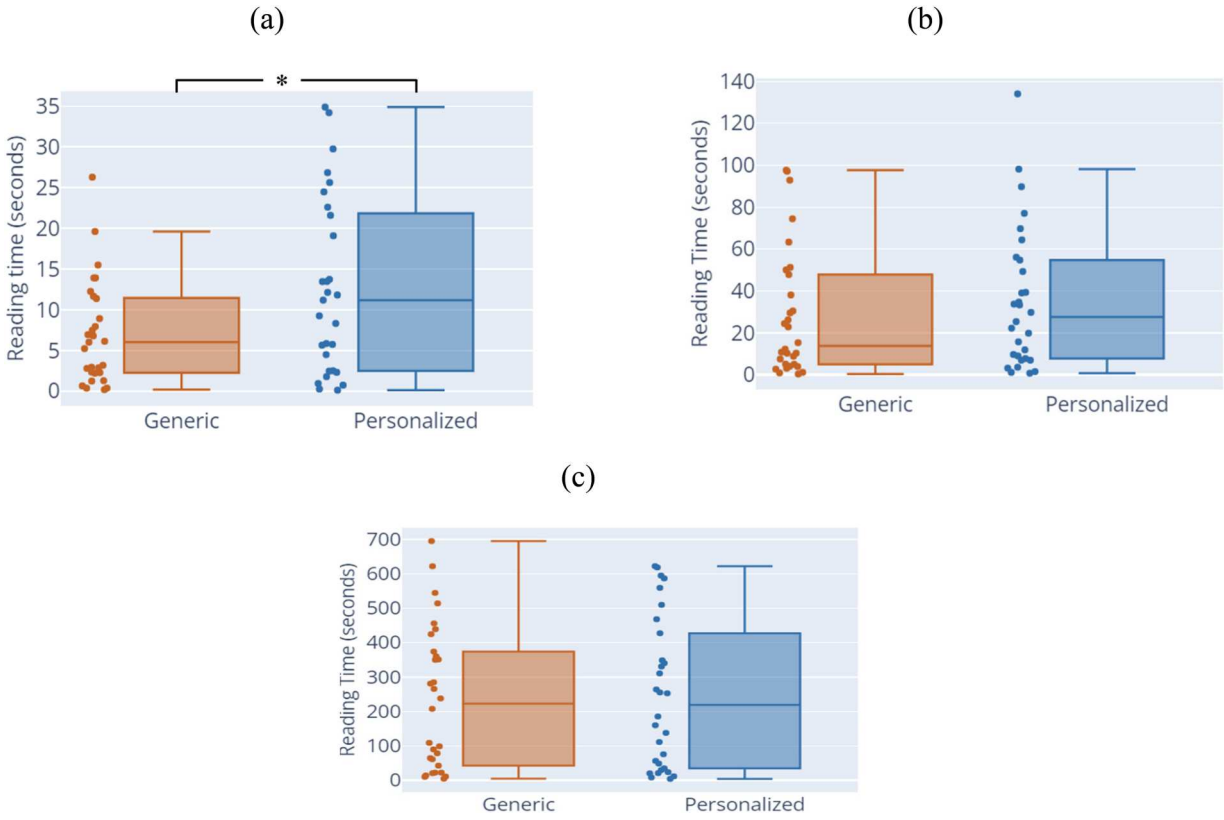

Figure 3. Boxplots showing (a) the reading times for the Introduction-section in the generic policy and an equally sized AOI in the personalised one, (b) the first page of the privacy policy, and (c) the full privacy policy. The asterisk denotes a significant difference ($p < .01$).

Here, participants spent a longer time reading the first page in the personalised condition ($M = 34.96s$, $SD = 33.3s$) than in the generic condition ($M = 28.26s$, $SD = 30.3s$; see Figure 3(b)), yet without statistical significance ($W = 187$, $p = .36$, $r = .402$).

H2.2 (*Participants' comprehension of the privacy policy's content will be higher in the personalised condition than in the generic condition.*) Since the most significant change in the personalised version occurred within the first section of the policy presented to participants, we compared the Cloze test results for this part as well as for the entire privacy policy. For the first section, only 7 out of 30 participants recalled more correct words from the personalised condition than the generic condition. Only 3 participants recalled more correct words from the personalised version when considering the entire policy. Wilcoxon signed-rank tests show that there is a statistically significant difference in correct word recalls between the personalised condition and the generic condition for both the first section (10 blanks to fill; Generic: $M = 4.7$, $SD = 1.6$, Personalised: $M = 3.9$, $SD = 1.8$; $W = 84$, $p = .031$, $r = .18$) and for all three sections (20 blanks to fill; Generic: $M = 12$, $SD = 2.9$, Personalised: $M = 9.1$, $SD = 3.8$; $W = 37.5$, $p < .001$, $r = .08$), indicating that participants recalled significantly more correct words in the generic condition than in the personalised condition. Thus, we can reject hypothesis H2.2 with high confidence. However, this finding should be interpreted as a condition-level difference (personalised vs. generic conditions) in Cloze performance, rather than as a strictly matched comparison of comprehension items.

H2.3 (*Participants will have a lower perceived workload when reading a personalised privacy policy than a generic one.*) Wilcoxon signed-rank tests applied to each of the six categories of the NASA TLX questionnaire, as well as to an overall non-weighted average index of the six different categories, did not show statistically significant results on either of the six categories or on the overall index (see Table 7). Therefore, H2.3 is rejected.

3.3.1 Interviews

During the interviews, only three participants (P13, P21, P23) out of thirty successfully identified the difference between the two conditions. After explaining the true purpose of our experiment, we asked participants for their perceptions. Overall, the feedback was positive,

Table 7. Average scores and standard deviations for the six dimensions of the NASA-TLX (on a scale from 0 to 21) with the associated *p*-value and test statistic *W* for Wilcoxon signed-rank tests.

Dimension	Generic Mean (SD)	Pers. Mean (SD)	W	*p*
Mental	8.20 (7.02)	8.50 (6.72)	126.0	.99
Physical	3.00 (4.31)	3.23 (4.35)	50.5	.90
Temporal	5.30 (5.19)	5.63 (4.44)	79.5	.34
Performance	16.57 (5.14)	17.23 (3.82)	107.5	.53
Effort	6.30 (5.94)	6.70 (6.06)	83.5	.64
Frustration	5.47 (5.51)	6.53 (6.69)	80.5	.55
Overall Index	7.47 (3.32)	7.97 (3.38)	151.0	.15

with most participants (n = 21) finding personalised reordering to be helpful. However, participants also expressed two main concerns (n = 14): issues with personalisation and issues with privacy policies. For the first cluster of concerns, participants expressed worries about personalisation being too intrusive and raised practical questions about how personalised privacy policies would function. They questioned whether personalisation would be applied consistently across devices and, in the case of shared devices, whose preferences would be displayed. Additionally, participants expressed concern that companies might exploit privacy preferences to the users' disadvantage.

For the latter cluster, feedback centred on how and why our personalised design did not address the existing issues with privacy policies or did not change participants' reading habits. Participants felt that even with personalisation, the privacy policy would still be too long, and users would still lack the option to download an app if they disagreed with the data processing activities disclosed in a privacy policy. Moreover, those who typically do not read privacy policies would still be unlikely to read or care about the tailored information. Among the participants who found personalisation helpful, 16 participants emphasised that it would allow them to focus on what is most important to them right from the start of a privacy policy. Another point highlighted was how personalised privacy policies could help users retain focus when their attention span is limited. Participants (n = 7) believed they concentrate best at the beginning of a privacy policy, so placing the most concerned section at the start would help them absorb the information more effectively. Furthermore, participants (n = 6) expressed their perception of personalisation as a way to facilitate reading, as they noted that they would be more likely to read the sections they care about most if those sections were placed at the beginning when their focus is sharpest.

Reading behaviours of the participants were nuanced. For those (n = 16) who shared their typical reading behaviours unprompted, half (n = 8) admitted not reading privacy policies in real-life. Three participants mentioned partial reading, influenced by factors such as the importance of specific sections of a policy (P21) or specific concerns, like the use of images (P31). The rest indicated that they would read a privacy policy in full, if the company or the app is (un)trustworthy, a paid service, a sensitive app like banking, or appears shady; or when they need to provide their real data instead of fake data. These circumstances would motivate participants to scrutinise privacy policies more closely. Otherwise, participants often chose not to read privacy policies because they felt compelled to accept the conditions if they wanted to use the service, with no ability to modify the terms.

We asked participants for suggestions to make privacy policies more engaging, aside from personalisation, especially in light of how they described the language used in privacy policies as difficult, complex, and containing technical terms. The top three recommendations were shortening the text, highlighting, and using simpler language. Participants also suggested collapsible content, summarised headings, a table of contents, and the use of icons or ratings, although one participant noted that an 'F' rating without context is unhelpful to users. Other ideas included forced interactions to ensure users go through the entire policy and layered privacy policies. P17 explained the need for a layered privacy policy: '*And because I think there are two purposes [for privacy policies]: the legal purpose and the customer purpose. And they should be really separated. There should be a short version and then one [version] that you can read if you want that for it to be optional*'.

Three participants also recommended using filter questions to reduce the information presented in a privacy policy, similar to the pre-experiment questionnaire we conducted. However, unlike in our approach, they preferred receiving only relevant information. As P1 explained, using the example of minors:

> Or if you are installing, you could already tell them that you have to be 18 years old or 13 years old. Otherwise, you cannot go through the installation. Because you are already coming down to the privacy policy, you see this 13-year-old thing. It's a bit weird.

Further recommendations included providing clear explanations, summarised information, and improved visualisations of privacy policies (e.g. images, videos, or dashboards). Participants expressed the need for videos to be kept short, ideally between 30 s and two minutes. ChatGPT was recommended for explanations, creating summaries, flagging concerning data practices, or displaying important information to users. A few

participants (n = 8) advocated for an initial layer of privacy disclosure through app marketplaces like the App Store. Lastly, participants (n = 4) also expressed their desire to have active control over their privacy, rather than simply being informed about terms without the ability to make changes. They suggested mechanisms such as checkboxes, similar to cookie banners, where they could exercise their choices regarding a company's data collection, processing, and sharing practices. Additionally, they expressed the need for just-in-time notices that prompt users to consent to or refuse specific activities as they occur.

3.4. Discussion

The results of the eye-tracking experiment present a mixed picture of user engagement with personalised privacy policies. While most participants subjectively found our personalisation approach helpful, their eye-tracking data, comprehension, and perceived workload results do not confirm a higher engagement beyond the beginning of the privacy policy.

Participants read the first few lines in a personalised policy, which correspond to the length of the Introduction section in the generic version, which is significantly longer. This effect might be attributed to the fact that the information presented was of particular interest to the participants because it matched their primary privacy concern. This would mean that our personalised approach only initially increases user engagement. An alternative explanation is that this is a novelty effect, whereby participants expected a privacy policy to begin with the 'Introduction' section and therefore continued reading due to the discrepancy between their expectations and the actual content. The presented way of personalising privacy policies might thus be too subtle to bring a substantial change in user engagement. Thus, future research could investigate combined approaches where a legally binding privacy policy is combined with a more interactive form of disclosure, such as privacy labels (Kelley et al. 2009), summaries that are presented before the text of each section (Woodring, Perez, and Ali-Gombe 2024), or sequential context (Masotina and Spagnolli 2022), to name just a few. However, privacy labels have drawn complaints for their use of incomprehensible terms and lack of details to support informed decision-making (Balash et al. 2024). Any combined approach to disclosure should thus capitalise on the strengths of each composing method without exacerbating their weaknesses. It is also possible that no single disclosure method can fully satisfy every user's needs. Therefore, in addition to the ongoing pursuit of technological and design solutions that enhance user knowledge and engagement, which aim to ultimately improve privacy control and management, other strategies should also be considered. One strategy could involve integrating online privacy literacy education that reinforces disclosure methods by providing users with a stronger foundational understanding of privacy. Online privacy literacy, defined by Trepte et al. (2015, p.339) , consists of '*a combination of factual or declarative ('knowing that') and procedural ('knowing how') knowledge about online privacy*'. Current research seems to focus on the connection between privacy literacy and an individual's social privacy protection practices, i.e. how they protect their privacy while using a service or interacting with fellow users (Bartsch and Dienlin 2016; Choi 2022); we propose a different direction for future research into whether online privacy literacy could influence people's assessment of disclosures presented to them and their decisions in consenting to the data practices of service providers.

The findings from the eye-tracking data further confirm previous studies on people's reading behaviours regarding privacy policies in terms of the general skipping and skimming, as well as individual differences in reading behaviours. In our study, participants read the privacy policies on average much longer (generic: 235.4s, personalised 247.4s; 6453 words) than in previous studies (M = 53s for 451 words [Steinfeld 2016]; M = 73s for 7977 words [Obar and Oeldorf-Hirsch 2020]). Yet, given that adults' average silent reading speed is 238 words per minute (Brysbaert 2019), participants would have needed around 27 min to properly read the full policy, meaning that they did not spend a meaningful amount of time on the privacy policies. We also observed large differences in reading time between the participants regardless of the conditions in the experiment, ranging from around five seconds to 11 min, which is in line with previous research (Obar and Oeldorf-Hirsch 2020; Steinfeld 2016). These findings highlight the inherent tension between maintaining the legal requirements of privacy policies and achieving user-friendly and more engaging text. In addition, they demonstrate the need for privacy policies and privacy disclosure in general to factor in differences in individual reading behaviour.

Participants have performed significantly better in the Cloze tests for the generic condition than the personalised condition for the first sections. The difference in performance was likely because it was probably easier to deduce the answers for the first section from the context in the generic condition than from the personalised condition, considering it contained many country names, or due to users' previous knowledge and

familiarity with the Introduction sections in other privacy policies.

Participants in both conditions have a similar degree of perceived workload. Our results could be attributable to the total length of the policies, which arguably vanishes the positive effect of reordering the sections based on users' primary concerns. In addition, objective workload measures such as Pupillometry (Kiefer et al. 2016; Kosch et al. 2023) could be considered as an option for future work to capture more subtle differences.

3.5. Limitations

In both conditions, the same privacy policy from TikTok was displayed, which may have caused a carryover effect. However, in the two conditions, the privacy policies were associated with different apps which were presented with visually different designs to mitigate this effect. The reordering of sections in personalised privacy policies led to different paragraph breaks from those in the generic ones, and each prototype app had a distinct colour scheme and design, further suggesting the impression of two distinct privacy policies. However, we acknowledge that the results of the Cloze tests might have been affected by the use of the same privacy policy twice. The Cloze tests were further limited due to concerns regarding their comparability. For this reason, the Cloze findings should be interpreted cautiously as an exploratory, condition-level measure of comprehension. Future work could strengthen the design by using parallel difficulty-matched comprehension forms, a between-subject design or the use of two different but comparable privacy policies to provide a comparison for the overall eye-tracking experiment.

The privacy policies were intentionally presented with minimal design elements, such as the absence of colours or highlighting with bold fonts. While this approach was chosen to control for confounding effects, it may have reduced participants' motivation to engage with the policies. Additionally, the reordering of sections in personalised privacy policies resulted in a somewhat varied design among participants with different primary concerns, which might have influenced the participants' reading behaviour. Furthermore, the sample population may have been subject to potential biases, as it is comprised predominantly of young men with a high educational background. Moreover, we opted for a paginated policy instead of putting all text on one page with scrolling because this format allows more precise analysis of the eye-tracking data. We acknowledge that this presentation of the policy and the laboratory environment may have led some participants to show an adapted gaze behaviour and to read the privacy policies more carefully than they would do in a more natural setting. To mitigate such effects to some extent, we instructed participants to follow the app installation 'as they normally would'. The privacy policy was displayed on a desktop screen for the sake of more accurate eye-tracking data collection. This may limit the ecological validity of the findings, as many users in real life encounter privacy policies also on smartphones or tablets. Future research could explore whether users' engagement with privacy policies differs across device contexts.

4. Conclusion

In this paper, we explored an approach to personalising privacy policies that simultaneously balances privacy policies' dual functions as legal documents mandated by law and as information sources for users. We evaluated this approach in two studies: to investigate whether personal characteristics could be used as proxies to predict a person's most concerned category of information in a privacy policy, and to determine whether our proposed personalisation approach would increase user engagement and comprehension while reducing workload.

Our first study indicates that personal characteristics alone are not sufficient to predict a person's most concerned category of information in a privacy policy. Using multinomial logistic regression, our initial findings show that neuroticism is a significant predictor. However, neuroticism was not significant after the Holm adjustment to control Type I errors due to multiple comparisons. In addition, the results from comparing section headings and one-sentence summarised statements further suggest that for users to make *informed* decisions about their most concerned category of information, they need at least a basic awareness of the content in each section of a privacy policy. Additionally, we recommend that privacy policies from different service providers follow a consistent structure to make it easy for users to locate their concerned categories of information across various services. If future research chooses to use inferences for personalisation in privacy disclosure or privacy management in general, it needs to make sure that such inferences account for multiple factors, moving beyond mere personal characteristics.

Our second study reveals that personalising the order of a privacy policy did not, in general, increase user engagement and comprehension or reduce perceived workload. However, we found that our personalisation approach leads to higher engagement with the first

few sentences of a privacy policy. Further research could study how this engagement might be retained longer. Or conversely, further research is needed to investigate how the content of privacy policies can be communicated through concise, legally compliant summaries that align with users' limited attention span and reading time.

One possible explanation of the null findings of our second study is that the challenge may stem from data controllers' disclosure compliance practices, within which personalisation needs to operate. The norm of our current paradigm often associates the quantity of information with good transparency and thus good disclosure to users (Oehler and Wendt 2017) as well as with good compliance (Ding and Huang 2024). However, the resulting complex and lengthy privacy policies, as shown in the second study, not only limits the effectiveness of measures such as personalisation but also disregards the reality that average users have limited attention spans they are willing to devote to reading privacy disclosures. In addition, the null findings suggest that the usability problems of privacy policies cannot be fully addressed through presentation improvements alone if compliance is expected with exhaustive disclosure, which is at the expense of effective and usable transparency for users. From this perspective, our findings indicate that current compliance-oriented disclosure practices may constrain the practical success of user-centred privacy designs. We recommend that relevant stakeholders, such as regulators and data protection authorities, provide clearer guidance and greater legal certainty for more user-centred transparency methods such as personalised transparency.

Feedback from the interviews shows a reason for users' withdrawal from privacy policies: Participants felt they had little control over the data collection and processing activities described in privacy policies. They perceived themselves as being in a 'take it or leave it' situation – if they wanted to use a service, they had no choice but to consent, or else they couldn't use it. As a result, whether to engage with the privacy policy or not did not alter their circumstances.

The discrepancy between participants' positive perceptions of the personalised reordering design and the null objective findings of our eye-tracking experiment may reflect a gap between perceived relevance and usefulness of legally compliant personalisation and the actual user engagement with the design for privacy disclosure. Many participants in the interviews viewed personalised reordering positively because, in principle, the design makes a privacy policy feel more useful and easier to navigate. However, as this design aims to preserve the full legally required text and change only the order of content, it may not have been sufficient to substantially alter reading behaviour for users, if users do not read in the first place.

Beyond the immediate findings of our studies, these results also raise broader questions about whether more modular and customisable forms of privacy choices could increase the practical relevance of privacy disclosures. This lack of modularity in choices has been tackled by previous scholars (e.g. Das et al. 2018), with some proposing from the legal perspective a right to customisation (Tamò-Larrieux, Zihlmann, and Garcia 2021). Under the right to customisation, users could selectively consent to specific data processing practices outlined in a privacy policy and opt out of those that do not align with their privacy preferences or concerns, such as sharing personal information with certain third parties. The service provider would then offer a version of the service per these preferences. In this scenario, a privacy policy would serve as more than just a compliance tool for companies or a shield for liability; it would become a meaningful document for individuals, incentivising them to engage with it and exert control over their data. It would also clear the legal impediment to applying existing HCI designs and prototypes in the real world, allowing users to have granular consent and better control, e.g. through the usage of personalised privacy assistants (PPAs; i.e. '*intelligent agents capable of learning the privacy preferences of their users over time, semi-automatically configuring many settings, and making many privacy decisions on their behalf*') (Personalized Privacy Assistant Project, accessed in 2024). PPAs could serve as one possible approach when the door of compliance is widened, envisioned to streamline privacy disclosures with personalisation, provide personalised recommendations, and facilitate user decision-making (Chang and Barber 2023; Das et al. 2018; Liu et al. 2016; Marky et al. 2024; Morel, Iwaya, and Fischer-Hübner 2025; Xu et al. 2026). Additionally, studies have examined various levels of automation that PPAs can employ, including fully automated decision-making by the agent (Colnago et al. 2020). If the right to customisation were to be implemented, a personalised privacy assistant could be instrumental in motivating and enabling users to better understand the data processing activities and control their personal information before they sign up for a service. This would ensure that these activities align with and truly reflect a user's privacy concerns and preferences. However, as discussed in Xu et al. (Xu, Rossi, and Tamò-Larrieux 2025), such technologies should ensure their accuracy, respect user agency and take accountability measures.

In addition to agentic systems such as PPAs, other privacy designs utilising personalisation have also shown promising results in usable privacy research (Fischer-Hübner and Karegar 2024), such as to enhance privacy communication and user decision-making (see, for example, Harbach et al. 2014; Peer et al. 2020). However, when such designs are translated into real-world applications, they may encounter legal and regulatory barriers related to compliance. We therefore encourage policymakers, data protection authorities, and legal scholars to consider whether and under what conditions these approaches can be implemented lawfully for end users.

Last but not least, to gain a better understanding of how engagement with privacy policies is affected by the presence or absence of meaningful user control, we recommend that future research explore whether users' engagement with privacy policies changes when they are granted control over their personal information and its processing by a service provider, with immediate adjustments to the service based on their choices.

Author contributions

CRediT: **Meihe Xu:** Conceptualization, Data curation, Formal analysis, Funding acquisition, Investigation, Methodology, Resources, Software, Validation, Visualization, Writing – original draft, Writing – review & editing; **Jannis Strecker-Bischoff:** Conceptualization, Data curation, Formal analysis, Investigation, Methodology, Resources, Software, Visualization, Writing – original draft, Writing – review & editing; **Clement Guitton:** Conceptualization, Formal analysis, Investigation, Methodology, Resources, Software, Validation, Writing – review & editing; **Kenan Bektas:** Conceptualization, Methodology, Resources, Software, Validation, Writing – review & editing; **Aurelia Tamo-Larrieux:** Conceptualization, Funding acquisition, Project administration, Supervision, Writing – review & editing; **Simon Mayer:** Project administration, Resources, Supervision, Writing – review & editing.

Disclosure statement

No potential conflict of interest was reported by the author(s).

References

Acquisti, Alessandro, Laura Brandimarte, and George Loewenstein. 2015. "Privacy and Human Behavior in the age of Information." *Science (New York, N.Y.)* 347 (6221): 509–514. https://doi.org/10.1126/science.aaa1465.

Adhikari, Kishalay, and Rajeev Kumar Panda. 2018. "Users' Information Privacy Concerns and Privacy Protection Behaviors in Social Networks." *Journal of Global Marketing* 31 (2): 96–110. https://doi.org/10.1080/08911762.2017.1412552.

Balash, David G., M. M. Ali, Chris Kanich, and Adam J. Aviv. 2024. "I Would Not Install an app with This Label": Privacy Label Impact on Risk Perception and Willingness to Install iOS Apps. 413–432. Retrieved January 29, 2025 from https://www.usenix.org/system/files/soups2024-balash.pdf.

Bansal, Gaurav, and David Gefen. 2010. "The Impact of Personal Dispositions on Information Sensitivity, Privacy Concern and Trust in Disclosing Health Information Online." *Decision Support Systems* 49 (2): 138–150. https://doi.org/10.1016/j.dss.2010.01.010

Bansal, Gaurav, Fatemeh Mariam Zahedi, and David Gefen. 2016. "Do Context and Personality Matter? Trust and Privacy Concerns in Disclosing Private Information Online." *Information & Management* 53 (1): 1–21. https://doi.org/10.1016/j.im.2015.08.001.

Barlow, David H., Kristen K. Ellard, Shannon Sauer-Zavala, Jacqueline R. Bullis, and Jenna R. Carl. 2014. "The Origins of Neuroticism." *Perspectives on Psychological Science: A Journal of the Association for Psychological Science* 9 (5): 481–496. https://doi.org/10.1177/1745691614544528

Bartsch, Miriam, and Tobias Dienlin. 2016. "Control Your Facebook: An Analysis of Online Privacy Literacy." *Computers in Human Behavior* 56:147–154. https://doi.org/10.1016/j.chb.2015.11.022.

Baruh, Lemi, Ekin Secinti, and Zeynep Cemalcilar. 2017. "Online Privacy Concerns and Privacy Management: A Meta-analytical Review." *The Journal of Communication* 67 (1): 26–53. https://doi.org/10.1111/jcom.12276.

Bawden, D., and L. Robinson. 2020. "Information Overload: An Overview." In *Oxford Encyclopedia of Politics*, edited by E. Hannah. New York, NY: Oxford University Press. https://doi.org/10.1093/acrefore/9780190228637.013.1360.

Ben-Shahar, Omri, and Ariel Porat. 2021. *Personalized Law: Different Rules for Different People*. New York: Oxford University Press.

Bergström, Annika. 2015. "Online Privacy Concerns: A Broad Approach to Understanding the Concerns of Different Groups for Different Uses." *Computers in Human Behavior* 53:419–426. https://doi.org/10.1016/j.chb.2015.07.025.

Bhatia, Jaspreet, and Travis D. Breaux. 2018. "Empirical Measurement of Perceived Privacy Risk." *ACM Transactions Computer-Human Interact* 25 (6): 1–47. https://doi.org/10.1145/3267808.

Boerman, Sophie C., Sanne Kruikemeier, and Frederik J. Zuiderveen Borgesius. 2021. "Exploring Motivations for Online Privacy Protection Behavior: Insights from Panel Data." *Communication Research* 48 (7): 953–977. https://doi.org/10.1177/0093650218800915

Bormuth, John R. 1968. "Cloze Test Readability: Criterion Reference Scores." *Journal of Educational Measurement* 5 (3): 189–196. https://doi.org/10.1111/j.1745-3984.1968.tb00625.x

Braun, Virginia, and Victoria Clarke. 2019. "Reflecting on Reflexive Thematic Analysis." *Qualitative Research in Sport, Exercise and Health* 11 (4): 589–597. https://doi.org/10.1080/2159676X.2019.1628806

Broeck, Evert Van den, Karolien Poels, and Michel Walrave. 2015. "Older and Wiser? Facebook Use, Privacy Concern, and Privacy Protection in the Life Stages of Emerging, Young, and Middle Adulthood." *Social Media + Society* 1 (2): 2056305115616149. https://doi.org/10.1177/2056305115616149.

Brysbaert, Marc. 2019. "How Many Words Do We Read per Minute? A Review and Meta-analysis of Reading Rate."

Journal of Memory and Language 109:104047. https://doi.org/10.1016/j.jml.2019.104047

Busch. 2019. Implementing Personalized Law. The University of Chicago Law Review. University of Chicago. Law School. Retrieved from https://www.jstor.org/stable/26590557

California Consumer Privacy Act of 2018. Retrieved from https://www.oag.ca.gov/privacy/cpra.

Chang, Cheng, Huaxin Li, Yichi Zhang, Suguo Du, Hui Cao, and Haojin Zhu. 2019. "Automated and Personalized Privacy Policy Extraction under GDPR Consideration. In Wireless Algorithms." *Systems, and Applications* 11604: 43–54. https://doi.org/10.1007/978-3-030-23597-0_4.

Chang, Kai-Chih, and Suzanne Barber. 2023. "Personalized Privacy Assistant: Identity Construction and Privacy in the Internet of Things." *Entropy* 25 (5): 717. https://doi.org/10.3390/e25050717.

Chen, Baiqi, Tingmin Wu, Yanjun Zhang, Mohan Baruwal Chhetri, and Guangdong Bai. 2023. "Investigating Users' Understanding of Privacy Policies of Virtual Personal Assistant Applications." In Proceedings of the 2023 ACM Asia Conference on Computer and Communications Security (ASIA CCS '23), 65–79. https://doi.org/10.1145/3579856.3590335

Chen, Shi-Yi, Zhe Feng, and Xiaolian Yi. 2017. "A General Introduction to Adjustment for Multiple Comparisons." *Journal of Thoracic Disease* 9 (6): 1725. https://doi.org/10.21037/jtd.2017.05.34

Chen, Xiaogang, and Yuhui Zhang. 2021. "The Construct of Information Privacy Concerns in the Chinese Cultural Setting." *Nankai Business Review International* 12 (1): 42–55. https://doi.org/10.1108/nbri-12-2019-0071.

Choi, Soeyoon.. 2022. "Privacy Literacy on Social Media: Its Predictors and Outcomes." *International Journal of Human–Computer Interaction* 39 (1): 217–232. https://doi.org/10.1080/10447318.2022.2041892.

Colnago, Jessica, Yuanyuan Feng, Tharangini Palanivel, Sarah Pearman, Megan Ung, Alessandro Acquisti, Lorrie Faith Cranor, and Norman Sadeh. 2020. "Informing the Design of a Personalized Privacy Assistant for the Internet of Things." In Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (CHI '20), 1–13. https://doi.org/10.1145/3313831.3376389

Costa, P. T., and R. R. McCrae. 1999. "A Five-Factor Theory of Personality." *Handbook of Personality: Theory and Research* 2 (01): 1999.

Das, Anupam, Martin Degeling, Daniel Smullen, and Norman Sadeh. 2018. "Personalized Privacy Assistants for the Internet of Things: Providing Users with Notice and Choice." *IEEE Pervasive Computing* 17 (3): 35–46. https://doi.org/10.1109/MPRV.2018.03367733

Das, Sanchari, Jayati Dev, and Kaushik Srinivasan. 2018. "Modularity Is the Key a New Approach to Social Media Privacy Policies." In Proceedings of the 7th Mexican Conference on Human-Computer Interaction (MexIHC '18). https://doi.org/10.1145/3293578.3293589

Degeling, Martin, Christine Utz, Christopher Lentzsch, Henry Hosseini, Florian Schaub, and Thorsten Holz. 2019. "We Value Your Privacy … Now Take Some Cookies." *Informatik-Spektrum* 42 (5): 345–346. https://doi.org/10.1007/s00287-019-01201-1.

Ding, Xiaodong, and Hao Huang. 2024. "For Whom Is Privacy Policy Written? A new Understanding of Privacy Policies." *Computer law and Security Report* 55 (106072): 106072. https://doi.org/10.1016/j.clsr.2024.106072.

Drisko, James W., and Tina Maschi. 2016. *Content Analysis*. New York: Oxford University Press.

Ebers, Martin, and Karin Sein. 2024. "Data-Driven Technologies: Challenges for Privacy and EU Data Protection Law." In *Privacy, Data Protection and Data-Driven Technologies*, edited by M. Ebers and K. Sein. London: Routledge. https://doi.org/10.4324/9781003502791-2.

Ebert, Nico, Kurt Alexander Ackermann, and Peter Heinrich. 2020. "Does Context in Privacy Communication Really Matter? A Survey on Consumer Concerns and Preferences." In Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (CHI '20), 1–11. https://doi.org/10.1145/3313831.3376575

Ebert, Nico, Kurt Alexander Ackermann, and Björn Scheppler. 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 (CHI '21), 1–12. https://doi.org/10.1145/3411764.3445516

European Parliament and Council of the European Union. 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 (General Data Protection Regulation). Official Journal of the European Union 59, L119 (April 2016), 1–88. Retrieved from https://eur-lex.europa.eu/eli/reg/2016/679/oj.

Faja, Silvana, and Silvana Trimi. 2008. "Privacy Concerns in e-Commerce: An Empirical Investigation of Gender Differences." *International Journal of Electronic Business* 6 (4): 386–404. https://doi.org/10.1504/IJEB.2008.020676.

Fan, H., and M. S. Poole. 2006. "What Is personal." tion? Perspectives on the Design and Implementation of Personalization in Information Systems." *Journal of Organizational Computing* 16 (3-4): 179–202. https://doi.org/10.1080/10919392.2006.9681199.

Fischer-Hübner, Simone, and Farzaneh Karegar. 2024. *Addressing Challenges: A Way Forward. In: The Curious Case of Usable Privacy. Synthesis Lectures on Information Security, Privacy, and Trust*, 133–160. Cham: Springer. https://doi.org/10.1007/978-3-031-54158-2_5.

Gellert, Anna S., and Carsten Elbro. 2013. "Cloze Tests May Be Quick, but Are They Dirty? Development and Preliminary Validation of a Cloze Test of Reading Comprehension." *Journal of Psychoeducational Assessment* 31 (1): 16–28. https://doi.org/10.1177/0734282912451971.

Greene, Benjamín. 2001. "Testing Reading Comprehension of Theoretical Discourse with Cloze." *Journal of Research in Reading* 24 (1): 82–98. https://doi.org/10.1111/1467-9817.00134

Groß, Thomas. 2021. "Validity and Reliability of the Scale Internet Users' Information Privacy Concerns (IUIPC)." *Proceedings on Privacy Enhancing Technologies* 2021 (2): 235–258. https://doi.org/10.2478/popets-2021-0026.

Harbach, Marian, Markus Hettig, Susanne Weber, and Matthew Smith. 2014. "Using Personal Examples to Improve Risk Communication for Security & Privacy Decisions." In Proceedings of the SIGCHI Conference on

Human Factors in Computing Systems. https://doi.org/10.1145/2556288.2556978

Hart, Sandra G., and Lowell E. Staveland. 1988. "Development of NASA-TLX (Task Load Index): Results of Empirical and Theoretical Research." In *Advances in Psychology*. Vol. 52, edited by Peter A. Hancock and Najmedin Meshkati, 139–183. North-Holland: Elsevier. https://doi.org/10.1016/S0166-4115(08)62386-9.

Hoofnagle, Chris Jay, Jennifer King, Su Li, and Joseph Turow. 2010. How Different Are Young Adults from Older Adults When It Comes to Information Privacy Attitudes and Policies? Available at SSRN 1589864. https://doi.org/10.2139/ssrn.1589864

Ibdah, Duha, Nada Lachtar, Satya Meenakshi Raparthi, and Anys Bacha. 2021. ""Why Should I Read the Privacy Policy, I Just Need the Service": A Study on Attitudes and Perceptions toward Privacy Policies." *IEEE Access* 9:166465–166487. https://doi.org/10.1109/ACCESS.2021.3130086

International Labour Organization. 2024. Classification of Education (ISCED). Retrieved July 12, 2025 from https://ilostat.ilo.org/methods/concepts-and-definitions/classification-education/.

Janic, Milena, Jan Pieter Wijbenga, and Thijs Veugen. 2013. Transparency Enhancing Tools (TETs): An Overview. In 2013 Third Workshop on Socio-Technical Aspects in Security and Trust, 18–25. https://doi.org/10.1109/STAST.2013.11

John, Oliver P., E. M. Donahue, and R. L. Kentle. 1991. *The Big Five Inventory – versions 4a and 5*. Berkeley: University of California, Berkeley, Institute of Personality and Social Research.

John, Oliver P., Laura P. Naumann, and Christopher J. Soto. 2008. "Paradigm Shift to the Integrative big Five Trait Taxonomy." *Handbook of Personality: Theory and Research* 3 (2): 114–158.

Junglas, Iris A., Norman A. Johnson, and Christiane Spitzmüller. 2008. "Personality Traits and Concern for Privacy: An Empirical Study in the Context of Location-Based Services." *European Journal of Information Systems* 17 (4): 387–402. https://doi.org/10.1057/ejis.2008.29.

Kelley, Patrick Gage, Joanna Bresee, Lorrie Faith Cranor, and Robert W. Reeder. 2009. "A "Nutrition Label" for Privacy." In Proceedings of the 5th Symposium on Usable Privacy and Security (SOUPS '09), 1–12. https://doi.org/10.1145/1572532.1572538

Kiefer, Peter, Ioannis Giannopoulos, Andrew Duchowski, and Martin Raubal. 2016. "Measuring Cognitive Load for Map Tasks through Pupil Diameter." In *Geographic Information Science. GIScience 2016. Lecture Notes in Computer Science, vol. 9927*, edited by J. Miller, D. O'Sullivan, and N. Wiegand, 323–337. Cham: Springer. https://doi.org/10.1007/978-3-319-45738-3_21.

Kitkowska, Agnieszka, Yefim Shulman, Leonardo A. Martucci, and Erik Wästlund. 2023. "Designing for Privacy: Exploring the Influence of Affect and Individual Characteristics on Users' Interactions with Privacy Policies." *Computers & Security* 134:103468. https://doi.org/10.1016/j.cose.2023.103468.

Kitkowska, Agnieszka, Mark Warner, Yefim Shulman, Erik Wästlund, and Leonardo A. Martucci. 2020. "Enhancing Privacy through the Visual Design of Privacy Notices: Exploring the Interplay of Curiosity, Control and Affect." In Sixteenth Symposium on Usable Privacy and Security (SOUPS 2020), 437–456.

Knijnenburg, Bart P., Reza Ghaiumy Anaraky, Daricia Wilkinson, Moses Namara, Yangyang He, David Cherry, and Erin Ash. 2022. User-Tailored Privacy. In Modern Socio-Technical Perspectives on Privacy. Springer, Cham, 367–393. Retrieved from https://library.oapen.org/bitstream/handle/20.500.12657/52825/978-3-030-82786-1.pdf?sequence=1#page=365.

Korzaan, Melinda L., and Katherine T. Boswell. 2008. "The Influence of Personality Traits and Information Privacy Concerns on Behavioral Intentions." *Journal of Computer Information Systems* 48 (4): 15–24. https://doi.org/10.1080/08874417.2008.11646031.

Kosch, Thomas, Jakob Karolus, Johannes Zagermann, Harald Reiterer, Albrecht Schmidt, and Paweł W Woźniak. 2023. "A Survey on Measuring Cognitive Workload in Human-Computer Interaction." *ACM Computing Surveys* 55:13s. https://doi.org/10.1145/3582272.

Kununka, Sophia, Nikolay Mehandjiev, Pedro Sampaio, and Konstantina Vassilopoulou. 2017. "End User Comprehension of Privacy Policy Representations." In *End-User Development. IS-EUD 2017. Lecture Notes in Computer Science, vol. 10303*, edited by S. Barbosa, P. Markopoulos, F. Paternò, S. Stumpf, and S. Valtolina, 135–149. Cham: Springer. https://doi.org/10.1007/978-3-319-58735-6_10.

Lankton, Nancy K., D. Harrison McKnight, and John F. Tripp. 2017. "Facebook Privacy Management Strategies: A Cluster Analysis of User Privacy Behaviors." *Computers in Human Behavior* 76:149–163. https://doi.org/10.1016/j.chb.2017.07.015.

Lee, Hwansoo, Siew Fan Wong, Jungjoo Oh, and Younghoon Chang. 2019. "Information Privacy Concerns and Demographic Characteristics: Data from a Korean Media Panel Survey." *Government Information Quarterly* 36 (2): 294–303. https://doi.org/10.1016/j.giq.2019.01.002.

Li, Shu Yan, and Prasad Siba Borah. 2018. "An Empirical Study on the Changes of Internet Privacy Concern: Differences between American and Chinese Cultures." *European Journal of Business and Management* 10 (21): 95–106.

Liu, Bin, Mads Schaarup Andersen, Florian Schaub, Hazim Almuhimedi, Shikun Zhang, Norman Sadeh, Alessandro Acquisti, and Yuvraj Agarwal. 2016. "Follow My Recommendations: A Personalized Privacy Assistant for Mobile App Permissions." In *Proceedings of the Twelfth USENIX Conference on Usable Privacy and Security* (SOUPS '16), 27–41.

Luzak, Joasia. 2021. "Tailor-made Consumer Protection: Personalisations Impact on the Granularity of Consumer Information." In *Legal Design*, edited by Marcelo Corrales Compagnucci, Helena Haapio, Margaret Hagan, and Michael Doherty, 107–132. Edward Elgar Publishing. https://doi.org/10.4337/9781839107269.00013.

Malhotra, Naresh K., Sung S. Kim, and James Agarwal. 2004. "Internet Users' Information Privacy Concerns (IUIPC): The Construct, the Scale, and a Causal Model." *Information Systems Research* 15 (4): 336–355. https://doi.org/10.1287/isre.1040.0032.

Marky, Karola, Alina Stöver, Sarah Prange, Kira Bleck, Paul Gerber, Verena Zimmermann, Florian Müller, Florian

Alt, and Max Mühlhäuser. 2024. "Decide Yourself or Delegate-User Preferences Regarding the Autonomy of Personal Privacy Assistants in Private iot-Equipped Environments." In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, 1–20.

Masotina, Mariavittoria, and Anna Spagnolli. 2022. "Transparency of Privacy Notices and Contextualisation: Effectively Conveying Information without Words." *Behaviour & Information Technology* 41 (10): 2120–2150. https://doi.org/10.1080/0144929X.2022.2077234

Masur, Philipp K. 2020. "How Online Privacy Literacy Supports Self-data Protection and Self-determination in the age of Information." *Media and Communication* 8 (2): 258–269. https://doi.org/10.17645/mac.v8i2.2855.

Mayring, Philipp. 2014. *Qualitative Content Analysis: Theoretical Foundation, Basic Procedures and Software Solution*. Klagenfurt.

McDonald, A. M., and L. F. Cranor. 2008. The Cost of Reading Privacy Policies. *Isjlp*. Retrieved from https://heinonline.org/hol-cgi-bin/get_pdf.cgi?handle=hein.journals/isjlpsoc4&section=27.

Milne, George R., and Mary 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.

Miltgen, Caroline Lancelot, and Dominique Peyrat-Guillard. 2014. "Cultural and Generational Influences on Privacy Concerns: A Qualitative Study in Seven European Countries." *European Journal of Information Systems* 23 (2): 103–125. https://doi.org/10.1057/ejis.2013.17

Mohamed, Norshidah, and Ili Hawa Ahmad. 2012. "Information Privacy Concerns, Antecedents and Privacy Measure use in Social Networking Sites: Evidence from Malaysia." *Computers in Human Behavior* 28 (6): 2366–2375. https://doi.org/10.1016/j.chb.2012.07.008.

Moody, Daniel L. 2004. "Cognitive Load Effects on End User Understanding of Conceptual Models: An Experimental Analysis." In *Advances in Databases and Information Systems. ADBIS 2004. Lecture Notes in Computer Science, vol. 3255*, edited by A. Benczúr, J. Demetrovics, and G. Gottlob, 129–143. Berlin: Springer. https://doi.org/10.1007/978-3-540-30204-9_9.

Morel, Victor, Leonardo Horn Iwaya, and Simone Fischer-Hübner. 2025. "AI-driven Personalized Privacy Assistants: A Systematic Literature Review." *IEEE Access: Practical Innovations, Open Solutions* 13:160982–161002. https://doi.org/10.1109/access.2025.3609188.

Obar, Jonathan A., and Anne Oeldorf-Hirsch. 2020. "The Biggest lie on the Internet: Ignoring the Privacy Policies and Terms of Service Policies of Social Networking Services." *Information, Communication and Society* 23 (1): 128–147. https://doi.org/10.1080/1369118X.2018.1486870.

Oehler, Andreas, and Stefan Wendt. 2017. "Good Consumer Information: The Information Paradigm at Its (Dead) end?" *Journal of Consumer Policy* 40 (2): 179–191. https://doi.org/10.1007/s10603-016-9337-5.

O'Neil, Dara. 2001. "Analysis of Internet Users' Level of Online Privacy Concerns." *Social Science Computer Review* 19 (1): 17–31. https://doi.org/10.1177/089443930101900103

Ortloff, Anna-Marie, Maximiliane Windl, Valentin Schwind, and Niels Henze. 2020. "Implementation and In Situ Assessment of Contextual Privacy Policies." In *Proceedings of the 2020 ACM Designing Interactive Systems Conference* (DIS '20), 1765–1778. https://doi.org/10.1145/3357236.3395549

Osatuyi, Babajide. 2015. "Personality Traits and Information Privacy Concern on Social Media Platforms." *Journal of Computer Information Systems* 55 (4): 11–19. https://doi.org/10.1080/08874417.2015.11645782.

Park, Yong Jin. 2013. "Digital Literacy and Privacy Behavior Online." *Communication Research* 40 (2): 215–236. https://doi.org/10.1177/0093650211418338.

Peer, Eyal, Serge Egelman, Marian Harbach, Nathan Malkin, Arunesh Mathur, and Alisa Frik. 2020. "Nudge Me Right: Personalizing Online Security Nudges to People's Decision-Making Styles." *Computers in Human Behavior* 109:106347. https://doi.org/10.1016/j.chb.2020.106347.

Peng, Chao-Ying Joanne, Kuk Lida Lee, and Gary M. Ingersoll. 2002. "An Introduction to Logistic Regression Analysis and Reporting." *The Journal of Educational Research* 96 (1): 3–14. https://doi.org/10.1080/00220670209598786.

Personalized Privacy Assistant Project, retrieved October 1, 2024 from https://privacyassistant.org/.

Preibusch, Sören. 2013. "Guide to Measuring Privacy Concern: Review of Survey and Observational Instruments." *Intl. Journal of Human--Computer Studies* 71 (12): 1133–1143. https://doi.org/10.1016/j.ijhcs.2013.09.002.

Prince, Christine, Nessrine Omrani, Adnane Maalaoui, Marina Dabic, and Sascha Kraus. 2021. "Are We Living in Surveillance Societies and Is Privacy an Illusion? An Empirical Study on Privacy Literacy and Privacy Concerns." *IEEE Transactions on Engineering Management* 7: 3553–3570. https://doi.org/10.1109/TEM.2021.3092702.

Proctor, Robert W., M. Athar Ali, and Kim-Phuong L. Vu. 2008. "Examining Usability of Web Privacy Policies." *International Journal of Human–Computer Interaction* 24 (3): 307–328. https://doi.org/10.1080/10447310801937999.

Reidenberg, Joel R., Travis Breaux, Lorrie Faith Cranor, Brian French, Amanda Grannis, James T. Graves, Fei Liu, Aleecia McDonald, Thomas B. Norton, and Rohan Ramanath. 2015. "Disagreeable Privacy Policies: Mismatches between Meaning and Users' Understanding." *Berkeley Technology Law Journal* 30 (1): 1–88. Retrieved from https://heinonline.org/HOL/P?h=hein.journals/berktech30&I=51.

Salonen, Ville, and Heikki Karjaluoto. 2016. "Web Personalization: The State of the Art and Future Avenues for Research and Practice." *Telematics and Informatics* 33 (4): 1088–1104. https://doi.org/10.1016/j.tele.2016.03.004.

Salvucci, Dario D., and Joseph H. Goldberg. 2000. "Identifying Fixations and Saccades in Eye-Tracking Protocols." In *Proceedings of the 2000 Symposium on Eye Tracking Research & Applications* (ETRA '00), 71–78. https://doi.org/10.1145/355017.355028

Sheehan, K. B. 2002. Toward a Typology of Internet Users and Online Privacy Concerns. *The information society*. Retrieved from https://www.tandfonline.com/doi/abs/10.1080019722402528182O7.

Standing Committee of the National People's Congress. 2021. Personal Information Protection Law of the People's Republic of China (PIPL). Retrieved from http://en.npc.gov.cn.cdurl.cn/2021-12/29/c_694559.htm.

Steijn, Wouter M. P., and Anton Vedder. 2015. "Privacy Concerns, Dead or Misunderstood? The Perceptions of Privacy amongst the Young and old." *Information Polity* 20 (4): 299–311. https://doi.org/10.3233/ip-150374.

Steinfeld, Nili. 2016. ""I Agree to the Terms and Conditions": (How) Do Users Read Privacy Policies Online? An eye-Tracking Experiment." *Computers in Human Behavior* 55:992–1000. https://doi.org/10.1016/j.chb.2015.09.038.

Strecker, Jannis, Simon Mayer, and Kenan Bektaş. 2025. "Towards Societally Beneficial Personalized Realities: A Conceptual Foundation for Responsible Ubiquitous Personalization Systems." In *Designing Interactive Systems Conference (DIS '25), July 05–09, 2025, Funchal, Portugal*, Vol. 23. New York, NY: ACM. https://doi.org/10.1145/3715336.3735709.

Tamò-Larrieux, A., Z. Zihlmann, K. Garcia, and S. Mayer. 2021. "The Right to Customization: Conceptualizing the Right to Repair for Informational Privacy." In *Privacy Technologies and Policy. APF 2021. Lecture Notes in Computer Science, vol. 12703*, edited by N. Gruschka, L. F. C. Antunes, K. Rannenberg, and P. Drogkaris. Cham: Springer. https://doi.org/10.1007/978-3-030-76663-4_1.

Tifferet, Sigal. 2019. "Gender Differences in Privacy Tendencies on Social Network Sites: A Meta-analysis." *Computers in Human Behavior* 93:1–12. https://doi.org/10.1016/j.chb.2018.11.046.

Trepte, Sabine, Doris Teutsch, Philipp K. Masur, Carolin Eicher, Mona Fischer, Alisa Hennhöfer, and Fabienne Lind. 2015. "Do People Know about Privacy and Data Protection Strategies? Towards the "Online Privacy Literacy Scale" (OPLIS)." In *Reforming European Data Protection Law*, edited by Serge Gutwirth, Ronald Leenes, and Paul de Hert, 333–365. Dordrecht: Springer Netherlands. https://doi.org/10.1007/978-94-017-9385-8_14

U. N. Trade and Development. 2023. Age Structure. Retrieved from https://hbs.unctad.org/age-structure/.

Vu, Kim-Phuong L., Vanessa Chambers, Fredrick P. Garcia, Beth Creekmur, John Sulaitis, Deborah Nelson, Russell Pierce, and Robert W. Proctor. 2007. "How Users Read and Comprehend Privacy Policies." In *Human Interface and the Management of Information. Interacting in Information Environments: Symposium on Human Interface 2007, Held as Part of HCI International 2007, Beijing, China, July 22-27, 2007, Proceedings, Part II*, 802–811.

Wagner, Isabel. 2023. "Privacy Policies across the Ages: Content of Privacy Policies 1996–2021." *ACM Transaction Privacy Security* 26 (3): 1–32. https://doi.org/10.1145/3590152

Williams, Rihana Shiri, Omer Ari, and Carmen Nicole Santamaria. 2011. "Measuring College Students' Reading Comprehension Ability Using Cloze Tests." *Journal of Research in Reading* 34 (2): 215–231. https://doi.org/10.1111/j.1467-9817.2009.01422.x

Windl, Maximiliane, Niels Henze, Albrecht Schmidt, and Sebastian S. Feger. 2022. "Automating Contextual Privacy Policies: Design and Evaluation of a Production Tool for Digital Consumer Privacy Awareness." In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (CHI '22), 1–18. https://doi.org/10.1145/3491102.3517688

Woodring, Justin, Katherine Perez, and Aisha Ali-Gombe. 2024. "Enhancing Privacy Policy Comprehension through EmphPrivacify: A User-Centric Approach Using Advanced Language Models." *Computers & Security* 145:103997. https://doi.org/10.1016/j.cose.2024.103997.

Xu, Heng, Sumeet Gupta, M. Rosson, and John Millar Carroll. 2012. "Measuring Mobile Users' Concerns for Information Privacy." In *ICIS 2012 Proceedings*. https://aisel.aisnet.org/icis2012/proceedings/ISSecurity/10.

Xu, Meihe, Clement Guitton, Aurelia Tamò-Larrieux, and Christoph Lutz. 2026. "Personalized Transparency: A Scoping Review of Approaches to Personalized Privacy Disclosures." *Computer Law & Security Review: The International Journal of Technology Law and Practice* 60:106265. https://doi.org/10.1016/j.clsr.2026.106265.

Xu, Meihe, Žiga Jug, and Aurelia Tamò-Larrieux. 2024. "A Cross-Cultural Analysis of Transparency: The Interplay of law, Privacy Policies, and User Perceptions." *International Data Privacy Law* 14 (3): 197–222. https://doi.org/10.1093/idpl/ipae011

Xu, Meihe, Arianna Rossi, and Aurelia Tamò-Larrieux. 2025. "The Future of Personalized Privacy Assistants: Gathering of Expert Opinions." *Digital Society* 4: 1–32. https://doi.org/10.1007/s44206-025-00232-4.

Yang, Zhihui, Ruiming Wang, Hui Chen, and Jiali Ding. 2015. "Personality and Worry: The Role of Intolerance of Uncertainty." *Social Behavior and Personality: An International Journal* 43 (10): 1607–1616. https://doi.org/10.2224/sbp.2015.43.10.1607.

Youn, Seounmi. 2009. "Determinants of Online Privacy Concern and Its Influence on Privacy Protection Behaviors among Young Adolescents." *The Journal of Consumer Affairs* 43 (3): 389–418. https://doi.org/10.1111/j.1745-6606.2009.01146.x.

Zhang, Donghuan, Min Fan, Lingyi Meng, and Xifu Zheng. 2022. "Neuroticism and Fear of COVID-19 during the COVID-19 Pandemic: Testing the Mediating Role of Intolerance of Uncertainty and Sense of Control among Chinese High School Students." *Frontiers in Psychology* 13: 1–10. https://doi.org/10.3389/fpsyg.2022.1010767.

Zhang, Shikun, Yuanyuan Feng, Yaxing Yao, Lorrie Faith Cranor, and Norman Sadeh. 2022. "How Usable Are iOS App Privacy Labels?." *Proceedings on Privacy Enhancing Technologies* 2022 (4): 204–228. https://doi.org/10.56553/popets-2022-0106.

Zhang, Shikun, Lily Klucinec, Kyerra Norton, Norman Sadeh, and Lorrie Faith Cranor. 2024. "Exploring Expandable-Grid Designs to Make iOS app Privacy Labels More Usable." *In Proceedings of the Twentieth USENIX Conference on Usable Privacy and Security* SOUPS 2024: 139–157.

Zhang, Xing, Shan Liu, Xing Chen, Lin Wang, Baojun Gao, and Qing Zhu. 2018. "Health Information Privacy Concerns, Antecedents, and Information Disclosure Intention in Online Health Communities." *Information & Management* 55 (4): 482–493. https://doi.org/10.1016/j.im.2017.11.003.

Zukowski, Tomasz, and Irwin Thomas Joseph Brown. 2007. "Examining the Influence of Demographic Factors on Internet Users' Information Privacy Concerns." In Proceedings of the 2007 Annual Conference of the South African Institute of Computer Scientists and Information Technologists on IT Research in Developing Countries, SAICSIT Conf. 2007, Port Elizabeth, South Africa, October 2-3, 2007, 197–204. https://doi.org/10.1145/1292491.1292514

Appendices

Appendix A. Study 1: prediction of most concerned category of information in a privacy policy

Table A1. Privacy Statements in different forms as used in Study 1

Condensed	Headings in a Privacy Policy	Summarised Privacy Statements
Information collected	What information we collect	Information about the data (e.g. profile information, user content, contacts, location information, cookies, payment information) that is being collected from the social media provider.
Information usage	How we use your information	Information about the purposes for which user data is used (e.g. personalising the experience, providing advertisement and sharing of data with third parties to provide advertisement, facilitating purchases, enforcing the terms of services).
Information sharing	How we share your information	Information about how third parties obtain access to the data of the user. Third parties can be cloud hosting or content delivery providers, partners or corporate groups of the social media providers (e.g. for single-sign-in by using a third party platform), or the public depending on the privacy settings chosen by the user.
Legal bases and information processing	Our legal bases and how we process your information	Information about the legal grounds upon which user data is being processed (e.g. certain data are collected based on a contractual obligation, other data is collected only with the user's consent).
Rights and choices	Your rights and choices	Information about the rights that users have with respect to how their data is being processed (e.g. information on how to make use of access rights, information on how to demand for certain data to be deleted).
Security and retention	Data security and retention	Information about what security measures are put in place to protect user data and how long the data is kept by the social media provider.
Global operations and transfers	Our global operations and data transfers	Information about how data is being used globally and transferred to entities processing data on behalf of the social media provider (e.g. information about the legal safeguards in place for such transfers).
Young users	Young users	Information targeted to users under 13 (e.g. information on how to obtain more age-appropriate content, or information for parental consent).
Updates	Privacy policy updates	Information on how the privacy policy can be updated and the notifications that this triggers.
Contact	Contact us	Information on how to contact the social media provider in case of questions related to privacy.

Appendix B. Study 1: additional results

Table A2. Classification table for classification accuracy on predicted vs. actual outcomes

	Predicted										
Observed	Information collected	Information usage	Information sharing	Legal bases and information processing	Rights and choices	Security and retention	Our global operations and data transfers	Young users	Updates	Contact	Percent correct
What information we collect	79	4	2	1	2	0	1	1	0	0	87.8%
How we use your information	16	6	1	0	2	0	2	0	0	0	22.2%
How we share your information	11	1	3	0	0	0	0	0	1	0	18.8%
Our legal bases and how we process your information	5	0	0	0	0	1	0	1	0	0	0.0%
Your rights and choices	9	0	1	0	7	1	1	0	0	0	36.8%
Data security and retention	6	1	0	0	0	2	1	0	0	0	20%
Our global operations and data transfers	10	1	2	0	1	0	2	0	0	0	12.5%
Young users	8	0	0	0	1	1	0	0	0	0	0.0%
Privacy policy updates	3	0	0	0	0	0	0	0	3	0	50%
Contact us	4	0	0	0	1	0	0	0	0	1	16.7%
Overall percentage	72.9%	6.3%	4.3%	0.5%	6.8%	2.4%	3.4%	1.0%	1.9	0.5%	49.8%

Appendix C. Study 1 and study 2: socio-demographic information of the participants

Table A3. Socio-demographic information of participants in Study 1

Category	Description	N	Marginal Percentage
Most concerned statement with summarised statements	What information we collect	90	43.5%
	How we use your information	27	13.0%
	How we share your information	16	7.7%
	Our legal bases and how we process your information	7	3.4%
	Your rights and choices	19	9.2%
	Data security and retention	10	4.8%
	Our global operations and data transfers	16	7.7%
	Young users	10	4.8%

(*Continued*)

Table A3. Continued.

Category	Description	N	Marginal Percentage
	Privacy policy updates	6	2.9%
	Contact us	6	2.9%
Level of Education	Below high school	8	3.9%
	High school and post-high school	60	29%
	Short-cycle and bachelor's	84	40.6%
	Master's and doctoral	55	26.6%
Age	18–24	57	27.5%
	25–39	123	59.4%
	60–64	27	13.0%
Gender	Men	105	50.7%
	Women	102	49.3%
Total		207	

Table A4. Socio-demographic information of participants in Study 2

Category	Description	N	Marginal Percentage
Most concerned statement with summarised statements	What information we collect	12	40%
	How we use your information	4	13.3%
	How we share your information	4	13.3%
	Our legal bases and how we process your information	1	3.3%
	Your rights and choices	1	3.3%
	Data security and retention	3	10%
	Our global operations and data transfers	2	6.7%
	Young users	0	-
	Privacy policy updates	0	-
	Contact us	3	10%
Level of Education	High school and post-high school	6	20%
	Short-cycle and bachelor's	15	50%
	Master's and doctoral	9	30%
Major	Economics	6	20%
	Accounting and Finance	4	13.3%
	Strategy and International Management	3	10%
	Business Administration	3	10%
	International Affairs	2	6.7%
	Quantitative Economics and Finance	2	6.7%
	Computer Science	2	6.7%
	Management	2	6.7%
	Other	6	20%
Age	18–24	15	50%
	25–34	15	50%
Gender	Men	20	66.7%
	Women	9	30%
	Prefer not to say	1	3.3%
Total		30	

Appendix D. Study 2: semi-structured interview guide

In the semi-structured interview, we asked all participants the main questions (1) to (6). Depending on their answers, we adapted the follow-up questions.

1. When you were answering the two fill-in-the-blank tests, how did you decide which words to write?
2. When you installed the two apps, did you read their privacy policies?
    a. If Yes, then: Did you notice any differences between them? Can you describe what those differences were? Follow-up: Do you prefer one privacy policy over the other? If so, what are your reasons for this preference? And if not, what are your reasons?
    b. If No, then: Can you share why you chose not to? Follow up: Even though you didn't read the privacy policies, did anything about them stand out to you in any way? For example, did you notice any differences between the two privacy policies?
3. We personalised the order of information in one privacy policy based on your preferences from the pre-experiment survey. Did you notice this personalisation? If Yes, then: Did it affect your understanding or perception of the privacy policy? How so?
4. Do you think this personalised approach is helpful? Why or why not?
5. What could make privacy policies more engaging or easier to read for you? Do you have any specific suggestions?
6. Do you think other ways of information disclosure are better alternatives than privacy policies for you?

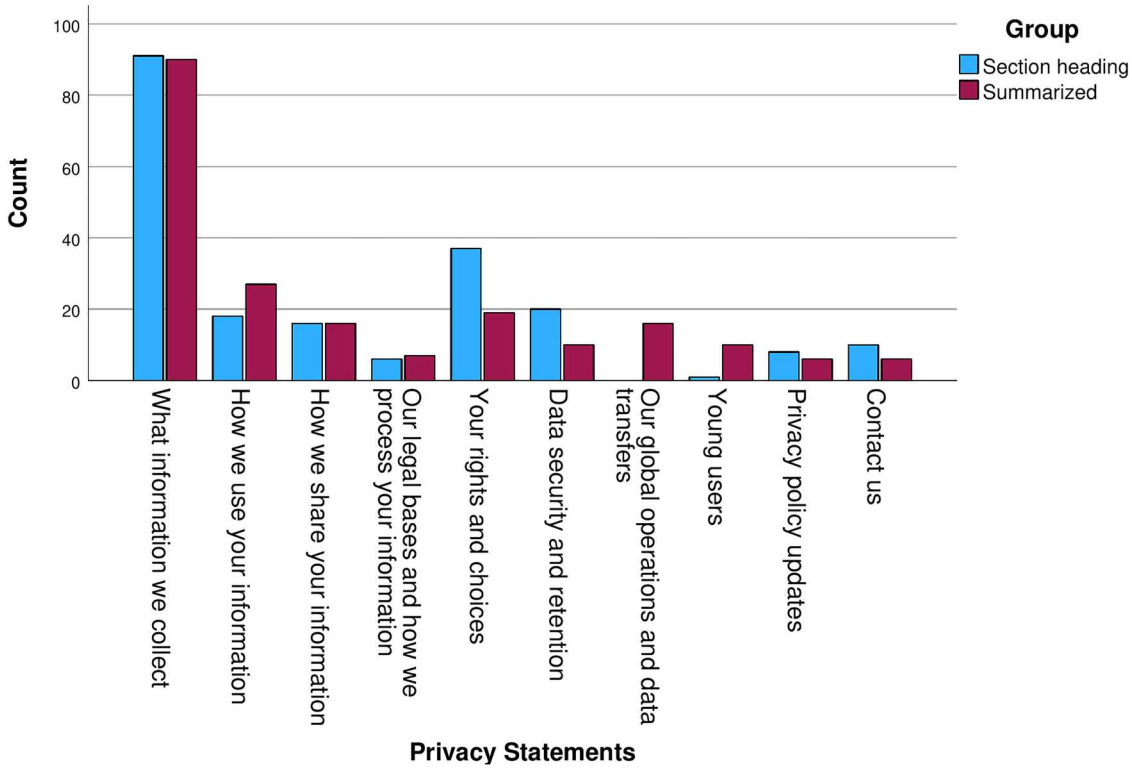

Figure A1. Comparison of participants' most concerned category of information: headings vs. summarised privacy statements (N = 207).