

Analysing the Influence of Loss-Gain Framing on Data Disclosure Behaviour:

A Study on the Use Case of App Permission Requests

Kerstin Bongard-Blanchy

Luxembourg Media and Digital
Design Center
Belvaux, Luxembourg
kerstin.bongard-blanchy@uni.lu

Jean-Louis Sterckx

KU Leuven
Leuven, Belgium
jean-louis.sterckx@live.be

Arianna Rossi

LIDER Lab, Dirpolis Institute, Scuola
Superiore Sant'Anna
Pisa, Italy
arianna.rossi@santannapisa.it

Anastasia Sergeeva

University of Luxembourg
Esch sur Alzette, Luxembourg
anastasia.sergeeva@uni.lu

Vincent Koenig

University of Luxembourg
Esch sur Alzette, Luxembourg
vincent.koenig@uni.lu

Salvador Rivas

University of Luxembourg
Esch sur Alzette, Luxembourg
salvador.rivas@uni.lu

Verena Distler

University of Luxembourg
Luxembourg
University of the Bundeswehr Munich
Munich, Germany
verena.distler@unibw.de

ABSTRACT

This paper examines the effect of the dark pattern strategy “loss-gain framing” on users’ data disclosure behaviour in mobile settings. Understanding whether framing influences users’ willingness to disclose personal information is important to (i) determine if and how this technique can subvert consent and other privacy decisions, (ii) prevent abuse with appropriate policies and sanctions, and (iii) provide clear evidence-based guidelines for app privacy engineering. We conducted an online user study (N=848), in which we varied the framing of app permission requests (i.e., positive, negative, or neutral framing) and examined its impact on participants’ willingness to accept the permission, their evaluation of the trustworthiness of the request and their perception of being informed by it. Our findings reveal effects on disclosure behaviour for request types that users cannot easily understand. In this case, negative framing makes users more likely to disclose personal information. Contrary to our expectations, positive framing reduces disclosure rates, possibly because it raises users’ suspicion. We discuss implications for the design of interfaces that aim to facilitate informed, privacy-enhancing decision-making.

CCS CONCEPTS

• **Security and privacy** → **Usability in security and privacy**; • **Human-centered computing** → **Empirical studies in HCI**.

KEYWORDS

Usable privacy and security, Human-computer interaction, Dark patterns, Empirical research

ACM Reference Format:

Kerstin Bongard-Blanchy, Jean-Louis Sterckx, Arianna Rossi, Anastasia Sergeeva, Vincent Koenig, Salvador Rivas, and Verena Distler. 2023. Analysing the Influence of Loss-Gain Framing on Data Disclosure Behaviour: A Study on the Use Case of App Permission Requests. In *The 2023 European Symposium on Usable Security (EuroUSEC 2023)*, October 16–17, 2023, Copenhagen, Denmark. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3617072.3617108>

1 INTRODUCTION

Our digital experiences have become a playground for designs that exploit cognitive biases to steer people into decisions that are not necessarily in their best interest through the use of so-called dark patterns or deceptive design patterns. The term *dark patterns* refers to interface design techniques that manipulate the information flow or modify the decision space [44]. These designs are omnipresent in our digital experiences, as a recent report authored by the European Commission exposes [21]. Dark patterns can have far-reaching consequences on user privacy since they may cause individuals to unwittingly and invisibly share personal data with thousands of third parties. Indiscriminate data sharing can expose them to various types of harm, including behavioural manipulation, psychological or even physical detriment, discrimination, time loss, and financial risks [16, 50]. On a collective level, the use of



This work is licensed under a [Creative Commons Attribution International 4.0 License](https://creativecommons.org/licenses/by/4.0/).

EuroUSEC 2023, October 16–17, 2023, Copenhagen, Denmark
© 2023 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0814-5/23/10.
<https://doi.org/10.1145/3617072.3617108>

dark patterns may hamper competition and impact user trust in the digital economy [44, 50].

Dark patterns are increasingly under the spotlight and cause frustration and great concern from individuals and institutions alike. On both sides of the Atlantic, policymakers reinforce existing prohibitions on the use of manipulative techniques on digital services (e.g., in the Digital Services Act, Digital Markets Act, Data Act proposal and the AI Act proposal, as well as the California Privacy Rights Act, Colorado Privacy Act and Connecticut Data Privacy Act, among the others). Watchdogs have started to fine the use of privacy-invasive dark patterns that circumvent consent, transparency and other data protection obligations [27]. Moreover, official guidelines that interpret applicable laws in terms of technology design are being issued. For example, in 2023, the European Data Protection Board (EDPB) published a taxonomy of dark patterns that are likely to violate the General Data Protection Regulation [20]. Among the various deceptive strategies, the EDPB condemns the implementation of “emotional steering”, defined as “using wording or visual elements [...] in a way that conveys information to users in either a highly positive outlook, making users feel good, safe or rewarded, or a highly negative one, making users feel anxious, guilty or punished.” (p. 19). The EDPB warns that emotional steering influences “the emotional state in a way that is likely to lead users to act against their data protection interests” (p. 20). Such a strategy is commonly called loss-gain framing [13] and occurs when people’s decisions are influenced by how options are linguistically presented (“framed”) rather than by the inherent characteristics of the options [63].

A concrete case where framing can affect user privacy decisions is permission dialogues in mobile applications, which serve as a control mechanism to help users oversee and safeguard access to data and resources on their smartphones. Users encounter such requests when installing or using a smartphone application. While app permissions differ in design and formulations across operating systems, their common denominator is that they offer basic information about why they are requested and to what extent applications have access to personal data after being authorised [9]. In this model, it is exclusively in the hands of the users to control the access by app developers and numerous, often opaque, third parties to their personal data (e.g., intimate pictures, financial information, private conversations). Thus, mobile device users are called to assess whether the requested access to a specific resource is appropriate and proportional with respect to the app functionality and then take a decision with broad impacting consequences.

This paper presents an empirical study that investigates the effect of framing on users’ decision to accept or decline an app permission request. We conducted an experiment with 848 participants to test the effect of three types of framing (negative - neutral - positive) for three permission types: storage, camera, and location, in the context of a fictive navigation app. Our research makes the following contributions:

- We found empirical evidence that negative framing can steer users into accepting app permissions, while positive framing may decrease acceptance when it is unclear to the user why such permissions are asked.

- We found no statistically significant effect of framing on the extent to which people perceive to be informed by an app permission request, and partial evidence that positive framing might decrease people’s perceived trustworthiness of the permission request.
- We found that higher perceived trustworthiness of the app is associated with higher acceptance rates.
- We discuss the implications of our results for user-centred privacy permission design.

2 RELATED WORK

2.1 Privacy dark patterns and loss-gain framing

Within the broad typology of dark patterns, certain dark patterns seek to steer users’ privacy decisions or modify their decision space and can lead to material and non-material harms [27]. Such privacy-related dark patterns are usually found either in “entry” requests (e.g., subscription) and consent interactions, in user settings, or in “exit” requests (e.g., opt-outs, profile erasure) [27]. Since the GDPR has strengthened transparency and consent requirements, researchers have denounced that positive and negative language in cookie banners attempt to steer website visitors towards more data disclosure for advertising and profiling purposes [31, 56]. Experimental studies have demonstrated that many design elements in consent interfaces [28] and cookie banners can influence user choices towards more data sharing [6, 7, 24, 42, 49].

Dark patterns are commonly associated with visual design strategies, but the use of language can also be an element of manipulation. For instance, the framing-based “confirmshaming” uses emotional wording to inspire a sense of guilt in the users and thereby push them to perform an action [14], for example, subscribe to a service, or to refrain from cancelling that subscription. In addition, complicated, overly technical or legalistic language in privacy policies and consent requests can overload users with information, confuse or deceive them, and thereby limit their abilities to make an informed decision [12, 15, 56]. To avoid deceptive and misleading linguistic expressions, the GDPR explicitly mandates the use of plain, non-misleading language in any privacy-related information and request directed to users (i.e., the requirements of transparency, consent and data protection by design [10]). This is why the EDPB [20] warns that misleading information, i.e., a difference between the information that is provided and the available actions, may nudge users to act in unintended manners and therefore weaken their rights and freedoms by infringing one of the overarching principles of EU data protection law, namely the fairness of data processing. They also argue that providing conflicting information and ambiguous wording leaves individuals unsure of what they should do, preventing them from exercising a real choice.

Another way of employing language to manipulate user decisions towards a predictable outcome exploits the cognitive bias named “framing effect”. Loss-gain framing refers to a presentation of options that overemphasises either their positive consequences (i.e., the gains) or the negative ones (i.e., the losses) [13]. The framing effect was initially uncovered in the context of risky choices about monetary outcomes or the loss of human lives [63], and it can steer human decision-making in various contexts, including medical [23], political [34] and organisational [58] contexts.

Framing is ubiquitous in websites and applications: experimental data show that framing of deals is widely used on e-commerce sites [43] and that the wording of such messages can steer users into accepting the deal [40]. This strategy is particularly concerning regarding decisions about one's data. In a study examining 407 cookie banners, Santos et al. [56] found that 30% of cookie banners used positively framed messages which direct the user's attention towards the benefits of accepting cookies, such as a better user experience, while the risks of extensive online tracking are withheld. Similarly, 2% of banners applied negatively framed messages that emphasise the negative consequences of rejecting cookies, such as loss of functionality. The authors hypothesise that such framing of consequences may influence user behaviour towards accepting online tracking.

An effect of framing on users' disclosure behaviour has already been demonstrated. Adjerid et al. [3] found effects ranging from 10% to 14% of increased disclosure behaviour for sensitive information when a high data protection notice preceded the question compared to a low protection privacy notice. Other studies have shown that the labelling of data disclosure choices (namely, "privacy" versus "app settings"; and "allow" versus "prohibit") has a significant impact on the extent to which people decide to share their personal information online in contexts with a high privacy risk [4, 54]. However, these studies did not investigate the distinct effect of loss-gain framing on users' disclosure behaviour. Other studies did not find a considerable effect of positive-negative framing neither regarding users' privacy choices in cookie consent banners [7] nor altering their level of privacy concern following the exposure to privacy notices [17, 22]. However, in Berens et al. [7] and de Gluck et al. [22], the manipulated phrases were buried in a long text, at the risk of being overlooked by the study participants. Indeed, Berens et al. [7] noted that two-thirds of the study participants self-reported not reading the text. Such nuances indicate a challenge in studying and demonstrating the framing effect in long textual documents, where other elements may be at play.

2.2 Trust judgements and uncertainty

Privacy decisions are not exclusively based on objective information about a company's data practices but also on other elements. For example, trust in the service provider is key to users' willingness to use and adopt a technology [25]. Beyond the notoriety and public image of the provider, it has been shown that the design of interface elements can affect whether interfaces are perceived as trustworthy [52, 53, 61]. Trust is furthermore affected by the quality and quantity of the proposed content and information [35] and, even more specifically, explanations of recommendations provided by the system [29, 33, 48]. Such elements act as proxies for users' trust judgements and may not depend on the actual trustworthiness of a company, product or application.

Uncertainty seems to be an additional element that influences users' perception of the system's privacy and thus their decisions [2]. Through the constant evolution of information technologies and the complex and nuanced trade-offs associated with privacy decisions, users are left with incomplete information about possible privacy-relevant outcomes and their consequences [1]. As individuals face these layers of complexity, uncertainty arises, which elicits the use

of heuristic thinking and potentially leads to biases in decision-making. A recent online study examined the uncertainty that users associate with the app permission requests of popular mobile applications [11]. The authors investigated people's certainty about app permission requests for eight apps with their associated permissions (48 combinations in total). In particular, the study considered two aspects of the uncertainty that users can face: first, users may not understand why the app needs access to a particular resource and therefore asks for permission (understandability); second, even if they understand why the application requests such permission, they may not know if granting access is relevant to their personal use of the app or not (clarity on relevance). The results showed that users were uncertain about the necessity of accepting app permissions for 56% of the requests. The study also identified app permission requests that users were highly certain about (considering both understandability and clarity on relevance), such as the location access permission for a navigation app like Google Maps, and others that users were much less certain about, such as camera for apps like Spotify or calendar for Tiktok.

3 RESEARCH OBJECTIVE

Building on the association of established permission requests with categories of (un-)certainty in [11], our objective was to understand whether uncertainty can be a relevant factor in privacy decision-making [2]. In particular, since the literature presents conflicting results about the effects of framing, we wanted to examine whether loss-gain framing can push users to authorise apps to access various resources, similarly to the study described in [17]. In addition, we also sought to test whether framing influenced the perceived trustworthiness of the app and users' perception of being adequately informed about why the app makes such a request. We additionally factored the (un)certainty users may feel about why an app asks for permission in. Thus, we address the following research questions and hypotheses:

RQ1: How does framing of app permission requests influence users' disclosure behaviour?

H1.1: The *negative* framing of app permission requests affects people's disclosure behaviour towards *accepting* in the case of uncertainty.

H1.2: The *positive* framing of app permission requests affects people's disclosure behaviour towards *accepting* in the case of uncertainty.

RQ2 How does framing of app permission requests influence people's perception of being informed and their trust in the request?

H2.1a: The *positive* framing of app permission requests influences people's *perception of being informed*.

H2.1b: The *negative* framing of app permission requests influences people's *perception of being informed*.

H2.2a: The *positive* framing of app permission requests influences the *perceived trustworthiness* of the request.

H2.2b: The *negative* framing of app permission requests influences the *perceived trustworthiness* of the request.

H2.3a: People's perception of being informed is associated with their *disclosure behaviour*.

H2.3b: The perceived trustworthiness of the request is associated with their *disclosure behaviour*.

4 METHOD

4.1 Sampling and demographics

We conducted an online survey in November 2022 hosted on LimeSurvey¹ and distributed through Prolific². We collected answers from adult participants using the sampling options for current location in the UK, English fluency and gender balance. The survey took a maximum of five minutes. Participants received a compensation of £8.66/hr. The study received approval from the University of Luxembourg’s ethics board. Through a between-subjects design, we collected responses from 851 participants. 3 participants’ answers were excluded as they were incomplete, resulting in a sample of 848 participants. In this sample 48% identified as male, 51% female, and 1% as non-binary. Their age ranged from 18 to 86 years (mean 41.2, SD 13.6). 30% had a high school diploma or lower, 49% vocational training or a Bachelor’s degree, and 21% a Master’s degree or higher. 47% participants were iOS users, 52% Android users and 1% declared using another operating system. 32% of participants used navigation apps at least once per month.

4.2 Survey material and structure



Figure 1: Mock-up of app store and home screen.

We selected suitable permission requests for this study based on the results of Bongard-Blanchy et al.’s [11] study. There, the navigation app and its permissions received strong (un)certainly scores across the different permission types. Thus, as a use case, we selected a fictive navigation app (“Atlasly”) and three app permission requests with differing certainty levels based on Bongard-Blanchy et al.’s [11] study (see Section 2.2), namely: camera (low-medium

certainty), storage (medium certainty), and location (high certainty). As a consequence, we created nine permission request mock-ups: three types (location, storage, camera) with three framing conditions (positive, neutral, negative).

First, participants answered demographic questions related to age, gender, education, mobile phone operating system and frequency of navigation app use. They were then presented with a fictitious yet realistic scenario that enabled them to consider their disclosure decision. As part of the scenario, participants were asked to imagine looking for a navigation app on their phone, then install and launch it (Figure 1). An attention check question served to ensure participants understood the instructions. In the following step, they were randomly shown one of the three app permissions (camera, storage, or location) with either positive, neutral or negative framing:

The negative framing read: “Atlasly would like to access the Location/Storage/Camera. **Declining this permission will deteriorate your user experience with the app.**”

The neutral framing read: “Atlasly would like to access the Location/Storage/Camera.”

The positive framing read: “Atlasly would like to access the Location/Storage/Camera. **Accepting this permission will enhance your user experience with the app.**”

Of the nine possible permission-framing combinations, each participant only saw one (see Figure 2).

The participants were asked to accept or decline this app permission request. We will refer to this decision as “Disclosure behaviour”. To finalise the survey, they were asked to rate two statements on a 5-point Likert scale ranging from strongly disagree to strongly agree:

I felt sufficiently informed to decide whether to accept or decline this app permission request. (People’s perception of being informed)

This permission request seemed trustworthy to me. (Perceived trustworthiness)

Both are ad-hoc measures since, to the best of our knowledge, no suitable standardised tool was available. We chose the above item for trustworthiness instead of existing measures for trust in technology (e.g., [26]) as those were only partly suitable for our context. For instance, the trust dimensions of benevolence and competence seemed less suitable for our context (e.g., no assumed perception of benevolence of the request).

In addition to the quantitative measurements, we complemented the study procedure with two answer fields to gather qualitative insights into people’s reasoning for their disclosure behaviour. The participants had to complete the following sentence: *I decided to accept or decline this app permission request because...* To conclude, they were invited to freely share their thoughts and experiences regarding such requests from apps (*What are your thoughts and experiences regarding such requests from apps?*). For the purpose of replicability and transparency, the questionnaire and the full dataset can be found here: https://osf.io/5wc3v/?view_only=469d0989fc3a4422bcfe2b2f5f4a651.

¹<https://www.limesurvey.org/>

²<https://www.prolific.co/>

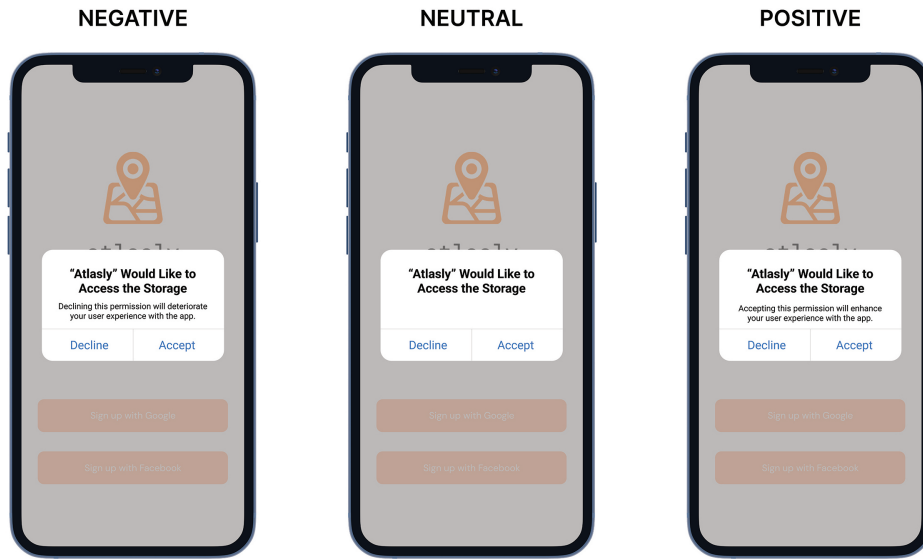


Figure 2: Mock-up of three of the nine tested app permission requests - here for the storage permission with the three different framings: negative, neutral, positive.

4.3 Data analysis

To answer RQ1 (i.e., effect of framing on users' disclosure behaviour), we extracted the percentage of participants who accepted or declined the permission for each of the nine conditions. We then ran a chi-square test with correction for multiple tests between the three framings (negative, neutral, positive) for the same permission type (location, storage, and camera). To answer RQ2 (i.e., effect of framing on people's perception of being informed and the perceived trustworthiness of the request), we computed the mean and standard deviation for the dependent variables: perception of being informed and perceived trustworthiness. We ran a Kruskal-Wallis equality-of-populations rank test with correction for multiple tests to control for the significance of differences between the three framing conditions. Following the approach proposed by Streiner [59], in case of multiple comparisons in post hoc tests, we report our results before and after Bonferroni correction. Results are reported as "weak" if they were significant before the applied correction but not after. A logistic regression was conducted to test the extent to which the gathered independent variables (demographics, test conditions, feeling of being informed and trust) were associated with the participant's probability of disclosing behaviour.

The participants' open-text responses to both questions were compiled and analysed via inductive coding, following the content analysis approach [45]. Given the straightforward questions and answers, we followed the practice recommended by McDonald [46] and had a single author thematically analyse the data. This resulted in responses categorised into the following topics: 1) user reaction and disclosure behaviour, 2) user assessment of the app permission requests, 3) users' attitude towards app permission requests, 4) users' privacy and security concerns, 5) perceived trustworthiness of app permission requests, and 6) the influence of framing. For

each topical code, the occurrence per person was quantified as follows: somebody mentioning a code several times counts as once, meaning that 10% occurrence corresponds to about 85 of the 848 participants mentioning the code at least once.

5 RESULTS

The 848 participants were randomly assigned to one of the nine conditions (between-subjects design). Table 1 shows the number of participants per condition. The following columns indicate the absolute number and percentage of the participants' acceptance rate and the means and SDs for perception of being informed and perceived trustworthiness.

5.1 Effect of framing and uncertainty on disclosure behaviour

In response to RQ1, we sought to determine if positive or negative framing affected the disclosure behaviour towards acceptance. To this end, we grouped the data by framing condition (negative, neutral or positive). We performed a chi-square analysis to see if the disclosure behaviour in any of the three framing condition groups differs from the others and obtained significant results $\chi^2(2, 848) = 8.294, p = .016$ (Figure 3).

The adjusted residuals of the results for the positive framing group suggest that the differences are significant. To investigate further, we performed an independent samples two-proportions test between the neutral and positive framing groups. The result shows a significant difference between the two groups ($z = 2.425, CI 0.20; .185, p (one-sided) = .008$). This result contradicts our hypothesis H1.2. Compared to neutral framing, positive framing affected the

Table 1: For the nine test conditions, we display the number of participants who encountered the condition, the absolute number and percentages of those that accepted the app permission request, and the mean and standard deviation for feeling informed and trustworthiness on a scale from -2 (strongly disagree) to 2 (strongly agree).

PERMISSION	FRAMING	nb	accepted	info. mean (SD)	trust mean (SD)
location	negative	84	94%	0.52 (1.23)	0.87 (0.93)
	neutral	114	96%	0.39 (1.23)	0.90 (0.87)
	positive	96	91%	0.74 (0.97)	0.97 (0.83)
storage	negative	94	57%	-0.36 (1.32)	-0.12 (1.13)
	neutral	90	43%	-0.66 (1.35)	-0.18 (1.20)
	positive	100	31%	-0.40 (1.50)	-0.55 (1.22)
camera	negative	103	42%	-0.28 (1.41)	-0.22 (1.26)
	neutral	86	37%	0.12 (1.44)	-0.26 (1.29)
	positive	81	32%	-0.16 (1.59)	-0.28 (1.25)

ALL	negative	neutral	positive
	105	110	133.00
declined	115.32	119.01	113.68
	-1.50	-1.30	2.90
	176	180	144
accepted	165.68	170.99	163.33
	1.50	1.30	-2.90

STORAGE	negative	neutral	positive
	40	51	69
declined	52.96	50.70	56.34
	-3.30	0.08	3.17
	54	39	31
accepted	41.04	39.30	43.66
	3.30	-0.08	-3.17

Figure 3: Right-tailed χ^2 test of connection between framings and accepting/declining the app permission request with observed and expected frequency, adjusted residuals. $\chi^2 (2, 848) = 8.294, p = .016$

disclosure behaviour towards declining. To obtain a more fine-grained impression, we looked independently into each permission request type (location, storage, and camera) and present the results in the following.

LOCATION	negative	neutral	positive
	5	5	9
declined	5.43	7.37	6.20
	-0.23	-1.15	1.41
	79	109	87
accepted	78.57	106.63	89.80
	0.23	1.15	-1.41

Figure 4: Right-tailed χ^2 test on location permission for three framings with observed and expected frequency, adjusted residuals. $\chi^2 (2, 294) = 2.1965, p = 0.333$

Location. The location permission request obtained high acceptance rates above 90% across all three framings (Table 1). A right-tailed chi-square test did not reveal a statistically significant difference between the framing conditions $\chi^2 (2, 294) = 2.1965, p = .333$ (Figure 4). To test if positive and negative framing of app permission requests affects people’s disclosure behaviour towards acceptance in the case of uncertainty (H1.1 and 1.2), we separately examined the two uncertainty conditions (i.e., storage and camera) regarding the participants’ disclosure behaviour under the three framing conditions.

Figure 5: Right-tailed χ^2 test on storage permission for three framings with observed and expected frequency, adjusted residuals. $\chi^2 (2, 284) = 13.7832, p = .001$

Storage. For the storage permission request, we found the highest acceptance rate (57%) when the request was framed negatively, a lower rate for neutral framing (43%), and the lowest rate for positive framing (31%)¹. The right-tailed chi-square test shows a significant difference between observed and expected frequencies $\chi^2 (1, 284) = 13.7832, p < .001, effect size (w) = 0.184$ (Figure 5), indicating that the observed frequencies are significantly different from what would be expected if the null hypothesis were true. For a posthoc analysis, we transformed the adjusted residuals into the p-values using the right-tailed probability of the chi-squared distribution function. We did not find significant differences between expected and actual frequencies for the neutral condition ($p = .98$).

However, the differences in acceptance rate were significant both for negative and positive framing (in particular, $p < .001$ for negative framing and $p = .002$ for positive framing), and both values are significant under Bonferroni adjustment of p-level = .017. The pairwise results reveal significant differences between positive and neutral conditions ($z = -1.760, CI -.274;.015, p (one-sided) = .039$). However, the direction of effect from positive framing is contrary to our hypothesis, suggesting that positive framing in the case of storage permission decreases permission acceptance. We also obtained significant results of differences in neutral and negative conditions ($z = -1.914, CI -.280;.015, p (one-sided) = .028$), but both results are not standing against the Bonferroni p-value adjustment of .017. However, the data demonstrated significant differences between positive and negative framing ($z = -3.710, CI -.398;-.127, p < .001$) (significance retained with a Bonferroni-adjusted p-level of .017). This confirms the existence of a decreasing acceptance

rate effect in the positive framing condition and an increasing acceptance rate effect in the negative framing condition.

CAMERA	negative	neutral	positive
	60	54	55
declined	64.47	53.83	50.70
	-1.16	0.05	1.18
	43	32	26
accepted	38.53	32.17	30.30
	1.16	-0.05	-1.18

Figure 6: Right-tailed χ^2 test on camera permission for three framings with observed and expected frequency, adjusted residuals. $\chi^2(2, 270) = 1.805, p = .406$

Camera. For the camera permission request, we obtained the highest acceptance rate (42%) when the request was framed negatively, a lower rate for neutral framing (37%), and the lowest for the positive framing condition (32%) 1. The results of the right-tailed chi-square test $\chi^2(2, 270) = 1.805, p = .406$ do not show a significant association between framing condition and acceptance rate (Figure 6), suggesting that differences are likely by chance.

Results regarding RQ1. Our results do not confirm our hypothesis H1.2 and only partially confirm H1.1. We found some evidence that **negative framing increased the acceptance rate for the storage app permission (low-medium uncertainty), which supports H1.1. Contrary to H1.2,** we found some evidence that **positive framing decreased the acceptance rate** for the storage app permission. However, we did not find a significant effect of the framing on disclosure behaviour in the case of the camera app permission (high uncertainty) (Table 2).

5.2 Effect of framing and uncertainty on the perception of being informed and trustworthiness

	trust		informed	
	H(2)	p-value	H(2)	p-value
location	0.45	p = 0.80	3.80	p = 0.15
storage	8.50	p = 0.014	2.84	p = 0.24
camera	0.09	p = 0.95	3.61	p = 0.16

Figure 7: Result of Kruskal-Wallis analysis of perceived trustworthiness and perception of being informed ratings across permission request types

To answer RQ2, we analysed how permission request framing influences the perception of being informed and the perceived trustworthiness of the request. We performed a Kruskal–Wallis equality-of-populations rank test across the rating means for the three framing conditions (negative, neutral, positive). The results (Figure 7) showed no significant effect of framing on the participants’ perception of being informed ($H(2) = 1.197, p = .550$) nor the perceived trustworthiness of the permission request ($H(2) = 2.412, p = .299$). To obtain a more fine-grained impression, we looked inde-

	odds r.	z	p> z	95% conf.int.	
trust	2.06	6.96	0.00	1.48	2.64
condition (sto pos)	-0.17	-0.40	0.69	-1.03	0.69

Figure 8: Mediation analysis for permission acceptance [Dep Var], and feeling trust and condition (baseline storage neutral) [Indep. Var]. Significant results in bold. Observations = 190, LR $\chi^2(2) = 116.80; Prob > \chi^2 = 0.000; Pseudo R^2 = 0.4670$

pendently into each permission request type (location = high certainty, storage and camera = low/medium certainty). We performed a Kruskal–Wallis equality-of-populations rank test for the three permissions across the framing conditions. The results showed no significant effect of framing on users’ perception of being informed. We found no effect of framing on the perceived trustworthiness of the location and camera permission request, either.

In the case of the storage permission request, a significant difference in trustworthiness was found when comparing neutral and positive framing ($H(2) = 8.500, p = .014$). The posthoc analysis using Mann-Whitney paired tests showed significant differences in the level of perceived trustworthiness between positive and negative framing ($U = 31.320, SE = 11.473, p = .006$) and positive and neutral framing ($U = 25.602, SE = 11.604, p = 0.27$). The difference between positive and negative framing is significant after applying a Bonferroni-adjusted p-value = .017. This result suggests that positive framing makes the permission request appear less trustworthy than negative and potentially also neutral framing. This result aligns with the effect on disclosure behaviour; namely, positive framing led to a decreased acceptance of the storage app permission request.

A mediation analysis confirmed that for the storage permission request, the effect of positive framing on disclosure was fully mediated by the observed reduction in trust. From Model 1, we conclude that the marginally significant odds of disclosure decrease by 41% for positive framing relative to neutral framing. In Model 2, including Trustworthiness and the framing condition, we found that the odds of disclosure for positive framing, relative to the neutral category, nearly disappeared to 16% and remain only marginally significant (Figure 8).

Results regarding RQ2 H2.1-2.2. We found **no significant effect of negative or positive framing on the perception of being informed**, thus refuting hypotheses 2.1a and 2.1b]. We found **weak evidence that positive framing decreases the perceived trustworthiness** of app permission requests, thus partly confirming hypothesis 2.2a. We also found that **negative framing decreases the perceived trustworthiness**, thereby confirming hypothesis 2.2b (Table 2).

5.3 Predictors of disclosure behaviour

Finally, in response to RQ2, we analysed whether the perception of being informed and the perceived trustworthiness of the request are associated with disclosure behaviour (H2.3a, H2.3b). A logistic regression, including the participants’ demographics, indicates that only the condition (permission type plus framing) and perceived trustworthiness were significantly associated, or predictive, of the participants’ disclosure behaviour, $LR \chi^2(6) = 493.97; p <$

	odds r.	z	p> z	95% conf.int.	
age	1.00	-0.35	0.73	0.98	1.01
education	0.97	-0.29	0.77	0.71	1.29
gender (female)	0.74	-1.40	0.16	0.49	1.13
condition (loc neg)	0.84	-0.06	0.81	0.20	3.45
condition (sto neg)	0.15	-3.34	0.00	0.05	0.45
condition (cam neg)	0.07	-4.77	0.00	0.02	0.20
condition (sto neu)	0.07	-4.67	0.00	0.02	0.21
condition (cam neu)	0.05	-5.00	0.00	0.02	0.17
condition (loc pos)	0.31	-1.82	0.07	0.09	1.10
condition (sto pos)	0.05	-5.02	0.00	0.02	0.17
condition (cam pos)	0.04	-5.58	0.00	0.01	0.12
informed	0.88	-1.40	0.16	0.74	1.05
trust	4.97	12.07	0.00	3.83	6.44

Figure 9: Log. Regression for permission acceptance [Dep Var] and age, education, gender (baseline male), condition (baseline location neutral), feeling informed, feeling trust [Indep. Var]. Significant results in bold. Obs. = 839, LR $\chi^2(6) = 555.91$; Prob > $\chi^2 = 0.000$; Pseudo R2 = 0.4895

.001; condition and perceived trustworthiness account for close to 43% of the variance in disclosure behaviour (Figure 9). We find that higher levels of perceived trustworthiness lead to nearly a six-fold (5.89) increase in the odds of accepting the app permission requests. Demographics seem not correlated to or predictive of people’s disclosure behaviour.

Results regarding RQ2 H2.3a and H2.3b. The perception of being informed does not affect disclosure behaviour. We find that **perceived trustworthiness is associated with higher acceptance rates.**

5.4 Self-reported reasons for disclosure choices

In this section, we report the participants’ answers to the two open-ended questions concerning (i) the reasons why they accepted or declined the app permission and (ii) their general thoughts and experiences regarding such permission requests from apps. The answers integrate the quantitative results by providing insights into the people’s reasoning that underlies their disclosure behaviour. As a general observation, participants’ imagined app usage was the most cited primary factor influencing their decision, which confirms their immersion in the test scenario. In the following, the qualitative results are presented according to the six topical categories that were extracted.

Privacy concerns. The participants understood that app permission requests serve to protect app users from privacy and security risks. Of our participants, 14% explicitly expressed privacy and security concerns related to app permissions. Only 3% stated they do not care about privacy. While 13% were wary of being exposed to such consent decisions, 4% of the participants saw them as a privacy-enhancing mechanism, and 8% expressed satisfaction with their current implementation.

Framing. The framing of the requests did not seem to draw the participants’ attention. Less than one per cent mentioned the messages. Among them, some found the messages very convincing, leading them to grant permission (three in the location negative framing, three in the camera negative framing, and four in the storage positive framing condition). Only four participants, in positive framing conditions (two for storage and two for camera), explicitly declined because the message framing was perceived as off-putting.

Accept by default. In our sample, 7% stated they always accept cookies, terms and conditions and similar requests out of convenience. Another 8% accepted the permission request because they had already encountered and accepted similar requests and hence did not see a reason to decline it this time. Some participants (7%) stated that they would withdraw and manage the permission later if it turned out unnecessary for their app use. Others also believed the permission becomes inactive when they stop using the app.

Decline by default. A non-negligible number of app users, 19% in our sample, distrust app providers and cautiously protect their data. They declared that they usually decline first and see if this hampers the app usage. These users prefer permission prompts in context and fine-grained permission settings that grant access only during app use. Some of these users reported that they uninstall apps when they face an app permission request that appears unnecessary. They prefer to look for a less privacy-invasive equivalent. They also reported that they limit the number of apps on their phones to avoid granting permissions excessively.

Reflected decision. Nearly 10% of the participants thought the app would not work if they did not grant all permissions. 44% believed the permissions were somewhat necessary for a proper app functioning. In the general comment on app permissions, about 20% stated that their disclosure behaviour depended solely on the app type and the purpose of the permission, thus pointing to a somehow reflected disclosure choice. About 16% reported reasons for their disclosure behaviour in the experiment related to their app use. Another 15% reported reflective behaviour, but they did not reach a satisfying conclusion and were uncertain about their choice. The participants who reported their reflections understood why the app would ask for access to location and camera and what it was needed for. However, this was less the case for storage. Some incorrectly assumed that granting access to location equals allowing GPS to function, that access to the camera is necessary to localise the phone, and that navigation apps need access to storage to search for a place.

Perception of being informed. 10% of the participants explicitly pointed to a lack of information in the permission requests. Participants wished to know more about why the permission was requested (17%) and what would happen to their data (5%).

Trustworthiness. For 4% of the participants, the privacy choice does not happen on the level of app permissions but is associated with the act of app installation. They said that when they decide to install an app, they decide to trust its provider and that this trust is then also given to the permission requests. 3% reported that they trust known app providers more than unknown ones. To assess whether an unknown app is trustworthy, 2% reported

	Hypothesis	Result	Explanation
1.1	<i>Negative framing affects people’s disclosure behaviour towards accepting in the case of uncertainty.</i>	Confirmed for one uncertain condition	Negative framing increased app permission acceptance in the storage condition (some evidence).
1.2	<i>Positive framing affects people’s disclosure behaviour towards acceptance in the case of uncertainty.</i>	Not confirmed, statistically significant results contrary to expectations in one uncertain condition	Positive framing negatively affected the acceptance of app permission in storage condition (some evidence); positive framing negatively affected the acceptance of the app permission regardless of certainty condition.
2.1a,b	Framing and the uncertainty associated with app permission requests influence people’s perception of being informed.	Not confirmed	No significant effect of negative or positive framing on the perception of feeling informed.
2.2a	<i>Positive framing influences the perceived trustworthiness of the request.</i>	Confirmed for one uncertain condition	Positive framing decreased trustworthiness compared to neutral framing.
2.2b	<i>Negative framing influences the perceived trustworthiness of the request.</i>	Confirmed	Negative framing was perceived as significantly more trustworthy than positive framing.
2.3a	People’s perception of being informed is associated with their disclosure behaviour.	Not confirmed	The perception of being informed was not associated with higher acceptance rates.
2.3b	The perceived trustworthiness of the request is associated with their disclosure behaviour.	Confirmed	Perceived trustworthiness is associated with higher acceptance rates.

Table 2: Overview of the results

relying on app store ratings and user comments. 2% also reported gathering additional feedback from web searches or acquaintances before installing an app from an unknown provider. From 2% of the participants’ answers emerges that some people believe the app stores verify privacy-invasive features before allowing an app into the app store. The trust in the Apple store was even more pronounced than in Android.

6 DISCUSSION

6.1 Influence of framing and uncertainty on people’s disclosure behaviour

In the study that we built on [11], users were highly uncertain about the camera app permission request, relatively uncertain about the storage request, and certain about the location request. We found that negative framing increased the acceptance rate for the storage app permission: alerting users of the risk of a degraded user experience seems to have steered their decision to provide that permission. This result may tie in with dark pattern types, like confirmshaming, that use negative emotions such as anxiety or guilt to pressure people into performing a certain action [14]. The qualitative results further suggest that negative framing instils fear of improper functioning and the belief that installing the app without granting permission makes the app useless.

Contrary to what we expected based on the literature, e.g., [39, 56], positive framing resulted in a lower acceptance rate in our affected condition (storage). Indeed, the design of the prompts induces “goal framing” which, according to Levin et al. [36], is subject to loss aversion for which Tversky and Kahneman saw choice reversal

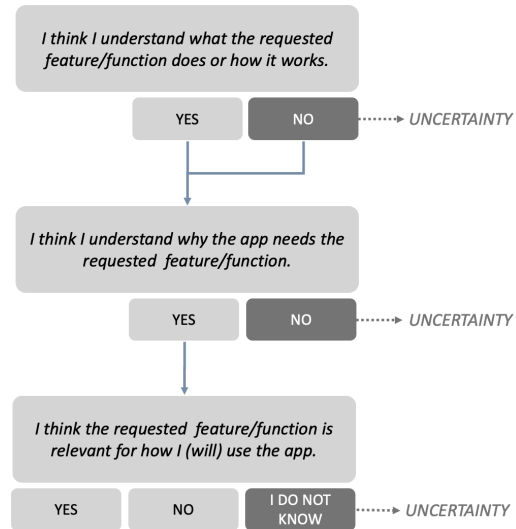


Figure 10: Steps where uncertainty can arise and play a part in app disclosure behaviour.

behaviour [63]. This effect has also been observed in recent studies [5, 41] in which negative framing increased, and positive framing reduced information disclosure. A study outside the loss-gain framing context also found that justifications regarding personal data disclosure lower disclosure rates [32]. This aligns with our participants’ comments that the positively framed message raised

their suspicion and might hence explain the increased declining behaviour.

While we found an effect of framing on disclosure behaviour for the storage permission, we found no effect in the other uncertain case, namely the camera app permission. We suspect people are less uncertain about the camera permission request than what we assumed following the results of [11]. A possible explanation may be that uncertainty does not only arise from people's understanding of why permission is asked for and their assessment of its relevance to their future app use: an additional factor may play a role. In our participants' open answers, we found that people's knowledge of the underlying technology also makes them more or less certain about their disclosure decision. For example, in the case of the camera, in [11], many users did not understand why this app permission was requested. However, the answers of our participants showed that they still have a good comprehension of what the camera function does. They might have been uncertain about the prompt, they were certain they did not want anyone to see their environment and were consequently certain that declining was the choice more in line with their preferences. Regarding storage, our participants had a less clear understanding of the associated functionalities. Some even had an erroneous understanding of the concerned feature. It was less evident what providing access to storage implies, so they were less certain about their disclosure choice. Building on [11], we propose a novel three-level uncertainty model in disclosure requests, as shown in Figure 10.

6.2 Influence of framing and uncertainty on people's perception of being informed and trustworthiness

Regarding the influence of framing and uncertainty on the perceived trustworthiness, we found a statistically significant effect in one condition: the perceived trustworthiness of the storage app permission request was influenced negatively when the request was framed positively. This shows that even dark patterns can backfire and end up damaging the trustworthiness of the systems that implement them. Moreover, such an insight can be relevant for other commercial and non-commercial actors in various settings, for instance, in the design of cookie consent requests (where research shows a predominance of positively framed messages [31, 55]), as well as other consent designs in mobile and IoT settings. Finally, accounting for the possible impact of the framing of benefits and risks on perceived trustworthiness could also be relevant for conceiving data-sharing requests in research contexts, like the biomedical one. Context-specific research would, however, be necessary to evaluate whether our findings can be generalised and to determine other factors at play on perceptions and decisions in such settings.

Regarding the effect of framing and uncertainty on people's perception of being informed, we hypothesised that the experimental conditions employing framing techniques might lead to an erroneous user perception of being better informed when compared to the condition without framing, because the message was longer and included the consequences of accepting or declining the permission requests. Yet, we did not find this result in our study, possibly because the length of our permission prompts did not vary sufficiently in the three conditions. As shown in the qualitative

results, one-third of the participants wished that the permission prompts included more information on why personal data is harvested. It seems relevant to investigate whether permission requests that include more information could lead to a false sense of being better informed or protected, thus potentially manipulating users into accepting the request. Recent comparable research [60] that attempted to make responses to cookie permission requests more informed showed that providing additional data protection details did not have an empowering effect on the participants - on the contrary, it decreased their vigilance and made them feel more secure without any reason.

Our results expose a tension that is well-known in the privacy domain. Providing complete information is considered a valuable instrument of self-determination as it offers the necessary arguments to choose in favour or against a certain option. Without such necessary details, a privacy decision cannot be truly informed - and this is also reflected in the wishes of our participants. However, experimental evidence shows that providing additional information may instil a false sense of security and thus be potentially misleading, thereby nullifying the beneficial effects of transparency toward users. There might be an under-exploited potential for crafting explanations to help users understand technical processes. For example, research investigating how to explain encryption to non-expert technology users found that text describing encryption positively affected understanding and perceived security [19].

6.3 Systemic challenges related to framing dark patterns in the privacy context

Ensuring an adequate level of user protection remains challenging if the app design practice continues to be primarily based on invasive data access that fuels a predominant business model entrenched in massive tracking, mainly for profiling and advertising purposes. "Best practices" concerning how to get users to accept permission requests are widely shared on the web and are also known in the marketing and consent management industry as "consent optimization" or "opt-in optimisation." For example, blog posts aimed at developers illustrate the "benefit explanation" (i.e., mentioning the gains for the users before showing the question) and encourage A/B testing to understand which message (or design, image, timing) is more likely to make users accept app permission requests [47]. That said, the intention of app developers is not necessarily malicious: a recent study [62] for example, shows that developers are conscious that an excessive number of permission requests is associated with a negative UX and shared concerns that unnecessary permissions may hamper trust and the app's reputation. However, they also pointed to their confusion on the scope of permissions and the use of third-party libraries for permission management that can cause them to ask for more permissions than strictly necessary.

While there are also examples of online resources encouraging developers to create informative permission requests, app developers have attempted to work around the restrictions. Increased standardisation, clearer policy and effective enforcement could help mitigate the risks of manipulation impacting individuals. In this regard, legislators seem to be concerned and are therefore shaping legislative answers that reflect their growing awareness of how online decision-making can be distorted by the design of technical

systems. For example, the Digital Services Act prohibits “repeatedly requesting that the recipient of the service make a choice where that choice has already been made, especially by presenting pop-ups that interfere with the user experience” (Art. 25 (3)); while the Data Governance Act greatly emphasises the necessity of standardisation for data sharing consent processes. Regulatory measures, however, seem insufficient since non-compliance with legal requirements is the norm rather than the exception. The gatekeepers (Google and Apple for their app stores) should thus exert greater oversight on the data practices of the apps they make available to dozens of millions worldwide. The situation is worrying: 80% of 2020’s top 10000 downloaded apps collect data for purposes unrelated to their functionality, and this tendency is growing [8]. The app stores could also decide to prioritise the search results and give higher visibility to equivalent privacy-friendly applications. An example is the suite of privacy-enhancing apps created by a German research group [51].

6.4 Strategies for better privacy permissions

Informative app permissions may be necessary to facilitate the creation of correct mental models on the functioning of an app. Still, many other factors are at play when users decide whether to accept or not. For example, considering the amount and the timing of permission requests per app scaled to the number of apps on a mobile (scaled to the number of devices one person may individually or collectively use and to the number of third parties possibly processing the data), one may wonder if it can be realistically expected that users cautiously weigh the pros and cons of each request they receive. It seems promising to investigate further how privacy permissions can be communicated in a way that improves user understanding and counters manipulative practices that benefit app developers and third parties.

Request in context. Privacy permission requests are used in various contexts beyond app permissions. For instance, websites and web apps request permission to access the location, the microphone or the camera through the browser. Recent efforts in the context of web permissions have highlighted the opportunity of using user interfaces with non-prompt, contextual permissions [65]. Authorisations to data access should be asked in context to make them more meaningful, including at the right timing and in an appropriate channel [57]. A digestible amount of information can clarify the pros and cons of acceptance, at least in those cases where asking the individual to choose is worthwhile for them, and an effective deliberation can occur.

Standardisation. Standardisation seems promising in the context of data requests. For instance, the wording used to ask users for consent could be standardised or at least regulated to ensure no loss-gain framing messages can be used, even though it can be cumbersome to define exactly what a “neutral” presentation of choices may be and implement it. For instance, to avoid a biased positive or negative presentation of the options, both pros and cons could be equally provided - however, how to do it succinctly remains a challenge. To this end, the W3C recommends further user research into the granularity and presentation of purpose specifications to provide transparency while keeping friction for users low [65].

Graphics. Besides textual explanations, enhanced transparency could entail adding a graphical label to the notice informing whether the resource is critical to the app’s functioning or optional. Privacy ratings³ that communicate to users the risk inherited with the acceptance may also be a viable manner to facilitate more thoughtful data disclosure behaviour, even though habituation to the graphical label [64] and permission request nagging from apps could annihilate the desired effect.

Global permission setting. Withdrawing granted permissions should also be made easy and reminded to users periodically. A promising evolution on Android OS 11 or higher is the “auto-reset permissions from unused apps” that automatically resets the sensitive permissions that users have granted after a few months of non-use. That said, it may be unrealistic and autonomy-constraining for lay users to be forced to continuously take granular decisions that affect the app’s functioning and privacy, also interfering with their primary objective of using the app. Hence, a default setting in the operating system that grants access to necessary resources during app usage, limited to that period, seems the most user-friendly option. Alternatively, setting the permissions for certain types of data or purposes at the OS level that cascade down to similar cases may be a useful support to decision-making, especially if the decisions can be automatised through the involvement of personalised privacy assistants [38] that manage data use based on user’s preferences and behaviours. A previous study [37] showed that it is possible to cluster users according to their approach to allow/decline permission requests, and this could be valuable knowledge to create profiles and derive reasonable attitudes to respond to such requests. Standardisation and interoperability of vocabulary are necessary for automating these controls. Many international efforts, such as the W3C’s Data Privacy Vocabulary⁴, are underway.

6.5 Future work

Future work should investigate which other factors contribute to the effect of framing. Framing seems to have a non-negligible influence on privacy decisions, yet more research is needed to comprehensively understand contributing (or moderating) factors. This will allow operating systems and policy-makers to provide evidence-based instructions on the type of explanations to be used in a permission dialogue. More research is also needed to understand how informed user consent is, and alternatives to the present practice of permission requests need to be explored. A similar phenomenon to warning fatigue [64] might influence how users perceive and react to privacy permission dialogues. There are several practical implications if technology users do not meaningfully engage with privacy permission dialogues or read the information provided. First, the information conveyed by such privacy permission dialogues should be evaluated empirically to investigate whether users understand what they consent to. Second, if users indeed take away little from these permission dialogues, it seems promising to change the paradigm surrounding privacy permissions. An example could be to provide users with more high-level controls about privacy permissions at the operating system level of the phone. This

³E.g., <https://www.privacyrating.info/#/>

⁴<https://w3c.github.io/cg-reports/dpvcg/CG-FINAL-dpv-20221205/>

suggestion has the shortcoming that users would not provide authorisations within their use context. Future work should investigate the advantages and disadvantages of such approaches empirically.

7 LIMITATIONS

Our study only provided partial evidence of framing affecting disclosure behaviour. Previous studies showed a medium-to-low effect of different digital nudges, including framing, on users' data disclosure behaviour [30]. Findings also demonstrate that, regarding privacy notices, the effect of the framing can be small if users perceive the results of their privacy decisions as not significant [54]. A possible reason for small effects might be caused by the artificial use context that influences how study participants perceive privacy decisions and dark patterns. The tension between a realistic risk representation (e.g., participants using their own device and risking their own data) and practical and ethical concerns is common to most research in usable privacy and security [18]. In our case, we asked participants to situate themselves in a hypothetical situation. Therefore user behaviour might have been different from when they use their own devices. While this methodological choice provides a controlled experimental setup, the opportunity for replicability and data protection, it does not provide the same perception of risk as in real-life situations. More work is needed to design methodologies to study dark patterns in context, ideally in a long-term set-up, to also account for habituation effects. This paper builds on a previous study [11], which, like the present study, was conducted with a UK-based sample. The results might not apply to other contexts.

8 CONCLUSIONS

This study sheds light on loss-gain framing in app permission requests, a potentially manipulative strategy that can subvert users' privacy decision-making. Contrasting results exist in literature about the effects of loss-gain framing in dark patterns and in official guidelines that prohibit their use which encouraged us to seek a clear answer: legislative efforts that attempt to regulate UI/UX design and system design are growing, also thanks to the contribution of empirical studies in usable privacy that shed light on the effects of controversial practices. Our findings partially confirm that negative framing can steer users towards acceptance, whereas positive framing may have the opposite effect, i.e., discourage acceptance behaviour. We also found that positive framing affected the perceived trustworthiness of an app permission request negatively. Based on our results, we provided recommendations and reflections on how to design consent and app permissions that help counter the effect of dark patterns on disclosure behaviour in privacy-sensitive contexts.

ACKNOWLEDGMENTS

This publication is part of the DECEPTICON project supported by the Luxembourg National Research Fund (FNR) (grant no. IS/14717072). Author 3 has carried out the research while being employed at SnT, University of Luxembourg and wishes to acknowledge BRIEF - Biorobotics Research and Innovation Engineering Facilities financed by NextGenerationEU under grant number "IR0000036" - CUP J13C22000400007. The last author acknowledges support by dtcc.bw - Digitalization and Technology Research Center of the

Bundeswehr. dtcc.bw is funded by the European Union - NextGenerationEU.

REFERENCES

- [1] Alessandro Acquisti, Idris Adjerid, Rebecca Balebako, Laura Brandimarte, Lorrie Faith Cranor, Saranga Komanduri, Pedro Giovanni Leon, Norman Sadeh, Florian Schaub, Manya Sleeper, Yang Wang, and Shomir Wilson. 2017. Nudges for Privacy and Security: Understanding and Assisting Users' Choices Online. *ACM Comput. Surv.* 50, 3, Article 44 (2017), 41 pages. <https://doi.org/10.1145/3054926>
- [2] Alessandro Acquisti, H Heinz, and Jens Grossklags. 2005. Uncertainty, ambiguity and privacy. In *4th Annual Workshop on Economics and Information Security (WEIS)*.
- [3] Idris Adjerid, Alessandro Acquisti, Laura Brandimarte, and George Loewenstein. 2013. Sleights of privacy: Framing, disclosures, and the limits of transparency. In *Proceedings of the ninth symposium on usable privacy and security*. 1–11.
- [4] Idris Adjerid, Alessandro Acquisti, and George Loewenstein. 2019. Choice architecture, framing, and cascaded privacy choices. *Management Science* 65, 5 (2019), 2267–2290.
- [5] Paritosh Bahirath, Martijn Willemsen, Yangyang He, Qizhang Sun, and Bart Knijnenburg. 2021. Overlooking context: How do defaults and framing reduce deliberation in smart home privacy decision-making?. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–18.
- [6] Jan M Bauer, Regitze Bergström, and Rune Foss-Madsen. 2021. Are you sure, you want a cookie?—The effects of choice architecture on users' decisions about sharing private online data. *Computers in Human Behavior* 120 (2021), 106729.
- [7] Benjamin Maximilian Berens, Heike Dietmann, Chiara Krisam, Oksana Kulyk, and Melanie Volkamer. 2022. Cookie Disclaimers: Impact of Design and Users' Attitude. In *Proceedings of the 17th International Conference on Availability, Reliability and Security*. 1–20.
- [8] Bo Bian, Xinchun Ma, and Huan Tang. 2021. *The Supply and Demand for Data Privacy: Evidence from Mobile Apps*. Number ID 3987541 in 1. SSRN, Rochester, NY. <https://doi.org/10.2139/ssrn.3987541>
- [9] Tim Biggs. Accessed 2023-01-23. 'If in Doubt, Say No': Why Phone Apps Want Permission to Use Your Personal Data. Available online at <https://www.smh.com.au/technology/if-in-doubt-say-no-why-phone-apps-want-permission-to-use-your-personal-data-20191106-p53813.html>.
- [10] European Data Protection Board. 2018. *Opinion 5/2018. Preliminary Opinion on Privacy by Design*. The European Data Protection Supervisor (EDPS). https://edps.europa.eu/sites/edp/files/publication/18-05-31_preliminary_opinion_on_privacy_by_design_en_0.pdf
- [11] Kerstin Bongard-Blanchy, Jean-Louis Sterckx, Arianna Rossi, Verena Distler, Salvador Rivas, and Vincent Koenig. 2022. An (Un)Necessary Evil - Users' (Un)Certainty about Smartphone App Permissions and Implications for Privacy Engineering. In *2022 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)*. IEEE, Genoa, Italy. <https://doi.org/10.1109/EuroSPW55150.2022.00023>
- [12] Christoph Bösch, Benjamin Erb, Frank Kargl, Henning Kopp, and Stefan Pfattheicher. 2016. Tales from the Dark Side: Privacy Dark Strategies and Privacy Dark Patterns. *Proc. Priv. Enhancing Technol.* 2016, 4 (2016), 237–254.
- [13] David M. Boush, Marian Friestad, and Peter Wright. 2009. *Deception In The Marketplace: The Psychology of Deceptive Persuasion and consumer self-protection* (first edition ed.). Routledge.
- [14] Harry Brignull. Accessed 2023-01-18. Deceptive Design - Types of Deceptive Design. Available online at <https://www.deceptive.design/types>.
- [15] Michael Chromik, Malin Eiband, Sarah Theres Völkel, and Daniel Buschek. 2019. Dark Patterns of Explainability, Transparency, and User Control for Intelligent Systems.. In *IUI workshops*, Vol. 2327.
- [16] Danielle Keats Citron and Daniel J. Solove. 2022. Privacy harms. *BUL Rev.* 102 (2022), 793.
- [17] Anzo DeGiulio, Hanoom Lee, and Eleanor Birrell. 2021. "Ask App Not to Track": The Effect of Opt-In Tracking Authorization on Mobile Privacy. In *International Workshop on Emerging Technologies for Authorization and Authentication*. Springer, 152–167.
- [18] Verena Distler, Matthias Fassl, Hana Habib, Katharina Krombholz, Gabriele Lenzini, Carine Lallemand, Lorrie Faith Cranor, and Vincent Koenig. 2021. A Systematic Literature Review of Empirical Methods and Risk Representation in Usable Privacy and Security Research. *ACM Transactions on Computer-Human Interaction (TOCHI)* 28, 6 (2021), 1–50.
- [19] Verena Distler, Tamara Gutfleisch, Carine Lallemand, Gabriele Lenzini, and Vincent Koenig. 2022. Complex, but in a good way? How to represent encryption to non-experts through text and visuals - Evidence from expert co-creation and a vignette experiment. *Computers in Human Behavior Reports* 5 (2022), 100161. <https://doi.org/10.1016/j.chbr.2021.100161>
- [20] European Data Protection Board. 2023. Guidelines 3/2022 on Deceptive design patterns in social media platform interfaces: How to recognize and avoid them. Version 2.0. available online at

- https://edpb.europa.eu/system/files/2022-03/edpb_03-2022_guidelines_on_dark_patterns_in_social_media_platform_interfaces_en.pdf.
- [21] Directorate-General for Justice, Consumers (European Commission), Francisco Lupiáñez-Villanueva, Alba Boluda, Francesco Bogliacino, Giovanni Liva, Lucie Lechardey, and Teresa Rodríguez de las Heras Ballell. 2022. *Behavioural study on unfair commercial practices in the digital environment. Dark patterns and manipulative personalisation: final report*. Publications Office of the European Union, LU. <https://data.europa.eu/doi/10.2838/859030>
 - [22] Joshua Gluck, Florian Schaub, Amy Friedman, Hana Habib, Norman Sadeh, Lorrie Faith Cranor, and Yuvraj Agarwal. 2016. How short is too short? implications of length and framing on the effectiveness of privacy notices. In *Twelfth symposium on usable privacy and security (SOUPS 2016)*. 321–340.
 - [23] Jingjing Gong, Yan Zhang, Zheng Yang, Yonghua Huang, Jun Feng, and Weiwei Zhang. 2013. The framing effect in medical decision-making: a review of the literature. *Psychology, health & medicine* 18, 6 (2013), 645–653.
 - [24] PAJ Graßl, HK Schraffenberger, FJ Zuiderveen Borgesius, and MA Buijzen. 2021. Dark and bright patterns in cookie consent requests. *Journal of Digital Social Research* 3, 1 (2021), 1–38.
 - [25] Siddharth Gulati, Sonia Sousa, and David Lamas. 2017. Modelling trust: An empirical assessment. In *Human-Computer Interaction—INTERACT 2017: 16th IFIP TC 13 International Conference, Mumbai, India, September 25–29, 2017, Proceedings, Part IV* 16. Springer, 40–61.
 - [26] Siddharth Gulati, Sonia Sousa, and David Lamas. 2019. Design, development and evaluation of a human-computer trust scale. *Behaviour & Information Technology* 38, 10 (Oct. 2019), 1004–1015. <https://doi.org/10.1080/0144929X.2019.1656779>
 - [27] Johanna Gunawan, Cristiana Santos, and Irene Kamara. 2022. Redress for Dark Patterns Privacy Harms? A Case Study on Consent Interactions. In *Proceedings of the 2022 Symposium on Computer Science and Law*. 181–194.
 - [28] Hana Habib, Yixin Zou, Aditi Jannu, Neha Sridhar, Chelse Swoopes, Alessandro Acquisti, Lorrie Faith Cranor, Norman Sadeh, and Florian Schaub. 2019. An Empirical Analysis of Data Deletion and {Opt-Out} Choices on 150 Websites. In *Fifteenth Symposium on Usable Privacy and Security (SOUPS 2019)*. 387–406.
 - [29] Daniel Holliday, Stephanie Wilson, and Simone Stumpf. 2016. User trust in intelligent systems: A journey over time. In *Proceedings of the 21st international conference on intelligent user interfaces*. 164–168.
 - [30] Athina Ioannou, Iis Tussyyadiah, Graham Miller, Shujun Li, and Mario Weick. 2021. Privacy nudges for disclosure of personal information: A systematic literature review and meta-analysis. *PLoS one* 16, 8 (2021), e0256822.
 - [31] Georgios Kampanos and Siamak F Shahandashti. 2021. Accept all: The landscape of cookie banners in Greece and the UK. In *ICT Systems Security and Privacy Protection: 36th IFIP TC 11 International Conference, SEC 2021, Oslo, Norway, June 22–24, 2021, Proceedings*. Springer, 213–227.
 - [32] Bart P Knijnenburg and Alfred Kobsa. 2013. Making decisions about privacy: information disclosure in context-aware recommender systems. *ACM Transactions on Interactive Intelligent Systems (TiiS)* 3, 3 (2013), 1–23.
 - [33] Matthias Kraus, Nicolas Wagner, and Wolfgang Minker. 2020. Effects of proactive dialogue strategies on human-computer trust. In *Proceedings of the 28th ACM Conference on User Modeling, Adaptation and Personalization*. 107–116.
 - [34] Colin R Kuehnhans, Bruno Heyndels, and Katharina Hilken. 2015. Choice in politics: Equivalency framing in economic policy decisions and the influence of expertise. *European Journal of Political Economy* 40 (2015), 360–374.
 - [35] Kun Chang Lee and Namho Chung. 2009. Understanding factors affecting trust in and satisfaction with mobile banking in Korea: A modified DeLone and McLean’s model perspective. *Interacting with computers* 21, 5–6 (2009), 385–392.
 - [36] Irwin P Levin, Sandra L Schneider, and Gary J Gaeth. 1998. All frames are not created equal: A typology and critical analysis of framing effects. *Organizational behavior and human decision processes* 76, 2 (1998), 149–188.
 - [37] Jialiu Lin, Bin Liu, Norman Sadeh, and Jason I Hong. 2014. Modeling {Users’} Mobile App Privacy Preferences: Restoring Usability in a Sea of Permission Settings. In *10th Symposium On Usable Privacy and Security (SOUPS 2014)*. 199–212.
 - [38] Bin Liu, Mads Schaarup Andersen, Florian Schaub, Hazim Almuheimi, Shikun (Aerin) Zhang, Norman Sadeh, Yuvraj Agarwal, and Alessandro Acquisti. 2016. Follow My Recommendations: A Personalized Privacy Assistant for Mobile App Permissions. In *Twelfth Symposium on Usable Privacy and Security (SOUPS 2016)*. USENIX Association, 27–41. <https://www.usenix.org/conference/soups2016/technical-sessions/presentation/liu>
 - [39] Jamie Luguri and Lior Strahilevitz. 2019. Shining a light on dark patterns. *U of Chicago, Public Law Working Paper* 719 (2019).
 - [40] Jamie Luguri and Lior Jacob Strahilevitz. 2021. Shining a light on dark patterns. *Journal of Legal Analysis* 13, 1 (2021), 43–109.
 - [41] Eryn Ma and Eleanor Birrell. 2022. Prospective Consent: The Effect of Framing on Cookie Consent Decisions. In *CHI Conference on Human Factors in Computing Systems Extended Abstracts*. 1–6.
 - [42] Dominique Machuletz and Rainer Böhme. 2019. Multiple purposes, multiple problems: A user study of consent dialogs after GDPR. *arXiv preprint arXiv:1908.10048* (2019).
 - [43] Arunesh Mathur, Gunes Acar, Michael J Friedman, Elena Lucherini, Jonathan Mayer, Marshini Chetty, and Arvind Narayanan. 2019. Dark patterns at scale: Findings from a crawl of 11K shopping websites. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019), 1–32.
 - [44] Arunesh Mathur, Jonathan Mayer, and Mihir Kshirsagar. 2021. What Makes a Dark Pattern... Dark? Design Attributes, Normative Considerations, and Measurement Methods. *arXiv:2101.04843 [cs]* (Jan 2021). <https://doi.org/10.1145/3411764.3445610> arXiv: 2101.04843.
 - [45] Philipp Mayring et al. 2004. Qualitative content analysis. *A companion to qualitative research* 1, 2 (2004), 159–176.
 - [46] Nora McDonald, Sarita Schoenbeck, and Andrea Forte. 2019. Reliability and inter-rater reliability in qualitative research: Norms and guidelines for CSCW and HCI practice. *Proceedings of the ACM on human-computer interaction* 3, CSCW (2019), 1–23.
 - [47] Ella Mullan. Accessed 2023-02-01. iOS Push Notification Permissions: The Best Practices. Available online at <https://blog.hurree.co/blog/ios-push-notification-permissions-best-practices>.
 - [48] Florian Nothdurft, Tobias Heinroth, and Wolfgang Minker. 2013. The impact of explanation dialogues on human-computer trust. In *Human-Computer Interaction. Users and Contexts of Use: 15th International Conference, HCI International 2013, Las Vegas, NV, USA, July 21–26, 2013, Proceedings, Part III* 15. Springer, 59–67.
 - [49] Midas Nouwens, Ilaria Liccardi, Michael Veale, David Karger, and Lalana Kagal. 2020. Dark patterns after the GDPR: Scraping consent pop-ups and demonstrating their influence. In *Proceedings of the 2020 CHI conference on human factors in computing systems*. 1–13.
 - [50] OECD. 2022. *Dark commercial patterns*. Number 336 in OECD Digital Economy Papers. OECD Publishing, Paris.
 - [51] Karlsruhe Institut of Technology KIT. 2022. Privacy Friendly Apps - improved privacy protection on the smartphone. Available online at <https://secuso.aifb.kit.edu/english/105.php>.
 - [52] Sayantan Polley, Rashmi Raju Koparde, Akshaya Bindu Gowri, Maneendra Perera, and Andreas Nuernberger. 2021. Towards trustworthiness in the context of explainable search: A framework for research and design. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 2580–2584.
 - [53] Jens Riegelsberger, M Angela Sasse, and John D McCarthy. 2005. The mechanics of trust: A framework for research and design. *International Journal of Human-Computer Studies* 62, 3 (2005), 381–422.
 - [54] Sonam Samat and Alessandro Acquisti. 2017. Format vs. content: the impact of risk and presentation on disclosure decisions. In *Thirteenth Symposium on Usable Privacy and Security (SOUPS 2017)*. 377–384.
 - [55] Cristiana Santos, Natalia Bielewa, and Célestin Matte. 2020. Are cookie banners indeed compliant with the law??. *Technology and Regulation* 2020 (Dec 2020), 91–135. <https://doi.org/10.26116/techreg.2020.009>
 - [56] Cristiana Santos, Arianna Rossi, Lorena Sanchez Chamorro, Kerstin Bongard-Blanchy, and Ruba Abu-Salma. 2021. Cookie Banners, What’s the Purpose? Analyzing Cookie Banner Text Through a Legal Lens. In *Proceedings of the 20th Workshop on Privacy in the Electronic Society (Virtual Event, Republic of Korea) (WPES ’21)*. Association for Computing Machinery, New York, NY, USA, 187–194. <https://doi.org/10.1145/3463676.3485611>
 - [57] Florian Schaub, Rebecca Balebako, Adam L Durity, and Lorrie Faith Cranor. 2015. A design space for effective privacy notices. In *Eleventh Symposium On Usable Privacy and Security (SOUPS 2015)*. 1–17.
 - [58] Myeong-Gu Seo, Brent Goldfarb, and Lisa Feldman Barrett. 2010. Affect and the framing effect within individuals over time: Risk taking in a dynamic investment simulation. *Academy of Management Journal* 53, 2 (2010), 411–431.
 - [59] David L Streiner. 2015. Best (but oft-forgotten) practices: the multiple problems of multiplicity—whether and how to correct for many statistical tests. *The American journal of clinical nutrition* 102, 4 (2015), 721–728.
 - [60] Joanna Strycharz, Edith Smit, Natali Helberger, and Guda van Noort. 2021. No to cookies: Empowering impact of technical and legal knowledge on rejecting tracking cookies. *Computers in Human Behavior* 120 (Jul 2021), 106750. <https://doi.org/10.1016/j.chb.2021.106750>
 - [61] S Shyam Sundar and Jinyoung Kim. 2019. Machine heuristic: When we trust computers more than humans with our personal information. In *Proceedings of the 2019 CHI Conference on human factors in computing systems*. 1–9.
 - [62] Mohammad Tahaei, Ruba Abu-Salma, and Awais Rashid. 2023. Stuck in the Permissions With You: Developer & End-User Perspectives on App Permissions & Their Privacy Ramifications. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–24. <https://doi.org/10.1145/3544548.3581060> arXiv:2301.06534 [cs].
 - [63] Amos Tversky and Daniel Kahneman. 1981. The Framing of Decisions and the Psychology of Choice. *Science* 211, 4481 (1981), 453–458. <https://doi.org/10.1126/science.7455683> arXiv:https://www.science.org/doi/pdf/10.1126/science.7455683
 - [64] Anthony Vance, Jeffrey L. Jenkins, Bonnie Brinton Anderson, Daniel K. Bjornn, and C. Brock Kirwan. 2018. Tuning Out Security Warnings: A Longitudinal Examination of Habituation Through fMRI, Eye Tracking, and Field Experiments. *MIS Quarterly* 42, 2 (Feb 2018), 355–380. <https://doi.org/10.25300/MISQ/2018/14124>

- [65] W3C. 2022. *Report on the 2022 W3C Workshop on Permissions*. Technical Report. W3C. <https://www.w3.org/Privacy/permissions-ws-2022/report>