The Illusion of Privacy: Investigating User Misperceptions in Browser Tracking Protection

Maximiliane Windl
LMU Munich
Munich, Germany
Munich Center for Machine Learning (MCML)
Munich, Germany
maximiliane.windl@ifi.lmu.de

Roman Amberg
LMU Munich
Munich, Germany
r.amberg@campus.lmu.de

Thomas Kosch
HU Berlin
Berlin, Germany
thomas.kosch@hu-berlin.de

Figure 1: Our study explores user perceptions of web-browser tracking protection plugins under the three conditions no plugin, functional plugin, and placebo plugin, where users were only primed with narratives without altering the website. During a hotel booking task, users felt more protected with functional or placebo plugins despite no actual changes to the website, revealing the participant's inability to judge tracking protection effectiveness accurately. We generated the figure using DALL-E.

Abstract

Third parties track users' web browsing activities, raising privacy concerns. Tracking protection extensions prevent this, but their influence on privacy protection beliefs shaped by narratives remains uncertain. This paper investigates users' misperception of tracking protection offered by browser plugins. Our study explores how different narratives influence users' perceived privacy protection by examining three tracking protection extension narratives: no protection, functional protection, and a placebo. In a study (N=36), participants evaluated their anticipated protection during a hotel booking process, influenced by the narrative about the plugin's functionality. However, participants viewed the same website without tracking protection adaptations. We show that users feel more protected when informed they use a functional or placebo extension, compared to no protection. Our findings highlight the deceptive nature of misleading privacy tools, emphasizing the need for greater

CHI '25, Yokohama, Japan

ACM ISBN 979-8-4007-1394-1/25/04
https://doi.org/10.1145/3706598.3713912

transparency to prevent users from a false sense of protection, as such misleading tools negatively affect user study results.

CCS Concepts

• **Security and privacy** → *Privacy protections*; • **Human-centered computing** → **Human computer interaction (HCI)**.

Keywords

Privacy, Placebo, Protection, Web Extension, Plugins, Tracking

ACM Reference Format:
Maximiliane Windl, Roman Amberg, and Thomas Kosch. 2025. The Illusion of Privacy: Investigating User Misperceptions in Browser Tracking Protection. In *CHI Conference on Human Factors in Computing Systems (CHI '25), April 26–May 01, 2025, Yokohama, Japan.* ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/3706598.3713912

1 Introduction

Tracking on the internet involves monitoring and recording users' online activities, such as the websites they visit, the pages they view, and the links they click [28]. Advertisers, website owners, and data brokers use various methods to track users, including cookies, web beacons, browser fingerprinting, and tracking pixels, often unbeknownst to the user. These techniques allow website operators and data brokers to collect detailed information about users' behavior and preferences, often without the user's explicit consent [26]. As a result, users can be profiled based on their online behavior, leading to targeted advertising and potential misuse of personal information. Although methods for measuring the amount of internet tracking exist [10], users can rarely reflect on the quantity and implications of their tracked information [30]. There are several methods for users to maintain control over their privacy online, such as regularly clearing cookies, enabling *Do Not Track* settings, or utilizing VPNs [23]. In addition, tracking protection plugins have become popular companions on the internet. Examples include *Ghostery*, *uBlock Origin*, *Disconnect*, and *AdGuard*. These extensions promise to prevent user tracking online and, thus, protect user privacy. Such tracking protection plugins promise to safeguard the user's privacy, preventing unauthorized access and misuse of sensitive data, such as personal interests, health conditions, and political beliefs. Preventing tracking also helps avoid profiling and discrimination, such as differential pricing or targeted political ads [35].

However, users cannot perceive the extent of data tracking or the effectiveness of their privacy protections, thus relying on the protection narrative they receive from privacy protection extensions. Similar to placebo in medicine, this invisibility means that users rely on the narratives of privacy tools and plugins to safeguard their online activities based on the protection narrative they receive. While privacy protection extensions like ad blockers demonstrate visible functionality (i.e., blocking ads), tracking protection extensions lack such tangible indicators, leaving users dependent on the provided extension narratives. Previous work assessed the impact of narratives on user performance, finding that users believe that they achieve better results using systems that provide improvements through artificial intelligence. At the same time, no functionality was present [16, 17]. Boot et al. [3] pointed out that user expectations through narratives manipulate user satisfaction and self-assessed performance. This subjective user satisfaction can be further enhanced by presenting control interfaces, giving users a phantom perception of control [37]. Based on previous research, we draw a parallel to the functionality of tracking protection extensions: if protection plugins do not function properly or falsely claim to provide security through a narrative, they act as a placebo. This means users may believe they are protected from tracking, whereas in reality, their data is still being monitored and collected. This misconception can lead to complacency and may distort how users evaluate privacy-preserving systems in study settings. However, previous research did not investigate how narratives shape the perception of data-tracking protection extensions.

In this paper, we examine how users perceive their level of protection when using one of three different protection mechanisms with different tracking protection narratives: (1) No Protection Narrative, (2) Functional Protection Narrative, and (3) a Placebic Protection Narrative (see Figure 1). Participants were informed that the No Protection Narrative offered no tracking protection, while both the Functional Protection Narrative and Placebic Protection Narrative were presented as tracking protection plugins, explained through a narrative. However, only the Functional Protection Narrative actually implemented tracking protection, whereas the Placebic Protection Narrative claimed to offer protection but, in reality, did not offer any protection. In a within-subjects user study with 36 participants, each participant was primed with the respective tracking protection narrative during a hotel booking process, a common task that is associated with providing personal data and data tracking activity [40], where the participants assessed their anticipated and perceived protection before and after the interaction. Furthermore, participants viewed the same website independent of the tracking protection extension used, and we only manipulated the narrative of the currently used plugin. Our results indicate that the perception of privacy protection is significantly enhanced when users are primed with a Functional Protection Narrative or Placebic Protection Narrative compared to a No Protection Narrative. We discuss how promised privacy protection mechanisms may deceive users when these mechanisms are ineffective and advocate for greater functional transparency in privacy protection extensions, noting that studies regarding privacy-preserving systems may be affected by placebic narratives, which need to be carefully controlled.

Contribution Statement

The contribution of this paper is twofold: (1) We demonstrate that participants' perceptions of privacy protection are significantly shaped by the narrative of privacy-preserving systems they are presented with, showing the importance of including a placebo condition in privacy research and (2) we discuss strategies to minimize placebo effects in user studies examining the perceived effectiveness of privacy-preserving systems.

2 Related Work

We report prior work on the use and perception of tracking protection extensions. Next, we summarize findings on how users evaluate interfaces after being primed with a placebo system description.

2.1 Tracking Protection Extensions

Developers and researchers have explored privacy protection by blocking connections to third-party servers and preventing data storage on users' devices. Browser extensions became popular for their easy integration into web browsers [21]. Tracking protection extensions enhance privacy by blocking tracking scripts, cookies, pixels, and web beacons. They prevent device fingerprinting by obfuscating browser and device details and blocking referrer headers. These extensions often include privacy-friendly defaults, such as sending "Do Not Track" signals and clearing cookies after sessions. Popular plugins, such as uBlock Origin[1] or Disconnect[2] provide default settings to maximize privacy without requiring user knowledge. At the same time, tracking protection extensions provide visualizations to inform users about their privacy safety. However, Schaub et al. [30] showed that users are uncertain about the reflected metrics and have difficulty assessing the privacy protection levels. Despite their benefits, tracking-protection browser extensions face scrutiny for their extensive access to users' visited websites, raising concerns about potential privacy risks and information leakage. For example, Starov et al. [32] examined privacy leakage in the 10,000 most popular Google Chrome extensions. They found that many extensions unintentionally leak sensitive user data, like browsing history and search queries, often due to how they handle third-party content. To mitigate this issue, the authors developed BrowsingFog, a browser extension that obfuscates a user's browsing habits, protecting their private information from being inferred by history-stealing trackers. A study by Kariryaa et al. [14] showed that users are unaware of the privileges they provide to the extension providers. The study also found that users trust extension developers but have a limited understanding of how extensions work. This supports previous findings by Schaub et al. [30], who showed that users struggle to interpret the feedback from these extensions, limiting their understanding of the protection provided. Tracking protection extensions can also be used to fingerprint users. Gulyas et al. [11] found that 54.86% of users with detectable extensions are unique, and this increases to 89.23% for those with both an extension and a login, highlighting privacy risks. Despite this, the study suggests the benefits of privacy extensions outweigh the risks and recommends countermeasures. Overall, the balance between providing robust tracking protection and ensuring the extension does not become a source of privacy vulnerability is a critical issue that research must investigate.

2.2 Tracking and Privacy Awareness

The previous section highlighted the importance of informing users about how tracking protection extensions enhance privacy. Users often struggle to assess the level of privacy protection provided, as the extensions' actions are unclear [30]. Many extensions, including uBlock Origin and Disconnect, offer metrics to quantify privacy improvements, and past research has explored ways to make these metrics more understandable for users. Starov and Nikiforakis [33] introduced PrivacyMeter, a browser extension that calculates a privacy score for websites based on their privacy practices relative to other sites, addressing limitations of existing anti-tracking extensions. They found that PrivacyMeter effectively evaluates websites by covering various privacy practices and providing accurate measurements with minimal performance impact. The study highlights the potential of crowdsourcing for privacy research while stressing the need to balance user anonymity with protection against malicious clients. Similarly, Takano et al. [36] presented MindYourPrivacy, a system designed to visualize third-party web tracking and identify packets that infringe on users' privacy. An evaluation of MindYourPrivacy showed that visualizing web tracking significantly enhances users' awareness of privacy issues. Bhattacharjee et al. [2] explored how visualization can facilitate transparency in privacy implications for various stakeholders within the data ecosystem. The study highlights existing gaps and research opportunities in privacy-preserving data visualization, emphasizing the necessity for collaboration among stakeholders. It suggests that developing privacy-focused techniques, policies, and visualization tools is essential to balance privacy and utility in data sharing. In summary, past research has explored how to help users better understand how privacy extensions protect their data. Despite these efforts, users still struggle to assess the effectiveness of these tools [30], often relying on trust in the extensions' functionality. This trust can lead to overreliance and inflated expectations, especially when the extensions fail to perform as intended. We discuss how such user expectations can distort the perception of a system's functionality.

2.3 The Placebo Effect of Interactive Systems

A medical placebo, such as a sugar pill that contains no active ingredients, can enhance a patient's subjective condition [7] without involving any active substance or specific procedure. This placebo can relieve pain [18] or aid in treating various ailments [1], thus offering effective medical treatment without a mechanism specific to the illness. The key factor in the placebo effect is the patient's belief in the placebo's effectiveness, which results in a positive assessment after treatment [24, 34]. Consequently, the placebo treatment must manipulate the user's expectations towards an improvement as a prerequisite for being successful [12]. Increased user expectations towards a novel system may change the subjective and objective perception through a novelty effect. Wells et al. [39] showed that the perceived novelty is a prominent emotional belief significantly influencing the adoption of information technology innovations. In the context of human-computer interaction, studies in gaming experience pioneered investigating placebo effects. Providing users with a game description that adapts the game difficulty according to the players' performance changes the subjective gaming experience [5]. Spiel et al. [31] showed how adaptive difficulty, based on performance and eye movements, affects gameplay experience in TETRIS. The results suggest that eye-movement-based adaptive difficulty does not significantly impact player performance, but the way adaptive difficulty is presented can influence players' game experience and perceived competence. In summary, previous research suggests that placebos do exist in interfaces where users cannot fully gauge a system's functionality. Consequently, users face difficulties in assessing their own competence and the effectiveness or reliability of interfaces. Recent research investigated placebo

[1] https://ublockorigin.com
[2] https://disconnect.me/disconnect

effects in AI systems as users face similar challenges regarding functional transparency. Kosch et al. [17] showed that, although users were interacting with a non-adaptive system throughout the study, they believed to perform better when using an allegedly adaptive AI-based system. Villa et al. [38] investigated how placebic cognitive augmentation systems affect the risk-taking behavior of users, finding that participants were taking on more risk during sham augmentations[3]. Pataranutaporn et al. [27] investigated how participants perceived voice agents with neutral, malevolent, and benevolent intent. The results showed that participants assessed the benevolent voice agent as helpful, although participants were interacting with a neutral voice agent throughout all conditions. Kloft et al. [16] found that participants performed better when they believed an AI-enhanced their task interface, even when no AI was present. Previous research showed that the efficacy of placebos in interactive systems depends on the narrative. In this context, Bosch et al. [4] showed that the narrative about using a specific display refresh rate already skews the performance perception of participants.

2.4 Summary and Research Gap

Prior research explored the functionalities and limitations of tracking protection extensions, showing that users are challenged with understanding the privacy metrics and exact functionality of such extensions [11, 14, 30], letting user rely on the narrative of the extension. Efforts to enhance transparency through visualizations and metrics have demonstrated mixed success, as users often rely on trust rather than fully understanding the protections offered [33, 36]. Additionally, studies on placebo effects have shown that user expectations and perceptions are significantly distorted through narratives of an allegedly functional system, even when the system provides no actual benefit [16, 17, 38]. In these studies, the narrative represented the only variable that was manipulated which had an impact on the participants' subjective perception. However, there is limited research on how the narratives surrounding tracking protection extensions shape users' privacy perceptions. These narratives may influence users' sense of privacy, potentially affecting their actual privacy in real-world settings or skewing research findings that evaluate perceived levels of privacy protection. This study addresses this gap by examining how tracking protection narratives influence users' perceived privacy through the following research question **(RQ): How do tracking protection narratives influence users' perceptions of privacy protection, personalization, and security during online interactions, regardless of their actual functionality?**

3 Methodology

In the context of privacy protection extensions, users may trust the narrative of additional protection through bespoke plugins, yet the extensions may not provide protection. We formulate the following hypotheses to answer our research question.

H1: Users perceive a higher level of privacy protection when using functional and placebic tracking protection extensions.

[3]Sham augmentations refer to changes in a system that appear as if they enhance functionality or performance but, in reality, do not have any real effect.

Table 1: The questions we asked to determine the perceived tracking protection. Q1 – Q3 were asked before interaction with the system to assess if users believed the narratives. We asked the questions Q4 – Q6 post-interaction to evaluate if the user still believed the narrative. We asked all questions on 100-point sliders.

ID	Question
Q1	I think my privacy will be protected.
Q2	I think I will see personalized content
Q3	I think I will feel secure while browsing.
Q4	I felt like my privacy was protected.
Q5	I felt like I saw personalized content.
Q6	I felt secure while browsing.

H2: Users perceive less personalized content when using functional and placebic tracking protection extensions.

H3: Users will feel more secure when using functional and placebic tracking protection extensions.

We designed our study to explore the influence of the protection narrative on user perceptions rather than focusing on observable functionality differences between placebo and functional plugins. Inspired by previous work [6, 17], the primary aim was to investigate whether the protection narrative, Functional Protection Narrative or Placebic Protection Narrative, could shape users' perceived sense of security compared to a No Protection Narrative.

3.1 Participants

We recruited 36 participants through university mailing lists, contact databases for study participation, and snowball sampling. The participants were between 18 and 60 years old ($M = 29.6, SD = 10.6$). Sixteen participants self-identified as female, 20 as male, and most (N = 20) were university students. We used the Internet Users' Information Privacy Concerns (IUIPC) questionnaire [19] to understand participants' general perception of privacy on a 7-point scale. We found an average of 6 ($SD = 1.1$) for Awareness, 5.4 ($SD = 1.5$) for Control, and 5.4 ($SD = 1.2$) for Collection, indicating a comparably high level of privacy concerns (cf. [13]). We compensated participants with €10.

3.2 Procedure

We welcomed the participants upon arrival and asked them to fill out an informed consent form. Then, we recorded their demographics, including age, gender, and occupation. After that, the participants provided their perceived privacy concerns using the IUIPC questionnaire. Furthermore, we asked participants about their knowledge of and current use of privacy protection plugins. Next, we told the participants how cookies, advertising, and tracking work and how privacy-protecting browser extensions try to block third parties from accessing those cookies. We introduced participants to the two browser extensions through their respective narratives: We introduced the Functional Protection Narrative as an established and the Placebic Protection Narrative as a novel tracking protection extension through a short description

of its purpose and functionality. Although the Functional Protection Narrative was actually uBlock Origin, we refrained from disclosing the plugin's name to the participants to prevent any potential biases that might have resulted from people knowing and using the plugin. Consequently, we also removed its name and icon. Additionally, we informed participants that there was a baseline condition that did not include a tracking protection extension (i.e., No Protection Narrative). During the next section, participants were required to answer three single-choice questions to determine if they fully understood the experiment's goal.

The study task was to find and book a hotel that fulfilled certain requirements as quickly as possible. We randomized the task goal, such as whether the hotel should be close to the city center or offer free parking. Participants conducted three booking sessions with each browser plugin in a counterbalanced order. We repeated each plugin's narrative before each condition. We also showed participants that the extension was running and how they could view blocked connections (see Figure 2). Only one plugin was displayed and active simultaneously for each condition. We cleared all browser data between conditions to prevent participants from assuming that prior interactions influenced results. We informed our participants about this reset procedure. Before each condition, the participants had to answer questions (see Table 1) regarding their perceived privacy within that specific condition. After each condition, participants had to fill in their perceived privacy protection again. Additionally, if participants had previously mentioned having experience with privacy-protecting browser extensions, we asked them about the differences between their extension and the one provided within the study. Overall, the study lasted one hour. After completing all tasks, the participants were informed about the study's true intention and received the option to retract their data from the study without losing their compensation. We conducted a pilot study with three participants (two female and one male, aged 21 – 23, all students) to test the study design and procedure. Based on their input, we made minor adjustments, such as providing a mouse alongside the touchpad and refining questions for clarity. We obtained ethical clearance for the study from our institutional review board.

3.3 Task

We instructed participants to select and book a hotel that fulfilled certain criteria. We chose a hotel booking task as it represents a realistic and familiar scenario for most users. Further, as the task revolves around selecting and completing a booking, it is inherently associated with explicitly providing personal data and collecting user data to enhance personalization. Hotel booking also includes multiple implicit tracking activities, such as accepting cookies or using location-based information. Previous work showed that third-party hotel booking websites collect and share personal data without explicitly telling users the implications of the data collection procedure, including data dissemination practices [40]. We decided to use a locally hosted website to have control over this environmental factor. This gave us the advantage of replicating the browsing experience exactly during each task. Additionally, this meant that no third party could collect user data. We ultimately decided to implement a website that closely resembled the popular third-party travel page booking.com[4] but with stripped functionality. We only implemented the functionality required for the study, such as integrating locally hosted tracking through Google Analytics to simulate tracking activity and enabling the booking process for a pre-selected set of hotels and cities. We used Google Analytics with a placeholder tracking ID without connecting to a Google Analytics account, meaning no data was collected through the service.

3.4 Apparatus

The setup consisted of a monitor, keyboard, and mouse. We used a browser to display the content of the booking app. For this study, the browser had to be perceived as neutral as possible and be capable of displaying both the questionnaire and the study website. Thus, we used Chromium in its pure form as its interface resembles Chrome, which was already well known by most users and, thus, reduced the impact on their perceived privacy. Additionally, Chromium came without Google services, allowing users to browse without potentially negative perceptions, which might have occurred using Chrome or Edge. Chromium's functional deficits, such as the inability to play YouTube videos, did not affect the study as we did not use or need these features. The plugin of the Placebic Protection Narrative simulated a randomized number of blocked connections between one and 15, similar to the plugin of the Functional Protection Narrative. Participants could verify the functionality by clicking on the extension's icon, where they could view the blocked connections of the Google Analytics service. We modified the openly available source code of uBlock Origin, replacing its icon with that of the placebo condition and removing its name to eliminate potential recognition bias.

3.5 Independent Variables

We use the Protection Narrative as the only independent variable with the three levels *No Protection Narrative*, *Functional Protection Narrative*, and *Placebic Protection Narrative* (see Figure 2).

3.5.1 No Protection Narrative. Participants were informed that no tracking protection extension was enabled. The browser did not display any protection mechanisms.

3.5.2 Functional Protection Narrative. The participants received a narrative that they used an established tracking protection extension. The plugin was functional and blocked connections from the Google Analytics mock-up service. We used the openly available code of uBlock Origin but modified it to remove the name and the icon to prevent recognition bias.

3.5.3 Placebic Protection Narrative. The placebic extension purported to have privacy protection functionality. We built the extension using the latest API for browser extension development, Manifest V3 [25]. We only used Chrome's standard libraries. To fulfill the study purpose, the extension simulated a functioning extension by displaying a number that visually replicated the behavior of the blocked items of the functional uBlock Origin extension by randomly displaying blocked connections between one and 15 per

[4] https://www.booking.com

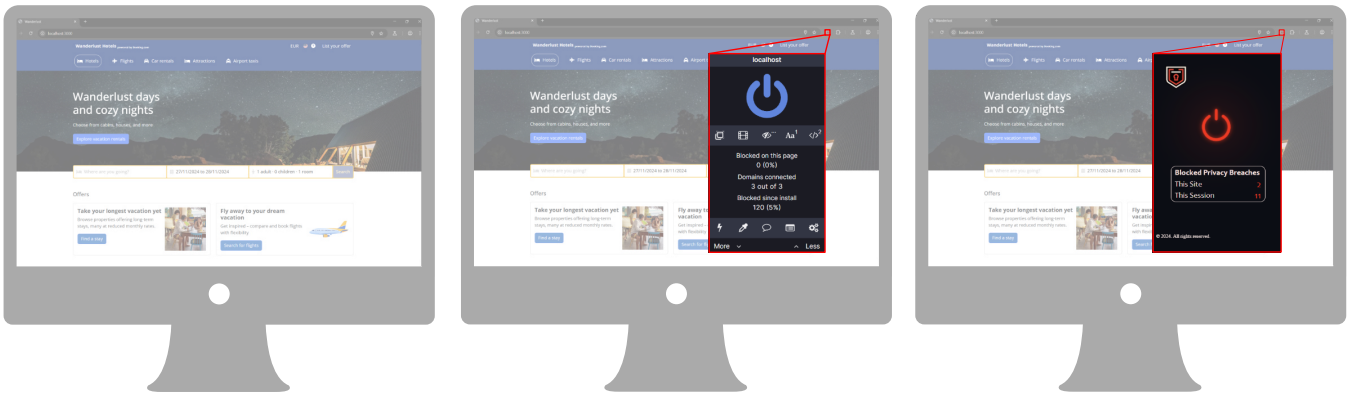

Figure 2: The three conditions employed in the study. The user interacted with the fictional booking website with (1) *No Protection Narrative* (no extension), (2) *Functional Protection Narrative* (modified uBlock Origin), and (3) *Placebic Protection Narrative* (our placebo extension). The participants could interact with the extensions during the booking process by clicking on the extension in the top right corner. Both extensions dynamically displayed blocked connections.

page. The popup design of the extension was inspired by a combination of AdBlock Plus[5] and uBlock Origin. We adopted the concept of the large power button from uBlock Origin, while AdBlock Plus inspired the statistics table. In conjunction with other extensions, we presented the on/off states with either a colored or a grayscale badge.

3.6 Dependent Variables

We measured the IUIPC [19] of each participant before the start of the study to assess the user's general attitude and concerns towards privacy. After we explained each protection plugin and told participants which extension they were using for the upcoming condition, we asked them about their confidence regarding privacy protection using 100-point sliders. We used visual analog scales (VAS) without ticks to prevent the responses from converging around the ticks (cf. [22]). Moreover, VAS has been shown to lead to more precise responses and, thus, higher data quality [8]. Finally, VAS collects continuous data, which allows for more statistical tests [29]. We measure the participant's confidence in the web extensions' protection capability after each condition again (Q4 – Q6, see Table 1).

4 Results

We used Python and R to analyze our data. To compare user expectations before and after each interaction with the plugins, we used a Friedman test to assess the main effects of the within-subject factors (i.e., No Protection Narrative, Functional Protection Narrative, Placebic Protection Narrative) with the effect size Kendall's W[6]. In case of a significant effect, we conducted post hoc analyses using Wilcoxon signed-rank tests with Bonferroni correction to test for statistical significance between the groups. The significance level was set at $\alpha = .05$. We report the p-values for each pairwise comparison if we find a significant main effect, along with the test statistics V[7]. Furthermore, we report the effect size r for each pairwise comparison[8]. Figure 3 and Figure 4 summarize the results in violin plots.

4.1 Prior Expectations

For Q1, the results indicated a significant main effect between the groups, $\chi^2(2, N = 36) = 42.99$, $p < .001$, $W = 0.60$. A post hoc test indicated a significant difference between No Protection Narrative and Functional Protection Narrative, $V = 14.0$, $p < .001$, $r = 0.84$, as well as between No Protection Narrative and Placebic Protection Narrative, $V = 37.0$, $p < .001$, $r = 0.78$. However, we did not find a significant effect between Functional Protection Narrative and Placebic Protection Narrative, $V = 176.0$, $p = .11$, $r = 0.41$. For Q2, the results indicated a significant main effect between the groups, $\chi^2(2, N = 36) = 23.62$, $p < .001$, $W = 0.33$. A post hoc test indicated a significant difference between No Protection Narrative and Functional Protection Narrative, $V = 564.0$, $p < .001$, $r = 0.60$, as well as between No Protection Narrative and Placebic Protection Narrative, $V = 568.0$, $p < .001$, $r = 0.62$. However, we did not find a significant effect between Functional Protection Narrative and Placebic Protection Narrative, $V = 390.0$, $p = .66$, $r = 0.15$. For Q3, the results indicated a significant main effect between the groups, $\chi^2(2, N = 36) = 26.63$, $p < .001$, $W = 0.37$. A post hoc test indicated a significant difference between No Protection Narrative and Functional Protection Narrative, $V = 47.5$, $p < .001$, $r = 0.75$, as well as between No Protection Narrative and Placebic Protection Narrative, $V = 44.0$, $p < .001$, $r = 0.76$, However, we did not find a significant effect between Functional Protection Narrative and Placebic Protection Narrative, $V = 238$, $p = .63$, $r = 0.25$.

[5] https://new.adblockplus.org

[6] $W < 0.1$: Negligible effect. $W < 0.3$: Weak effect. $W < 0.5$: Moderate effect. $W \geq 0.5$: Strong effect.

[7] The Wilcoxon signed-rank test was conducted using R (version 2024.09.1+394) with the test statistic being labeled as V in R, which is equivalent to W in this context.

[8] $r < 0.1$: Negligible effect. $r < 0.3$: Small effect. $r < 0.5$: Medium effect. $r \geq 0.5$: Large effect.

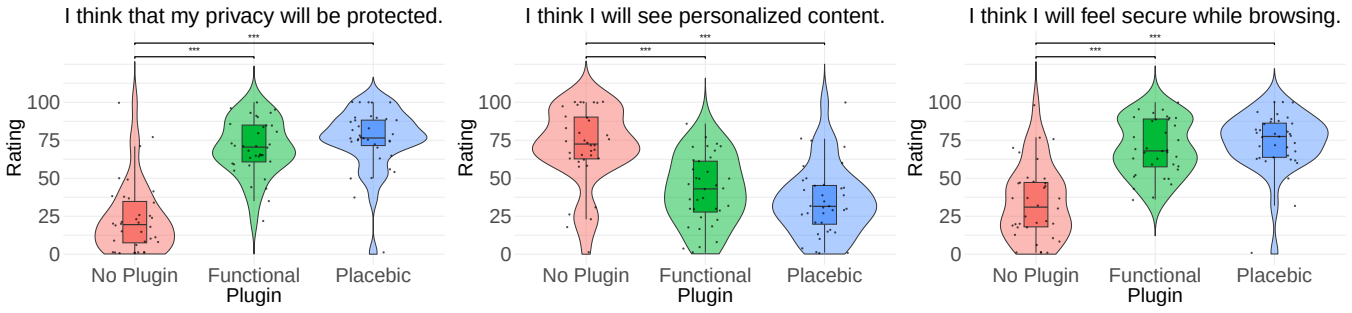

Figure 3: Violin plots of user expectations regarding the present privacy protection between the conditions No Protection Narrative, Functional Protection Narrative, and Placebic Protection Narrative before interaction. Asterisks denote significant differences.

4.2 Post-Expectations

For Q4, the results indicated a significant main effect between the groups, $\chi^2(2, N = 36) = 46.63$, $p < .001$, $W = 0.65$. A post hoc test indicated a significant difference between No Protection Narrative and Functional Protection Narrative, $V = 25.0$, $p < .001$, $r = 0.80$ between No Protection Narrative and Placebic Protection Narrative, $V = 20$, $p = .001$, $r = 0.82$, as well as between Functional Protection Narrative and Placebic Protection Narrative, $V = 118.0.0$, $p < .019$, $r = 0.56$. For Q5, the results indicated a significant main effect between the groups, $\chi^2(2, N = 36) = 8.80$, $p = .012$, $W = 0.12$. A post hoc test indicated a significant difference between No Protection Narrative and Placebic Protection Narrative, $V = 460.0$, $p = .016$, $r = 0.33$. However, we did not find a significant effect between No Protection Narrative and Functional Protection Narrative, $V = 411.0$, $p < .351$, $r = 0.20$ as well as Functional Protection Narrative and Placebic Protection Narrative, $V = 368.0$, $p < .354$, $r = 0.09$.

For Q6, the results indicated a significant main effect between the groups, $\chi^2(2, N = 36) = 23.82$, $p < .001$, $W = 0.33$. A post hoc test indicated a significant difference between No Protection Narrative and Functional Protection Narrative, $V = 76.0$, $p = .001$, $r = 0.67$, as well as between No Protection Narrative and Placebic Protection Narrative, $V = 74.5$, $p < .001$, $r = 0.68$. However, we did not find a significant effect between Functional Protection Narrative and Placebic Protection Narrative, $V = 214$, $p = .459$, $r = 0.31$.

5 Discussion

Our participants conducted three booking iterations using a tracking protection extension with No Protection Narrative, a Functional Protection Narrative, and a Placebic Protection Narrative. Our results show that users felt significantly better protected pre and post-interaction when using an extension with Functional Protection Narrative or Placebic Protection Narrative compared to No Protection Narrative. This implies that users who rely on the provided protection narrative are unable to assess the actual privacy protection. We discuss the implications of our results below.

5.1 The Presence of Tracking Protection Manipulates the Perceived Sense of Privacy Protection

We found that participants expect to be protected whenever a privacy protection extension is available. Interestingly, there were only significant differences between the extension with No Protection Narrative and Functional Protection Narrative as well as No Protection Narrative and Placebic Protection Narrative, indicating that an arbitrary presence of protections is enough for participants to believe that their privacy is protected. Thus, we **confirm H1**, indicating that users perceived higher privacy protection when using both functional and placebic tracking protection extensions compared to the no protection. Interestingly, participants also reported feeling more secure before and after the interaction, **confirming H3**, which posits that users feel more secure when using these extensions compared to the no-protection narrative. Furthermore, the participants expected to see less personalized content before interacting when using a functional or placebic protection narrative compared to a non-functional narrative. However, a significant effect between the no protection and placebic narrative remains after the interaction, with an overall decrease in the mean ratings. We expected this since the website did not adapt content, regardless of the employed privacy protection extension. Consequently, we **partially accept H2**, which hypothesized that users would perceive less personalized content when using functional and placebic tracking protection extensions.

5.2 Users Cannot Quantify Tracking and Privacy Protection

Our results echo the findings from Schaub et al. [30], where participants stated that they lean towards using web extensions but are unsure how they work and what they protect. We found that participants rarely used integrated visualizations to understand exactly what the extension blocked or protected. Yet, prior research has shown that privacy protection extensions enable the fingerprinting of users who disclose their identity and behavior [32] and that web extensions themselves can lead to security breaches that hurt user privacy rather than protecting it [26, 28]. Furthermore, participants

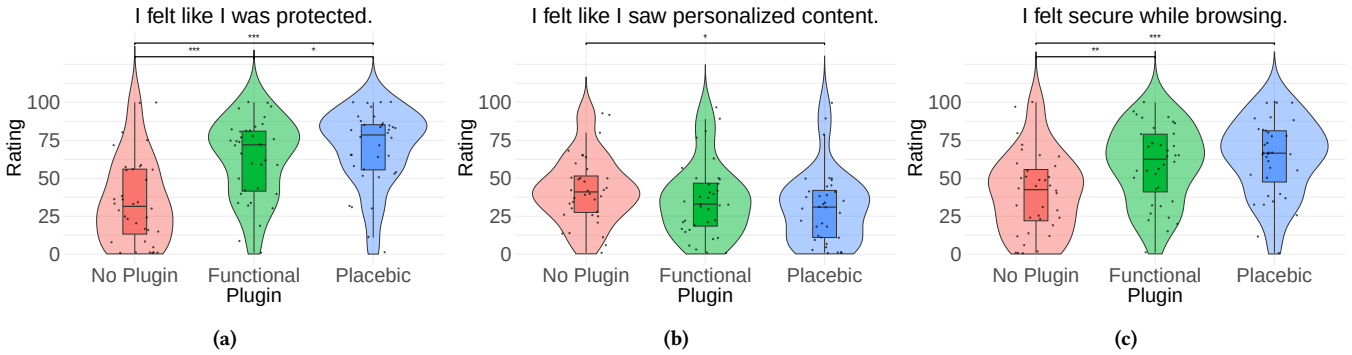

Figure 4: Violin plots of user expectations regarding the present privacy protection between the conditions No Protection Narrative, Functional Protection Narrative, and Placebic Protection Narrative after interaction. Asterisks denote significant differences.

disclosed that they did not see any changes to the website when using a privacy protection extension. We expected this, since we did not change any website content regardless of the used privacy protection extension. However, participants still rated their perceived protection and security higher before and after conducting the booking process. Hence, **future designs of privacy protection extensions should inform users about their exact intentions and highlight which privacy-related threats they diminish.** Furthermore, we **strongly recommend incorporating placebo conditions in studies assessing privacy-preserving systems and interfaces.** This ensures that **any observed improvement in privacy perception is not solely attributed to the study's narrative, thereby avoiding research outcomes influenced by participant misconceptions.** We state our recommendations below.

5.3 Studies Investigating Privacy Tools Must Include Placebo Conditions

Our results hold significance for research investigating new privacy protection methods using subjective user evaluations (cf. [2, 36]). We suggest that outcomes in prior privacy studies might stem from users' high expectations of the privacy protection method, even though the extension's functionality was assessed based on expectations rather than actual performance. Such placebo effects compromise privacy studies that evaluate new privacy protection mechanisms through subjective user feedback [16, 17]. We propose guidelines to control for placebo effects in subsequent research.

5.3.1 Measuring User Expectations Before and After Study Conditions. Controlling for placebo effects in privacy research is challenging. In medicine, the placebo effect can be so significant that many pharmacological trials control for this effect by comparing the effects of a new treatment (e.g., medicine or vaccine) with a control group that only receives a placebo. The researchers ensure that participants are unaware of whether they are in the experimental or control group so they can manage their expectations before the treatment. In contrast, user studies regarding novel privacy protection techniques face greater challenges with placebo control compared to pharmacological trials [9]. Participants in user studies can often tell if they are using a novel interface by recognizing its differences from a familiar one, leading them to assume the new interface might be better. **To explore the influence of expectancy on participants' subjective perceptions and objective performance, we recommend measuring the expectation regarding the level of protection before and after interaction with a privacy protection system.** Yet, a majority of prior literature in human-computer interaction neglected to control for user expectations in the past. **Asking for user expectations before and after interaction regarding the functionality of a privacy protection system provides insights into potential biases of subjective user measurements.** This ensures that the findings are not the result of pure user expectations.

5.3.2 Including Placebo-Control Conditions. Similar to medicine, we recommend including a placebo condition when conducting studies focusing on privacy protection. By incorporating an active control group or even a placebo control into the study design, researchers can assess if the new privacy system's usability surpasses user expectations. There is the possibility that participants were affected by a placebo when participants showed similar user expectations for the active and placebo-controlled condition. If the user expectations persist after the interaction, chances are high that users were primed through the study narrative. Nonetheless, using a placebo control may be inefficient, unethical, or impractical, especially when a placebo is easily recognizable or when dealing with vulnerable patients who need active treatment. In these situations, **researchers can statistically adjust for the effect of expectations, as demonstrated in this study. This involves measuring expectations before testing and then normalizing the bias during analysis.**

5.4 Limitations and Future Work

Although the framing of the extensions might have skewed privacy and security expectations more in favor of the Placebic Protection Narrative, we also observe a potential shortcoming in the extension design. By definition, a placebo should appear functional despite lacking true functionality. While our extension largely adhered to this principle, we included visual feedback with random

numbers and a fake blocked items statistic (see Figure 2). Despite both extensions showing the same number of blocked items during the search task, the synchronization issue when switching back to the survey page was caused by the Placebic Protection Narrative displaying a higher count on the questionnaire page than uBlock Origin. Future research should aim to align more closely with the definition of a placebo system as outlined by Kosch et al. [17]. An excellent example is the placebo extension used by Marella et al. [20].

Our participants were challenged to identify the differences between the Functional Protection Narrative and Placebic Protection Narrative. Although such an outcome is possible even under ideal conditions, our task design revealed two potential limitations that may have influenced these results. Firstly, the task might have been too brief to discern distinguishing features between the extensions. Additionally, participants had to complete a different search task for each condition, particularly searching for different cities, which might have hindered their ability to determine if the search results were personalized. Thus, we recommend that future studies in this area provide participants with more extensive exploration opportunities. This could affect their perception of placebo extensions. Secondly, the choice to have the experimenter, rather than the participants, activate and deactivate the extensions might have impacted participants' willingness to interact with the extensions. Moreover, participants were never explicitly encouraged to examine the extension popups. This could explain why most participants did not utilize the tools provided by the extensions, making it difficult for them to identify the placebo extension.

The majority of our participants were university students, who are generally considered to have above-average technology proficiency, potentially limiting the generalizability of our findings. However, even this tech-savvy group had difficulty distinguishing between the non-functional plugin and the working one. This indicates that less tech-savvy populations might face even greater challenges in differentiating functional from placebic technology. Future research should, therefore, include a more diverse sample to validate our findings and ensure their relevance across a broader range of user groups.

6 Conclusion

In this work, we investigated placebo effects with privacy-protecting browser extensions. We conducted a user study that compared participants' perceived privacy protection through different narratives provided by tracking protection extensions. By priming our participants with three distinct narratives, namely No Protection Narrative, Functional Protection Narrative, and Placebic Protection Narrative, we found that users felt more protected and secure when listening to a Functional Protection Narrative or Placebic Protection Narrative compared to a No Protection Narrative, although all participants viewed the same content. Furthermore, participants believed to see less personalized content when being primed with a Placebic Protection Narrative compared to a No Protection Narrative. Our study results not only showed a placebo effect but also confirmed previous research on the importance of education in terms of privacy and security, where users should not trust privacy protection extensions blindly [15, 19]. We see this lack of education as an issue that should be addressed promptly. Contrary to our significant findings on the perception of privacy protection, we found that the results of the perceptions of potentially visually personalized content do not show a placebo effect. We expected this result since the website did not adapt any content, regardless of the user tracking protection extension. This hints that participants were aware that nothing was adapted. Still, participants perceived a higher level of protection. Due to a lack of research regarding the placebo effect of privacy in human-computer interaction, we see a lot of potential for future work, such as investigating the validity of previous studies that evaluate a novel privacy protection system and long-term evaluations of placebic privacy-preserving interfaces. Furthermore, the lack of research is especially problematic in privacy, as a placebo effect can hurt the trust users place in protection mechanisms since users cannot assess the functionality. This could lead to users deliberately not using privacy-protecting applications or overestimating the protection, exposing them to more threats. Hence, we strongly encourage future privacy research to incorporate a placebo condition in their experiments.

Acknowledgments

This work is supported by the German Research Foundation (DFG), CRC 1404: "FONDA: Foundations of Workflows for Large-Scale Scientific Data Analysis" (Project-ID 414984028).

References

[1] Henry K. Beecher. 1955. The Powerful Placebo. *Journal of the American Medical Association* 159, 17 (Dec. 1955), 1602–1606. doi:10.1001/jama.1955.02960340022006

[2] Kaustav Bhattacharjee, Min Chen, and Aritra Dasgupta. 2020. Privacy-Preserving Data Visualization: Reflections on the State of the Art and Research Opportunities. *Computer Graphics Forum* 39, 3 (2020), 675–692. doi:10.1111/cgf.14032 _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.14032.

[3] Walter R. Boot, Daniel J. Simons, Cary Stothart, and Cassie Stutts. 2013. The Pervasive Problem With Placebos in Psychology: Why Active Control Groups Are Not Sufficient to Rule Out Placebo Effects. *Perspectives on Psychological Science* 8, 4 (July 2013), 445–454. doi:10.1177/1745691613491271 Publisher: SAGE Publications Inc.

[4] Esther Bosch, Robin Welsch, Tamim Ayach, Christopher Katins, and Thomas Kosch. 2024. The Illusion of Performance: The Effect of Phantom Display Refresh Rates on User Expectations and Reaction Times. In *Extended Abstracts of the 2024 CHI Conference on Human Factors in Computing Systems (CHI EA '24)*. Association for Computing Machinery, New York, NY, USA, 1–6. doi:10.1145/3613905.3650875

[5] Alena Denisova and Paul Cairns. 2015. The Placebo Effect in Digital Games: Phantom Perception of Adaptive Artificial Intelligence. In *Proceedings of the 2015 Annual Symposium on Computer-Human Interaction in Play*. ACM, London United Kingdom, 23–33. doi:10.1145/2793107.2793109

[6] Alena Denisova and Paul Cairns. 2019. Player experience and deceptive expectations of difficulty adaptation in digital games. *Entertainment Computing* 29 (March 2019), 56–68. doi:10.1016/j.entcom.2018.12.001

[7] Damien G. Finniss, Ted J. Kaptchuk, Franklin Miller, and Fabrizio Benedetti. 2010. Biological, clinical, and ethical advances of placebo effects. *The Lancet* 375, 9715 (Feb. 2010), 686–695. doi:10.1016/S0140-6736(09)61706-2 Publisher: Elsevier.

[8] Frederik Funke and Ulf-Dietrich Reips. 2012. Why Semantic Differentials in Web-Based Research Should Be Made from Visual Analogue Scales and Not from 5-Point Scales. *Field Methods* 24, 3 (2012). doi:10.1177/1525822X12444061

[9] Andrew L. Geers, Paul E. Weiland, Kristin Kosbab, Sarah J. Landry, and Suzanne G. Helfer. 2005. Goal Activation, Expectations, and the Placebo Effect. *Journal of Personality and Social Psychology* 89, 2 (Aug. 2005), 143–159. doi:10.1037/0022-3514.89.2.143

[10] Arthur Gervais, Alexandros Filios, Vincent Lenders, and Srdjan Capkun. 2017. Quantifying Web Adblocker Privacy. In *Computer Security – ESORICS 2017*, Simon N. Foley, Dieter Gollmann, and Einar Snekkenes (Eds.). Springer International Publishing, Cham, 21–42. doi:10.1007/978-3-319-66399-9_2

[11] Gabor Gyorgy Gulyas, Doliere Francis Some, Nataliia Bielova, and Claude Castelluccia. 2018. To Extend or not to Extend: On the Uniqueness of Browser Extensions and Web Logins. In *Proceedings of the 2018 Workshop on Privacy in the Electronic Society*. ACM, Toronto Canada, 14–27. doi:10.1145/3267323.3268959

[12] Astrid Hernández, Josep-E. Baños, Cristina Llop, and Magí Farré. 2014. The Definition of Placebo in the Informed Consent Forms of Clinical Trials. *PLOS ONE* 9, 11 (Nov. 2014), e113654. doi:10.1371/journal.pone.0113654 Publisher: Public Library of Science.

[13] Roberto Hoyle, Luke Stark, Qatrunnada Ismail, David Crandall, Apu Kapadia, and Denise Anthony. 2020. Privacy Norms and Preferences for Photos Posted Online. *ACM Trans. Comput.-Hum. Interact.* 27, 4, Article 30 (2020), 27 pages. doi:10.1145/3380960

[14] Ankit Kariryaa, Gian-Luca Savino, Carolin Stellmacher, and Johannes Schöning. 2021. Understanding Users' Knowledge about the Privacy and Security of Browser Extensions. In *Seventeenth Symposium on Usable Privacy and Security (SOUPS 2021)*. USENIX Association, 99–118. https://www.usenix.org/conference/soups2021/presentation/kariryaa

[15] Ankit Kariryaa, Gian-Luca Savino, Carolin Stellmacher, and Johannes Schöning. 2021. Understanding Users' Knowledge about the Privacy and Security of Browser Extensions. (2021).

[16] Agnes Mercedes Kloft, Robin Welsch, Thomas Kosch, and Steeven Villa. 2024. "AI enhances our performance, I have no doubt this one will do the same": The Placebo effect is robust to negative descriptions of AI. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. ACM, Honolulu HI USA, 1–24. doi:10.1145/3613904.3642633

[17] Thomas Kosch, Robin Welsch, Lewis Chuang, and Albrecht Schmidt. 2022. The Placebo Effect of Artificial Intelligence in Human–Computer Interaction. *ACM Transactions on Computer-Human Interaction* 29, 6 (Dec. 2022), 1–32. doi:10.1145/3529225

[18] Louis Lasagna, Frederick Mosteller, John M. von Felsinger, and Henry K. Beecher. 1954. A study of the placebo response. *The American Journal of Medicine* 16, 6 (June 1954), 770–779. doi:10.1016/0002-9343(54)90441-6 Publisher: Elsevier.

[19] Naresh K. Malhotra, 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 (2004). doi:10.1287/isre.1040.0032

[20] Aditya Marella, Chao Pan, Ziwei Hu, Florian Schaub, Blase Ur, and Lorrie Faith Cranor. 2014. Assessing privacy awareness from browser plugins. In *Proceedings of the Symposium on Usable Privacy and Security (SOUPS)*.

[21] David M. Martin, Richard M. Smith, Michael Brittain, Ivan Fetch, and Hailin Wu. 2001. The privacy practices of Web browser extensions. *Commun. ACM* 44, 2 (Feb. 2001), 45–50. doi:10.1145/359205.359226

[22] Justin Matejka, Michael Glueck, Tovi Grossman, and George Fitzmaurice. 2016. The Effect of Visual Appearance on the Performance of Continuous Sliders and Visual Analogue Scales. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (San Jose, California, USA) *(CHI '16)*. Association for Computing Machinery, New York, NY, USA, 5421–5432. doi:10.1145/2858036.2858063

[23] Maryam Mehrnezhad, Kovila Coopamootoo, and Ehsan Toreini. 2021. How Can and Would People Protect From Online Tracking? *Proceedings on Privacy Enhancing Technologies* 1 (Nov. 2021). doi:10.2478/popets-2022-0006 Publisher: De Gruyter Open.

[24] Guy Montgomery and Irving Kirsch. 1996. Mechanisms of Placebo Pain Reduction: An Empirical Investigation. *Psychological Science* 7, 3 (May 1996), 174–176. doi:10.1111/j.1467-9280.1996.tb00352.x Publisher: SAGE Publications Inc.

[25] Nikolaos Pantelaios and Alexandros Kapravelos. 2024. Manifest V3 Unveiled: Navigating the New Era of Browser Extensions. doi:10.14722/madweb.2024.23080 arXiv:2404.08310 [cs].

[26] Emmanouil Papadogiannakis, Panagiotis Papadopoulos, Nicolas Kourtellis, and Evangelos P. Markatos. 2021. User Tracking in the Post-cookie Era: How Websites Bypass GDPR Consent to Track Users. In *Proceedings of the Web Conference 2021*. ACM, Ljubljana Slovenia, 2130–2141. doi:10.1145/3442381.3450056

[27] Pat Pataranutaporn, Ruby Liu, Ed Finn, and Pattie Maes. 2023. Influencing human–AI interaction by priming beliefs about AI can increase perceived trustworthiness, empathy and effectiveness. *Nature Machine Intelligence* 5, 10 (Oct. 2023), 1076–1086. doi:10.1038/s42256-023-00720-7 Publisher: Nature Publishing Group.

[28] Sylvia E Peacock. 2014. How web tracking changes user agency in the age of Big Data: The used user. *Big Data & Society* 1, 2 (July 2014), 2053951714564228. doi:10.1177/2053951714564228 Publisher: SAGE Publications Ltd.

[29] Ulf-Dietrich Reips and Frederik Funke. 2008. Interval-level measurement with visual analogue scales in Internet-based research: VAS Generator. *Behavior Research Methods* 40, 3 (2008). doi:10.3758/BRM.40.3.699

[30] Florian Schaub, Aditya Marella, Pranshu Kalvani, Blase Ur, Chao Pan, Emily Forney, and Lorrie Faith Cranor. 2016. Watching Them Watching Me: Browser Extensions Impact on User Privacy Awareness and Concern. In *Proceedings 2016 Workshop on Usable Security*. Internet Society, San Diego, CA. doi:10.14722/usec.2016.23017

[31] Katta Spiel, Sven Bertel, and Fares Kayali. 2019. Adapting Gameplay to Eye Movements - An Exploration with TETRIS. In *Extended Abstracts of the Annual Symposium on Computer-Human Interaction in Play Companion Extended Abstracts*. ACM, Barcelona Spain, 687–695. doi:10.1145/3341215.3356267

[32] Oleksii Starov and Nick Nikiforakis. 2017. Extended Tracking Powers: Measuring the Privacy Diffusion Enabled by Browser Extensions. In *Proceedings of the 26th International Conference on World Wide Web*. International World Wide Web Conferences Steering Committee, Perth Australia, 1481–1490. doi:10.1145/3038912.3052596

[33] Oleksii Starov and Nick Nikiforakis. 2018. PrivacyMeter: Designing and Developing a Privacy-Preserving Browser Extension. In *Engineering Secure Software and Systems*, Mathias Payer, Awais Rashid, and Jose M. Such (Eds.). Vol. 10953. Springer International Publishing, Cham, 77–95. doi:10.1007/978-3-319-94496-8_6 Series Title: Lecture Notes in Computer Science.

[34] Steve Stewart-Williams and John Podd. 2004. The Placebo Effect: Dissolving the Expectancy Versus Conditioning Debate. *Psychological Bulletin* 130, 2 (2004), 324–340. doi:10.1037/0033-2909.130.2.324 Place: US Publisher: American Psychological Association.

[35] Christopher A. Summers, Robert W. Smith, and Rebecca Walker Reczek. 2016. An Audience of One: Behaviorally Targeted Ads as Implied Social Labels. *Journal of Consumer Research* 43, 1 (June 2016), 156–178. doi:10.1093/jcr/ucw012

[36] Yuuki Takano, Ohta Satoshi, Takeshi Takahashi, Ruo Ando, and Tomoya Inoue. 2014. *MindYourPrivacy: Design and implementation of a visualization system for third-party Web tracking*. doi:10.1109/PST.2014.6890923 Journal Abbreviation: 2014 12th Annual Conference on Privacy, Security and Trust, PST 2014 Pages: 56 Publication Title: 2014 12th Annual Conference on Privacy, Security and Trust, PST 2014.

[37] Kristen Vaccaro, Dylan Huang, Motahhare Eslami, Christian Sandvig, Kevin Hamilton, and Karrie Karahalios. 2018. The Illusion of Control: Placebo Effects of Control Settings. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM, Montreal QC Canada, 1–13. doi:10.1145/3173574.3173590

[38] Steeven Villa, Thomas Kosch, Felix Grelka, Albrecht Schmidt, and Robin Welsch. 2023. The placebo effect of human augmentation: Anticipating cognitive augmentation increases risk-taking behavior. *Computers in Human Behavior* 146 (Sept. 2023), 107787. doi:10.1016/j.chb.2023.107787

[39] John D. Wells, Damon E. Campbell, Joseph S. Valacich, and Mauricio Featherman. 2010. The Effect of Perceived Novelty on the Adoption of Information Technology Innovations: A Risk/Reward Perspective. *Decision Sciences* 41, 4 (2010), 813–843. doi:10.1111/j.1540-5915.2010.00292.x _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1540-5915.2010.00292.x.

[40] Haiyue Yuan, Matthew Boakes, Xiao Ma, Dongmei Cao, and Shujun Li. 2023. Visualising Personal Data Flows: Insights from a Case Study of Booking.com. In *Intelligent Information Systems*, Cristina Cabanillas and Francisca Pérez (Eds.). Springer International Publishing, Cham, 52–60.