

A Large-Scale Study of Cookie Banner Interaction Tools and Their Impact on Users' Privacy

Nurullah Demir

Institute for Internet Security
KASTEL Security Research Labs
Karlsruhe Institute of Technology

Tobias Urban

Institute for Internet Security
secunet Security Networks AG

Norbert Pohlmann

Institute for Internet Security

Christian Wressnegger

KASTEL Security Research Labs
Karlsruhe Institute of Technology

ABSTRACT

Cookie notices (or cookie banners) are a popular mechanism for websites to provide (European) Internet users a tool to choose which cookies the site may set. Banner implementations range from merely providing information that a site uses cookies over offering the choice to accepting or denying all cookies to allowing fine-grained control of cookie usage. Users frequently get annoyed by the banner's pervasiveness as they interrupt "natural" browsing on the Web. As a remedy, different browser extensions have been developed to automate the interaction with cookie banners.

In this work, we perform a large-scale measurement study comparing the effectiveness of extensions for "cookie banner interaction." We configured the extensions to express different privacy choices (e.g., accepting all cookies, accepting functional cookies, or rejecting all cookies) to understand their capabilities to execute a user's preferences. The results show statistically significant differences in which cookies are set, how many of them are set, and which types are set—even for extensions that aim to implement the same cookie choice. Extensions for "cookie banner interaction" can effectively reduce the number of set cookies compared to no interaction with the banners. However, all extensions increase the tracking requests significantly except when rejecting all cookies.

KEYWORDS

cookie banner, consent banner, cookies, privacy, web measurement, web privacy, robustness

1 INTRODUCTION

Websites make rich use of HTTP cookies for various means (e.g., user tracking). To provide (European) Internet users more control over the usage of cookies, many website providers embed so-called "cookie banners" on their webpages [8, 13, 17]. Some banners allow fine-grained control over the type of cookies to be used (e.g., rejecting advertising cookies), while others only inform users about their usage [44]. Cookie banners are meant to help users but are often designed in a way that nudges the user to accept all types of

cookies rather than to reject them ("Dark Patterns" [15, 27]), and their the omnipresence has started to annoy users [29].

Different tools help users cope with cookie banners by automatizing the interaction process [4, 24, 33, 34, 39]. These tools are usually implemented as a browser extension and often use rule-based approaches to identify and interact with banners. More specifically, the tools identify the banners and corresponding buttons on predefined patterns, similar to ad blockers that block URLs based on filter lists. Some tools offer the option to choose the cookie type the user consents to be set (e.g., "functional cookies" only). Users hence need to configure these tools according to their own privacy needs. However, the impact of these tools on the user's privacy still needs to be better understood. One challenge is that the interaction with these banners is neither standardized nor is it (legally) defined what the purpose of a cookie is or how it can be determined. Thus, while users rely on these tools for convenience and to implement their choice, it is still being determined if these tools even meet these expectations and which impact the tools have on users' privacy.

In this work, we perform a measurement study to understand the effects of six extensions for "cookie banner interaction." We analyze five tools used in the field and one custom extension developed for this study. In our experiments, we investigate (1) which and how many cookies a website sets, (2) the purpose of the used cookies, and (3) various deferred effects, such as the impact on tracking requests. With this study, we aim to understand the impact and potential benefits of the different tools for users' privacy. To do so, we visit 298k distinct pages on 30k websites once with each tool and uncover statistically significant differences in the analyzed extensions' effectiveness. We show that the number of cookies, category of used cookies, and individual cookies in terms of their key names differ based on the used extension.

Previous studies have either analyzed the design of different banners [8, 15, 41, 44] or have built new tools to interact with banners in a meaningful way [4, 34]. While we build upon prior work by utilizing some of the presented tools in this work, we investigate an entirely different problem: The effectiveness of different tools that automatically interact with cookie banners. To the best of our knowledge, this is the first work comparing the impact of "cookie banner interaction" tools on the user's privacy at a large scale.

In summary, the main contributions of this work are:

- **Large-scale measurement study.** We compare the effectiveness of six different tools that assist users in the interaction with

This work is licensed under the Creative Commons Attribution 4.0 International License. To view a copy of this license visit <https://creativecommons.org/licenses/by/4.0/> or send a letter to Creative Commons, PO Box 1866, Mountain View, CA 94042, USA.



Proceedings on Privacy Enhancing Technologies 2024(1), 5–20
© 2024 Copyright held by the owner/author(s).
<https://doi.org/10.56553/popets-2024-0002>

cookie banners and show that the tools can impact the usage of the setting of specific types of cookies by up to 640%.

- **Tool effectiveness varies largely.** The results show that although some tools claim to have similar capabilities (e.g., accepting functional cookies), several differences exist in the number and type of observed cookies.
- **Effects of interacting with cookie banners.** We find that statistically significant effects of the interaction with cookie banners exist. For example, the number of tracking requests increases by up to 60% if an extension allows the usage of all cookies.

The remainder of this paper is structured as follows: First, we introduce the terms used throughout this paper and provide background information (Section 2). In Section 3, we detail related research and describe how our paper differs from previous works. Section 4 describes the “cookie banner interaction” extensions used in our work (Section 4.2), the method to measure their effects (Section 4.4), and our approach to classify cookies (Section 4.5). Next, we present the results of our study by discussing the effects of the different tools (Section 5.2), by analyzing the types of used cookies (Section 5.3), and by looking at further dimensions like cookie usage patterns (Sections 5.4 to 5.7). Finally, we describe our work’s limitations and provide recommendations (Sections 6 and 7).

2 TERMINOLOGY AND BACKGROUND

In this section, we elaborate on common terminology used in the remainder of the paper.

Sites and Pages. In this work, we use the term *site* to depict the registerable part of a given domain—often referred to as “extended Top Level Domain plus one” (eTLD+1) [6, 10, 23, 42]. For example, given the URL `https://www.bar.com/` the eTLD+1 is `bar.com`, or the URL `https://foo.co.uk` the eTLD+1 is `foo.co.uk`. By *page* (or webpage), we mean a unique URL or, more specifically, the document (e.g., HTML or JavaScript) located at the particular URL.

Cookies. A *cookie* is a piece of text communicated via an HTTP header or set via JavaScript. More specifically, a cookie is a key-value pair set on a client by a visited website or third-party present on that website. Cookies were implemented to allow stateful communication in the otherwise stateless HTTP protocol [3]. Therefore, they are often used to manage sessions, store persistent client-side data, and often for user tracking purposes [42]. Various organizations proposed categories of cookies to account for the wide range of use cases and to bring some transparency to the usage of cookies (e.g., the *IAB Europe* proposed 17 categories of cookies [19]).

Cookie Banners. Websites often display so-called “cookie banners” that inform users about the usage of cookies and sometimes provide them some level of control over the type and amount of cookies a website may set. In the field, three types of such banners have been established [8, 44]: (1) banners that only inform users that a site uses cookies without control, (2) banners that allow users to accept or reject all cookies, and (3) banners that provide fine-grain control of different types of cookies a page may set. Different services have emerged that help website providers manage the communication of consent. Such services are often referred to as “consent management platforms” (CMP) [17, 26, 47]. In a nutshell, CMPs are standard “off-the-shelf” software that aim to automate the consent management process for websites.

3 RELATED WORK

The area of cookie banners or cookie notices and their impact on users’ privacy has been widely researched in recent years. Several works look at the structure and usability of such tools and investigate the impact of various factors, such as position, choice type, and content framing, on users’ interactions with cookie consent (e.g., [8, 44]). In contrast other studies [e.g., 17] aim to understand the impact of privacy laws by mapping the formation of the consent management provider ecosystem and measuring its adoption. Studies like Fouad et al. [13] investigate the legal compliance of the consent banners, while further works study the use of dark patterns when designing them [15, 41]. Our work distinguishes from these research streams as we investigate the effect of browser extensions that interact with cookie banners automatically or, more specifically, the impact these extensions have on users’ privacy.

Related to our research are works that look at the functionality of such banners on a technical level. Matte et al. [31] analyzed legal aspects in the storage of consent through cookie banners by crawling websites, and they found potential legal violations in the storage of user consent. Sanchez-Rola et al. [38] manually analyzed the effect of cookie notices and found that such banners often did not work as intended. Different works look at compliance issues with existing consent management tools [4, 26, 36] and find that websites often do not honor the users’ choices.

Other works focus on assisting users by automatically interacting with cookie banners. Nouwens et al. [34] present *Consent-O-Matic*, a tool that can interact with different cookie banners. Most recently, Bollinger et al. [4] present a machine-learning-based approach to classify cookies and to delete cookies of specific categories that the user can select. Hu et al. [18] also introduce a machine learning-based approach to classifying the purpose of a cookie. Klein et al. [25] conduct a large-scale crawl to explore the security impact of consenting to a cookie banner and find that users who consent to Web tracking are exposed to more security-sensitive data flows and are vulnerable to more client-side cross-site scripting (XSS) vulnerabilities. Jha et al. [22] present a measurement focusing on popular websites in Europe and the US to study the impact of privacy banners on the web. The authors developed a Web crawler, which can accept privacy policies, and compared the changes in websites before and after accepting the policies.

In contrast to the named works, this work aims to understand the effect of different “cookie banner interaction” browser extensions on users’ privacy. By analyzing the functionalities of these extensions, we aim to provide insights into the effectiveness and impact of these tools in protecting users’ privacy.

4 MEASUREMENT SETUP

The measurement framework provided by Demir et al. [9] is the basis for our setup, which allows orchestrating quasi-parallel web measurements using different browser setups. The framework consists of a master machine that starts multiple virtual machines (VM), each running a separate crawler with a distinct configuration. We use the individual configurations to implement nine different crawling profiles (cf. Section 4.4) for our experiments.

The *Google Cloud Platform* hosts our VMs, and all but one use an IP address associated with a European server (Frankfurt am Main;

DEU). Each profile is based on *OpenWPM* [12] (v0.20.0), uses the *Firefox* browser (v100.0 with the default user agent¹ and a screen resolution of 1920x1080), and each profile uses a different browser extension that interacts with a cookie banner (cf. Section 4.2). We configure *OpenWPM* to log the entire HTTP(s) traffic, all cookie interactions (e.g., creation or deletion), and data stored in the local storage of the browser [32]. Since *OpenWPM* natively does not support the extraction of this storage, we implemented a custom function that extracts the values from this storage. We make our crawling setup publicly available (cf. Appendix A).

A crawl of the Web can either be *stateless* or *stateful*. Stateful crawlers keep the state of the browser between page visits (like the browser of a real user would), while in stateless crawls, the browser is reset between two page visits. In our experiment, we perform a stateful crawl for each visited site, meaning we keep the browser's state when visiting a site's pages. The state is reset when a new site is analyzed. This design choice has major implications for our experiment because after interaction with the cookie banner on the first page (typically the landing page of a site), the banner might not show on other pages and, therefore, we can record the behavior of a site after the user made a choice. However, the order of visiting sites might impact the resulting cookie jar.

To allow fair comparability of the different profiles, we only include sites in our analysis if at least eight profiles successfully crawled all site pages in our dataset. More specifically, if at least two profiles did not crawl a given page, the entire site is dropped from our analysis. This filtering applies to 10,888 (36%) of the sites in our dataset that we excluded from our analysis. Since we visit the sites in parallel, this approach allows us to keep the data we compare consistent. More precisely, we only compare sites if the underlying data was collected from the same pages. We elaborate on this seemingly high exclusion rate in Appendix C.

4.1 Dataset

To build the set of pages to visit, we resort to the widely used quasi-standard Tranco list [30] generated on 29/08/2022². From the list, we randomly sampled a subset of sites to analyze based on their ranking. We divided the list into buckets and randomly drew sites from each bucket. We used the top 5,000 sites from the list and randomly selected 5,000 sites from each of the following rank 'buckets': 5,001–10k, 10,001–50k, 50,001–250k, 250,001–500k, and 500,001–1M. Thus, in total, we used 30,000 sites for our experiment. We use this sampling to understand if there are differences in the efficiency of the analyzed tools based on a site's popularity (rank).

Previous work has shown that subpages behave differently than the respective landing page [2, 9, 42]. Therefore, to build the final set of pages to visit, we visit each randomly selected site's landing page and collect 15 subpages (i.e., first-party links on the page) for each. We repeated the process recursively if the landing page did not hold enough links. We chose to use 15 pages since previous work showed that this number is a fair trade-off between crawling time and still capturing the behavior of a site. Overall, we visit 298k pages on the 30,000 sites utilizing the nine measurement profiles (on

Table 1: Overview of the used extensions that claim to interact with cookie banners automatically. The date indicates the last update of the extension. User numbers as of 02/23.

#	Name	Date	Firefox Users	Chrome Users	Num. Configs
I	I don't care about cookies	02/23	317,000 +	900,000 +	1
II	Consent-O-Matic	02/23	21,000 +	60,000 +	7
III	Ninja Cookie	04/21	6,000 +	40,000 +	1
IV	SuperAgent	06/22	3,000 +	20,000 +	5
V	CookieBlock	08/22	2,000 +	5,000 +	4

average 288k per profile). We provide a list of all analyzed pages and sites in the supplementary material of this paper (cf. Appendix A).

4.2 Cookie Banner Interaction Tools

To understand the impact of different tools that assist users in interacting with cookie banners, we perform a measurement study and compare the effects of other tools. Table 1 lists the extensions we analyze in this study. All analyzed extensions are supported by the most common browsers (e.g., Firefox, Chrome, or Safari). We chose these extensions since they are either popular (e.g., *I don't care about cookies* and *Ninja Cookie*) or were recently proposed by academic works as solutions to interact with cookie banners (e.g., *Consent-O-Matic* [35] and *CookieBlock* [4]). In our measurement, we install each extension in a separate *OpenWPM* instance, allowing us to compare each tool's effects individually. It should be noted that there is no specification (or understanding) of how cookie banners should work and not even a (legal) definition of how to define the type or purpose of a cookie. Yet, the extensions claim that they provide privacy or let users take control of their privacy and eliminate the need to interact with cookie banners. Therefore, it is interesting to understand their impact on users' privacy.

I don't care about cookies. The extension is by far the most popular browser extension in our measurement. The authors of the extension claim that it "*removes cookie warnings from almost all websites*" [24]. The tool's main intention is to be not to protect users' privacy but to eliminate the hassle of interacting with the banner. For example, the extension hides the HTML element(s) that compose the cookie banner. If that is not possible, the extension will interact with the banner by accepting all, some, or no cookies. The extension's description does not elaborate on how it determines its action. However, manual inspection and consultation with the developer showed that the extension usually accepts all cookies.

Consent-O-Matic. This extension was developed as part of a research project [35]. *Consent-O-Matic* was developed to interact with banners provided by different Consent Management Providers (CMPs) and aims to choose the most privacy-friendly setting by default. The extension works for a fixed set of CMPs and uses rules to interact with each banner. While, by default, the extension only accepts "necessary" cookies, it can also be configured to accept other types of cookies. We use this option to compare two extension configurations (cf. Section 4.4). When conducting our experiments, the extension supports 36 CMPs, which means if a website uses a banner not provided by any of those CMPs, the extension will not do anything. Previous work has shown that only roughly 11% of all websites use a cookie banner provided by a CMP [4, 5, 45].

¹Mozilla/5.0 (X11; Linux x86_64; rv:100.0) Gecko/20100101 Firefox/100.0
²<https://tranco-list.eu/download/X568N/1000000>

Table 2: Overview of the measurement profiles. Profiles #3 and #7 are custom extensions as described in Section 4.4. The last column marks tools that can process all banner types (and do not state any limitations of banner interaction) with a ✓.

#	Extension	Version	Location	Method	Cookie Policy	All Banners
1	I don't care about cookies	3.4.2	DE	rule-based	Hides banners and accepts/rejects cookies	✓
2	Consent-O-Matic	1.0.8	DE	rule-based	Reject all cookies	Only CMPs
3	Consent-O-Matic (<i>custom</i>)	1.0.8	DE	rule-based	Allow all cookies	Only CMPs
4	Ninja Cookie	0.2.7	DE	rule-based	Reject all cookies	✓
5	CookieBlock	1.1.0	DE	ML-based	Accepts functional cookies	✓
6	SuperAgent	2.6.0	DE	rule-based	Accepts functional and performance	✓
7	"Accept all" Extension (<i>custom</i>)	—	DE	rule-based	Accepts all cookies	✓
8	None	—	DE	—	No cookie banner interaction	—
9	None	—	US	—	No cookie banner interaction	—

Ninja Cookie. By default, this extension will interact with cookie banners and only allow the setting of essential cookies [33]. No public documentation exists on how *Ninja Cookie* identifies the cookie banners. However, by manually inspecting the source code, we found that it works on a heuristic, rule-based approach to identify and interact with a banner. According to the documentation, the tool also works for major CMPs and other banners. *Ninja Cookie* is the only extension that offers a premium subscription that provides additional features such as whitelisting. In our experiment, we use the basic subscription, which is free.

CookieBlock. This extension was also developed as part of a research project [4]. Since the work was published in 2022, the extension only comes with low installation numbers. We still integrated the extension into our experiment because it is the only tool not relying on a rule-based approach but using machine learning. Furthermore, the extension does not aim to interact with cookie banners but automatically deletes cookies that the user does not desire. By default, the extension allows setting "functional" cookies. In our experiments, we use this default configuration of *CookieBlock*.

Super Agent. The last extension in our corpus is *Super Agent* [39]. Like most solutions, this extension uses a rule-based approach to identify and interact with cookie banners. The extension's default settings allow the setting of "functional" and "performance" cookies. In our experiment, we use these default settings.

4.3 Manual Extension Verification

We begin by testing the functionality of the extensions to understand whether they interact with cookie banners and thus work as intended. To this end, we randomly sample 19 sites from each of the used buckets that show a cookie banner and one that does not show a banner. In the latter, we include testing if the extensions might break a page's function. First, we manually visited each of the selected sites and tested if they showed or did not show a cookie banner using a vanilla Firefox browser and a European IP address associated with a German university (*Westphalian University of Applied Sciences*) to gather ground truth. Then, we use the same setup (vanilla Firefox browser and a German IP address) and (1) manually install each extension separately, (2) visit the landing page of each site, and (3) check if the extension interacts with the banner.

Overall, we find that each of the identified extensions interacts, on average, with 12 (65%) (SD: 21%; max: 95% min: 48%) of all banners. Appendix B presents a more detailed overview of the manual verification of each extension. Since all extensions can

interact with some cookie banners, we include all of them in our subsequent experiments.

4.4 Profiles Used in the Measurement

In our experiment, we use nine profiles to analyze the effects of tools that assist users in interacting with cookie banners. For each of the previously introduced tools (cf. Section 4.2), we create an individual measurement profile, that is, a distinct *OpenWPM* instance that installs one of the extensions to the Firefox browser. In all but one profile, we use the default configuration for all extensions. In the only exception, we configure *Consent-O-Matic* to accept all cookies, which rejects all non-necessary cookies by default.

In addition to the named browser extensions, we used three additional browser profiles in our measurement. The first two profiles are plain *OpenWPM* installations that we use as reference values. These baselines allow us to compare how each website behaves if no extension is used, allowing us to understand the magnitude of the effect of the analyzed extensions. It is important to note that these baselines are not supposed to serve as a lower or upper bound but only as a reference to understand how a website utilizes cookies if we do not interact with a cookie banner. Like all other profiles, one is run from a German IP address (Frankfurt am Main; DEU), and the other is run from a US IP address (Council Bluffs, IA; USA). For the third profile, we use a custom rule set and a heuristic approach to 'click' the "accept all cookies" button. We provide details on this implementation in the following paragraph. This profile is supposed to provide an upper bound on the cookies set if users accept all of them. Table 2 shows an overview of all used profiles.

Custom "Accept All" Extension. To understand an upper bound of used cookies, we implemented a tool that aims to automatically "accept all" cookies of a page by triggering the respective button on a cookie banner. In the following, we briefly describe our approach. We implemented a custom command for *OpenWPM* that uses a two-step process to identify "accept all" buttons, following a heuristic signature-based approach: (1) Identify an accept-all button and (2) check if the button is used in a privacy-related context.

First, we analyzed the top 1,000 sites from the Tranco list to generate a signature list to identify accept-all buttons. We manually identified 483 sites that use a consent banner and collected 97 distinct accept-all button labels. While these numbers seem low at first sight, they are in line (and even exceed) findings on the average number of sites that use cookie banners [45]. Second, to determine if a button we identified in step one is used to manage consent,

we profit from privacy-related words presented by Degeling et al. [8]. After a page finishes loading, our extension tries to find the 'accept all' button in the DOM object using the created signature list. Afterward, we test if the parent elements of the identified buttons contain privacy-related keywords (Step 2). While this approach has several limitations, we still think it is viable to serve as an upper bound for our experiment. In our experiment, we could interact with cookie banners on 17% (5,173) of the sites using this approach³. Our extension is publicly available (cf. Appendix A).

4.5 Identifying the Purposes of Cookies

This work investigates the effectiveness of browser extensions that interact with cookie banners. To better understand what kind of cookies are set in the different profiles, we classify them according to their purpose. In this work, we use four types of cookie categories proposed by the *International Chamber of Commerce UK* [20], which have been used by prior work [4, 42]. These categories are: (1) "Strictly Necessary Cookies" needed to provide the basic functionality of a website, (2) "Performance Cookies" aggregate (anonymously) the user's usage of the website, (3) "Functionality Cookies" personalize the website's usage, and (4) "Targeting/Advertising Cookies" used to track users or to display personalized ads.

To identify the categories, we profit from *Cookiepedia* [7], a pre-categorized cookies database. *Cookiepedia* is run by *OneTrust*, a privacy management software company that also provides a CMP. The database provides for a given cookie name (i.e., the key) with additional data on the cookie. For example, for the cookie "`__cf_gads`," the database provides that it is used for "Targeting/Advertising" by *DoubleClick*. In our experiment, we query the database for additional data on each identified cookie. This process might be error-prone as cookie classes are assigned by hand but are—from our point of view—the best approximation of cookie usage today. We have been able to classify 57% of all observed cookies. Appendix H provides an overview of the cookies that could not be classified.

In our experiment, we only focus on cookies in the cookie jar after crawling the last page for a given site; note our stateful crawling approach. We do this because some extensions (e.g., *CookieBlock*) might delete cookies from a category that the users opted out to. Hence, some cookies might be set by a page and deleted by an extension. In these cases, we assume the extension works as intended since undesired cookies are not persistent on the client. We evaluate the final set of cookies present by comprehending the actions logged by *OpenWPM* (i.e., creation, updating, and deletion).

We use *EasyList* to understand the effects of the analyzed extensions on user tracking. The list is a crowd-sourced effort to identify web tracking and might be incomplete or inaccurate to some extent. However, we expect that this limitation has only a marginal impact on our results and does not endanger our experiment's correctness.

4.6 Statistical Comparison of Profiles

We use the *Jaccard index* to compare cookies on a page or site. The index is used to gauge the similarity of sets and is defined as follows: $J(A, B) = \frac{|A \cap B|}{|A \cup B|}$. By design, the index ranges from 0 to 1, where 1 denotes that the sets are equals and 0 that they have no element

³These number relates to *all* sites and not only those that use a cookie banner.

in common. Using the index, we can compare the similarity of the cookies set by a page and the impact of the used extension.

In our analysis, we use statistical tests to determine if an extension significantly impacts the usage of cookies. In particular, we use the *non-parametric permutational multivariate analysis of variance* (PERMANOVA) [1] with $\alpha = .05$ to find statistically significant differences between the measures of independent groups. In our case, these groups are, for instance, the used profiles. Furthermore, we utilize the permutation test [46] (10,000 random permutations) to understand if there is or is not a difference between treatment groups (i.e., the extensions). Finally, we use the *Kruskal-Wallis* test ($\alpha = .05$) to assess if samples originate from the same distribution. Due to their non-parametric nature and flexibility, these three tests allow us to robustly verify our hypothesis without relying on specific distributional assumptions.

5 RESULTS

This section provides an overview of the measurement results. First, we provide a general overview of our measurement (Section 5.1), then we analyze the effect of the used extensions in terms of the number of cookies (Section 5.2), type of used cookies (Section 5.3), changes in the usage patterns of cookies (Section 5.4), and impact of the rank of a site (Section 5.5). Finally, we investigate the impact of the extensions on objects in the *HTML Web Storage* (Section 5.6) and analyze potential subsequent effects of the extensions (Section 5.7).

5.1 General Measurement Overview

In our measurement, we successfully crawled 29,660 (99%) sites, and 2.6M pages. On average, each crawler successfully visited 288k (SD: 5,180; min: 281k; max: 297k) webpages. We successfully crawled 12 (SD: 7; min: 1; max: 16) pages per site, on average (the landing page plus at most 15 subpages). The sites our crawlers could not reach are not meant to be visited by a human (e.g., link shorteners or content delivery networks). As already described, we include only sites in the evaluation successfully crawled by at least eight profiles (cf. Section 4). In our measurement, that filtering corresponds to 18,445 (62%) of the sites. The timing difference between two page visits across all profiles is 8 seconds (SD: 231; min: 1; max: 843), on average. The resulting (raw) data has a size of roughly 1.5 TB, which we make openly available (cf. Appendix A).

Overall, we record 415k (SD: 153k; min: 181k; max: 770k) cookies per profile on average. Using *Cookiepedia*, we could classify 57% of all cookies and 38% of all distinct cookies in our dataset (cf. Section 4.5). Fig. 1 presents a general overview of the number of cookies set per page of each profile in our measurement setup. By visual inspection, one can see that some of the analyzed extensions

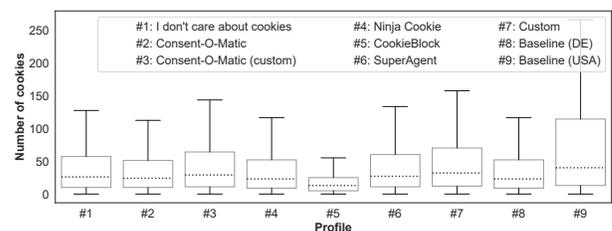


Figure 1: Cookies observed per page with different profiles.

Table 3: Overview of the high-level results per profile. The table lists (1) for each profile the number of third- and first-party cookies; (2) the different cookie categories: “Strictly Necessary” (Strict.), “Performance” (Perf.), “Functionality” (Func.), and “Targeting/Advertising” (Targ.); (3) the number of cookies in the local storage; and (4) number of observed tracking requests.

Nr.	Name	Group	Σ Cookies	1 st party c.	3 rd party c.	Strict.	Func.	Targ.	Perf.	Local storage	Track Req.
#1	I don’t care about cookies	AcceptAll	385,168	157,295	227,833	25,494	19,464	90,369	66,691	115,980	1,977,747
#2	Consent-O-Matic	RejectAll	369,264	155,883	213,346	25,562	18,764	83,438	64,526	113,266	1,891,632
#3	Consent-O-Matic (custom)	AcceptAll	424,089	165,639	258,412	26,917	21,730	100,607	70,371	124,697	2,100,955
#4	NinjaCookie	RejectAll	368,429	155,430	212,961	25,221	18,660	84,869	64,796	112,863	1,917,446
#5	CookieBlock	AcceptFunc	181,351	89,661	91,657	22,865	13,283	26,328	15,060	111,037	2,087,911
#6	SuperAgent	AcceptFunc	416,215	165,268	250,907	26,705	20,771	96,657	69,584	134,821	2,013,957
#7	Custom	AcceptAll	450,204	169,803	280,363	26,773	22,881	108,877	72,957	126,567	2,160,654
#8	None (EU)	Baseline	372,704	152,420	220,242	24,508	18,683	87,248	64,529	108,930	1,918,446
#9	None (US)	Baseline	76,540	183,232	586,267	32,073	34,044	186,613	89,657	152,513	2,542,748

substantially impact the setting of cookies in each profile. Thus, each extension impacts users’ privacy differently. The permutation test found a statistically significant (p -value < 0.001) impact of the profile on the number of cookies. However, the same impact was *not* observed when comparing profiles #2, #4, and #8 concerning the number of cookies per page. Profile #8 does not use any extensions but a plain, unmodified browser and hence does not interact with any cookie banner. Hence, we use it as our baseline (reference) for our study, run from the same location (the EU) as the profiles that use the extensions. Note that we do not assume that sites in the baseline profile use less (or more) cookies than sites in other profiles but that we can use the profile to understand how websites use cookies if the user does not interact with the cookie banner.

5.2 Effects of the Analyzed Extensions

Table 3 provides a high-level overview of the results measured with each profile. Overall, the results hint that the used extensions impact the presented metrics at different magnitudes. For example, we record that some essential metrics in terms of privacy (third-party cookies and targeting cookies) vary up to 640%. The average impact on the number of cookies is 83% (SD: 6; min: 0; max: 640). The permutation test found statistically significant (p -value < 0.001) impact of the profile (i.e., extension) on the number of total cookies, first-party cookies, third-party cookies, targeting and advertising cookies, usage of the local storage, tracking requests on page level (p -value < 0.001 for all named metrics). However, for the following metrics, we found *no* statistical impact for “strictly necessary” and “functional” cookies. The results indicate that the impact of the analyzed extensions in our measurement differs notably, which means that depending on the users’ requirements, different extensions need to be evaluated. Furthermore, the findings indicate that while all extension aim to fulfill a similar goal (i.e., blocking cookies for users), the rich landscape of different cookie banner implementations makes this an arguably hard and error-prone task. Similar to findings of previous work [9, 42], the measurement performed from the US shows considerably higher values in terms of cookies and tracking requests than the measurements performed from Europe. Furthermore, in the EU baseline profile (#8), the pages used fewer cookies than the profiles in which the extensions accepted (some) cookies and more than those that rejected most cookies. These observations indicate that our measurement approach works as intended, meaning extensions work on some sites and pages). Furthermore, we observe some expected behavior—which also shows

that our method works as intended—like increasing tracking requests if we accept all cookies or fewer targeting cookies if we refuse to use them. Based on the observed results, *CookieBlock* is the most effective tool for blocked cookies.

In the following, we analyze the profiles based on the expected functionality of the installed extension (column “group” in Table 3) and investigate the effects of these groups separately.

5.2.1 Accepting All Cookies. In our setup, we use three profiles (#1, #3, and #7; *AcceptAll* in Table 3) configured to accept all cookies on a page. Note that for the extension *I don’t care about cookies* (#1), it is unclear on which sites and pages it accepts or rejects cookies. Since the extension lists in its description that it primarily accepts cookies—when interacting at all—we include it in the “accept all” group. At first glance, grouping the extension “I don’t care about cookies” (#1) might seem counterintuitive since it is not documented how the extension interacts with the banners. However, after manually inspecting the plugin’s source code, we have classified it as “Accept all” because the used signatures overwhelmingly accept all cookies if the tool cannot hide the consent banners. To verify this evaluation, we contacted the developer of the extension, who confirmed our observation. As one would expect, the three profiles show an overall average increase of 10% in terms of cookies used by a page. The results show that, on average, a page in profile #1 uses 7 (SD: 11; min: 1; max: 433) cookies, 7.5 (SD: 12; min: 1; max: 433) cookies in profile #3, and 7.9 (SD: 13; min: 1; max: 433) cookies in profile #7. In comparison, in profile #8 (the EU baseline) a page uses 6.8 (SD: 11; min: 1; max: 848) cookies. Overall, the results show that the profiles in this group increase the number of first-party cookies by 4% and third-party cookies by 11% (1#: 4%; 3#: 11%; 7#: 19%), on average. Hence, the extensions in this group primarily impact third-party cookies.

Based on the numbers, the effect of the extensions seem to be similar. To better understand their impact on the used cookies, we compare the similarity of the cookies set by a page. Overall, a comparison of all profiles in this group with the baseline profile #8, using the Jaccard index (cf. Section 4.5), shows a medium similarity of .45 (SD: .46; min: 0; max: 1), on average. The same comparison, when excluding the baseline profile, shows a higher average similarity of .53 (SD: .46; min: 0; max: 1). We get the highest similarity when we test the similarity of #3 and #7 with an average of .64 (SD: .44; min: 0; max: 1). Thus, our analysis shows that even if all three extensions claim to allow similar types of cookies, there are

differences in the cookies present in the cookie jar. Therefore, the extensions interact differently with the banners or at least with different success. From a privacy perspective, this finding is concerning since it might be hard for users to evaluate the effect of an extension since sites set a similar amount of cookies, but the used cookies seem different. Counting the occurrence of potentially privacy-harming elements is a metric that ad blockers often use (i.e., blocked trackers) to evaluate effectiveness. Still, there might be better measures for cookie banner tools.

5.2.2 Accepting Functionality Cookies. In this section, we analyze the impact of the profiles that accept functional cookies (#5 and #6; *AcceptFunc* in Table 3). The two extensions impact the number of cookies a page uses differently. Compared to our baseline, profile #5 reduces the number of cookies per webpage by 40% but the extension used in profile #6 increases the number of cookies per webpage by 8%. On average, pages in profile #5 use 4 (SD: 6; min: 1; max: 433) cookies and pages in profile #6 use 7.2 (SD: 12; min: 1; max: 433) cookies. The notable difference in the effect of the compared tools can be attributed to how they work (cf. Section 4.2). *CookieBlock* (#5) uses an ML-approach and *SuperAgent* uses a rule-based approach. The results show that profile #6 increases the number of first-party cookies per page, on average, by 2% and third-party cookies by 7%. In contrast, profile #5 decreases the number of first-party cookies per page by 26% and of third-party cookies per page by 48%, on average. Note that #5 is the only profile showing a notable decrease in first-party and third-party cookies per page.

The distinctive differences between both profiles are also discernible in the similarity of the cookies set by the pages. The Jaccard index shows an average similarity of .48 (SD: .45; min: 0; max: 1) per page. Thus, only roughly half of the set cookies are present in both profiles. Comparing the two profiles with the baseline, the similarity test of the profiles #5, #6, and #8 shows a lower average similarity (mean: .39 (SD: .43; min: 0; max: 1)), which indicates that the extensions affect the used cookies. The results show that some websites seem to wait until a user interacts with a cookie banner before setting specific cookies. Still, approaches that actively delete cookies seem more effective than 'trusting' the page itself.

5.2.3 Rejecting All Cookies. Here, we inspect the impact of the extensions that assist users in rejecting all types of cookies (profiles #2 and #4; *RejectAll* in Table 3). The results show that, on average, the number of cookies per webpage is affected similarly by both extensions. Compared to the baseline, the number of cookies in profile #2 (6.8 cookies) and profile #4 (6.7 cookies) stays reasonably stable. Within the first- and third-party context, both profiles use a similar number of cookies. The permutation test found no statistically significant (p -value > 0.44) impact by the used extension on the average number of first-party and third-party cookies per page. Furthermore, the bootstrap confidence intervals for profile #2 and #4 are closely aligned, e.g., for third-party cookies we yield [21.1, 22.17] and [21.1, 22.13] for profile #2 and #4, respectively. This observation means that both extensions block the usage of cookies equally, which is expected as they deny a site to set any cookie.

Finally, we test the similarity of the cookies present in both profiles. Overall, we see a similarity of .70 (SD: .43; min: 0; max: 1), on average. The comparison of both profiles with the baseline profile shows a similarity of, on average, .58 (SD: .46; min: 0; max: 1). Thus,

even though the number of cookies in the store is similar, the set of used cookies differs. While we also see no high similarity in terms of used cookies for extensions that reject all cookies, we find that this type of extension leads to an overall decrease in used cookies and is similar to our baseline. This finding also highlights that statistically speaking, sites wait until a user chooses how cookies shall be used (the number of cookies remains similar to the baseline).

Lessons learned. Our observations show that some extensions can limit tracking and data collection by blocking cookies. In contrast, others may increase the number of cookies set by websites, including first-party and third-party cookies. This finding could result in more data collection and potentially negatively impact their privacy. Our findings revealed that the extension *CookieBlock*, in the category of "Accepting Functional Cookies", reduces the number of cookies drastically, while extensions in the "Rejecting All Cookies" category do not have any meaningful effect on the number of cookies. For users who utilize such tools to increase their privacy online, these findings mean that they should rely on a combination of tools (e.g., ad or tracking blockers) rather than only using a cookie banner tool.

5.3 Types of Used Cookies

In the previous section, we have shown that the analyzed extensions impact the cookie usage behavior of the investigated sites and pages. In this section, we dive deeper into the observed differences of used cookies to better understand the effectiveness of the extensions. Specifically, we analyze how each profile in our setup affects the cookies based on the cookie's purpose (cf. Section 4.5).

5.3.1 General Overview of Used Cookie Categories. Previously, we have analyzed in which context the pages set cookies and the impact of the analyzed extensions on them. In the following, we investigate for which purpose cookies are set and the effectiveness of the extensions to block specific types of cookies. Fig. 2 shows the distribution in the categories of used cookies across the nine profiles. Fig. 7 in Appendix D shows this distribution for each analyzed extension. A visual inspection of the figure shows that the extensions impact the used cookies differently. Indeed, the permutation test shows statistical significance (p -value < 0.001) in the used profile and the number of cookies in the categories. Compared to the baseline profile (#8), the results show in all profiles, except #5, that the number of cookies in a category per page increases (*Targeting/Advertising*: 14%; *Performance*: 5%; *Functionality*: 3%; *Strictly Necessary*: 5%). These observations apply to the first- and third-party contexts, but the increases we record are mainly in the third-party context. For instance, our analysis shows an increase of 2% for *Targeting/Advertising* cookies in the first-party context, while these cookies increase, on average, by 17% per page in third-party context. Notably, profile #7 increases the presents of third-party *Targeting/Advertising* cookies, on average by 45%. The results indicate that our method and the analyzed extensions work as intended because the number of cookies in different categories increases in specific profiles only (compared to the baselines).

We record small changes between the profiles regarding *Functionality* and *Strictly Necessary* cookies. However, the permutation test

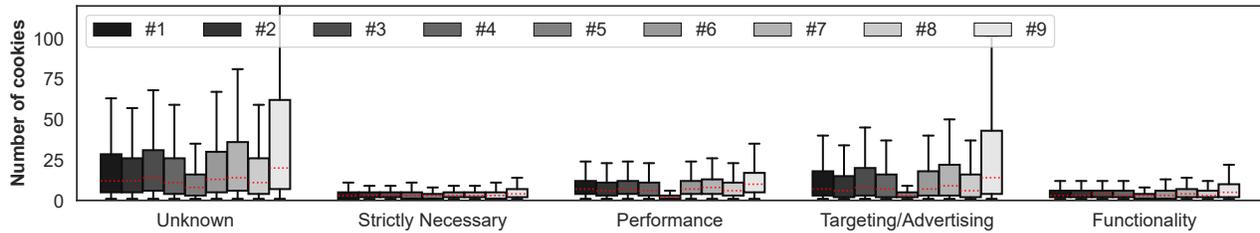


Figure 2: Number of cookies in the different profiles by cookie usage category. To increase readability, the y-axis is cut at 100. The upper whisker of the category “Unknown” in profile #9 is around 150.

reveals that the number of *Performance* and *Targeting/Advertising* cookies differ statistically significantly (p -value < 0.001) depending on the profile. In the baseline profile, we recorded, on average, 9.2 (SD: 11; min: 1; max: 102) tracking cookies per page. In profile #7, we measured an average increase of 28% in the usage of such cookies, and in profile #5, an decrease of 18% per page, on average. Our observation on *Unknown* cookies shows that the extensions have the same effect on these cookies as on targeting cookies. Thus, the results indicate that each extension impacts the number of cookies in each category differently. Extensions seem to have less impact on cookies that have potentially less impact on the users’ privacy (e.g., *Functionality*), which is a positive signal for users who aim to protect themselves against online tracking for different means.

Accepting All Cookies. For the extensions that accept all cookies (#1, #3, and #7), we see that the number of *Targeting/Advertising* cookies increases significantly (cf. Section 5.1) by 8% on average while the number in the other categories stayed stable (mean increase of 2%), compared to the baseline. Note that the increment mainly occurs in the third-party context. Overall, first- and third-party cookies increase by 4% and 11%, respectively. This (expected) increase in tracking cookies could negatively impact users’ privacy as third parties collect and analyze more data about them. To a lesser extent, this increase may result in increased loading times and data usage of pages [37]. Thus, users should weigh the benefits and drawbacks and decide based on their needs.

Accepting Functional Cookies. Overall, profile #5 reduces the number of observed cookies in all categories. On average, *Functionality* cookies go down by (9%), *Performance* cookies by (38%), *Targeting/Advertising* cookies by (36%), and the number of *Strictly Necessary* cookies stays stable (0%). The same analysis for profile #6 shows that the number of cookies per webpage stays almost similar. We record an average increases for *Performance* and *Strictly Necessary* cookies by (1%), for *Targeting/Advertising* cookies by (5%), and *Functionality* cookies stays stable (0%). The permutation test reveals a significant difference in the impact of all cookie types, except *Strictly Necessary*, between profiles #5 and #6. The results highlight the importance of carefully considering the privacy implications of different ‘cookie banner interaction’ extensions, even when they claim to provide similar functionalities. Therefore, it can be challenging for end users to select the extension that meets their privacy needs, as extensions with similar goals and functionalities can have a considerably different impact on users’ privacy.

Rejecting All Cookies. The profiles that reject all cookies (#2 and #4) have little impact on the number of cookies in each category.

Compared to the baseline profile (#8), we observed an increase in *Functionality* and *Strictly Necessary* cookies by 1% and a decrease for *Targeting/Advertising* cookies by 1%. This finding indicates that the extensions have almost no effect on the users’ privacy, but they eliminate the need to interact with (annoying) cookie banners). Furthermore, the results indicate that most sites wait until a user chooses before setting cookies.

5.3.2 Similarity of Observed Cookies and their Type. In the previous Section, we compared the number of cookies in each category in the different profiles. In the following, we compare, using the Jaccard index, the cookies set by different sites to understand which cookies are kept or removed by the extensions. To allow comparability, we compute the similarity of observed cookies within a functionality group of an extension (cf. Table 2). Hence, we only compare profiles where the pages can set the same types of cookies. Fig. 3 provides an overview of the similarity of the observed cookies in the different categories across the profile groups. “Functional” cookies have high similarity in all groups, including the one that should reject the usage of all cookies. The results of all other types of cookies show differences in the similarity of observed cookies and/or the effect of the different profile groups. Interestingly this also applies to “strictly necessary” cookies, which should be similar across all profiles.

On average, we find a high similarity of .75 across the groups that accept all cookies. However, within the group that only accepts functional cookies, we see an average similarity of .70. The group of “RejectAll” profiles shows an average similarity of .81. Hence, the extensions affect different cookies in these profiles, meaning other cookies are set and used on the analyzed pages. The observed discrepancy in similarity among the extension groups can likely be attributed to their different handling of cookies. On the one hand, the extensions within the “AcceptFunc” group use different approaches to accept cookies with varying success, which is evident by the number of cookies set between profiles #5 and #6. On the other hand, extensions within the “AcceptAll” group tend to trigger more third-party cookies, contributing to a lower similarity.

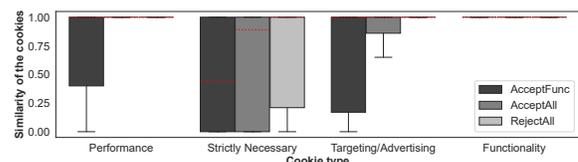


Figure 3: Similarity of the observed cookies for each cookie type by extension group.

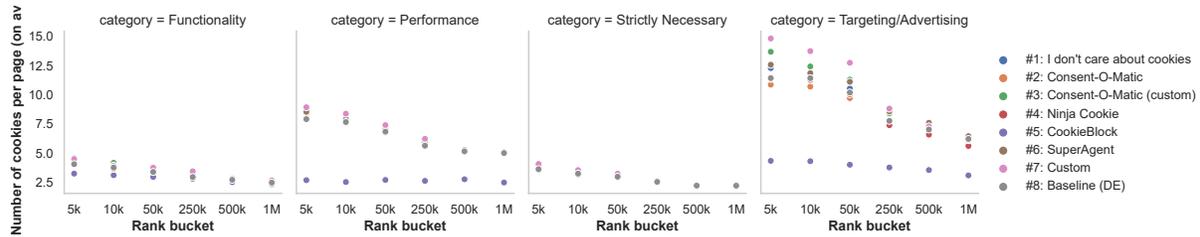


Figure 4: Number of cookies observed in each profile based on the popularity of a site and the type of a cookie.

Targeting/Advertising Cookies. Fig. 3 shows that the installed extensions affect targeting and advertising cookies quite differently. While we see a very high similarity (.83, on average) for extensions that reject all cookies, the similarity of other extension groups decreases. Notably, extensions that allow a more fine-grained control (e.g., group “AcceptFunc”) of cookies show a wide range in the similarity of cookies present on a page: .70 (SD: .44; min: 0; max: 1). Hence, the numbers support the claim that the usage of each extension seems to lead to a different set of cookies present in the browser, which makes it harder for users to assess their ability to help them to protect their privacy.

Strictly Necessary Cookies. The group of cookies without whom a page would not work correctly shows the most significant deviation (p -value < 0.001) in the similarity of all cookies based on the permutation test. The boxes range from low similarity ($\leq .2$) to perfect similarity (1) for all extension groups. On average, the results show a similarity of .64 for extensions that reject all cookies, .49 for extensions that accept all cookies, and .43 for extensions that only allow functional cookies. The results suggest that extensions have a more noteworthy impact on the similarity of strictly necessary cookies than targeting/advertising cookies. Thus, this disparity highlights the impact that extensions may have on a page’s functionality and suggests that there is a need for further investigation of the interaction between the extensions, the (real) purpose of necessary cookies, and the page’s functionality.

Lessons learned. Our analysis shows that different extensions for blocking cookies have varying effects on the types of cookies used by the analyzed websites. For example, extensions that accept all cookies tend to increase the number of ‘Targeting/Advertising’ cookies, while extensions that only accept functional cookies show different results. However, extensions that aim to reject all cookies have a minimal impact on the used cookies and, to some extent, on users’ privacy. This finding highlights the challenges users face if they seek to enhance their privacy while browsing the web, as it may take more work to determine the most suitable cookie interaction extension solely based on its stated purpose or description. The results of the similarity measurements suggest that users must carefully consider the functionality and implications of the extensions they choose to install and use them accordingly to balance the trade-off between privacy and functionality.

5.4 Observed Cookie Usage Patterns

In this section, we provide an overview of the usage of cookies (e.g., setting or changing a cookie). Our method allows us to distinguish between three operations (1) *added* if a new cookie is set, (2) *changed* if an existing cookie value is changed, and (3) *deleted* if a cookie has been deleted. Fig. 5 provides an overview of the different cookie access patterns resulting from using the different browser extensions over time. The figure shows the first and third-party cookie “usage patterns” over time after visiting a page. By “usage pattern,” we mean how a page uses cookies. For example, when does a page usually set first-party cookies or delete third-party cookies. The different extensions seem not to impact the usage pattern but affect the number of added, changed, or deleted cookies. Thus, in general, the users’ choice of which cookies might be used does not lead to different strategies for how pages utilize cookies.

Across all profiles, the number of *added* and *changed* cookies rises after the first few seconds after the browser visits the website. Thus, websites wait before setting cookies but use them whether users choose to accept cookies or not. However, by visual inspection, one can observe different cookie access patterns. For example, the number of added and changed cookies drops after their initial setting. After accessing a page, the peak is reached after 13 seconds (SD: 7; min: 0; max: 30) for first-party cookies, and after 15 seconds (SD: 6; min: 0; max: 30) for third-party cookies, on average. On average, operations on third-party cookies are processed 2 seconds later than first-party cookies for each profile. This difference indicates that the analyzed extensions actively lead to more operations on third-party cookies, especially during longer page visits. In profiles #5 (*CookieBlock*) and #9 (US profile), we see a notable increase in deletion operations, in contrast to all other profiles.

Lessons learned. Our results indicate that the different extensions impact the access patterns of websites to some extent. However, we did not observe notable differences, meaning that websites do not change their inner workings based on the users’ choice to use cookies. Nevertheless, the extensions impact the number of cookies used by each site.

5.5 Impact of a Site’s Rank

Next, we discuss the impact of a site’s rank on the effectiveness of the analyzed tools. More precisely, we analyze if the tools work better on popular (e.g., because they tend to use more CMPs) websites than on less popular sites. Fig. 4 provides an overview of the relation between the purpose of a cookie and the visited sites’ rank for the analyzed extensions. Overall, we see that popular websites use

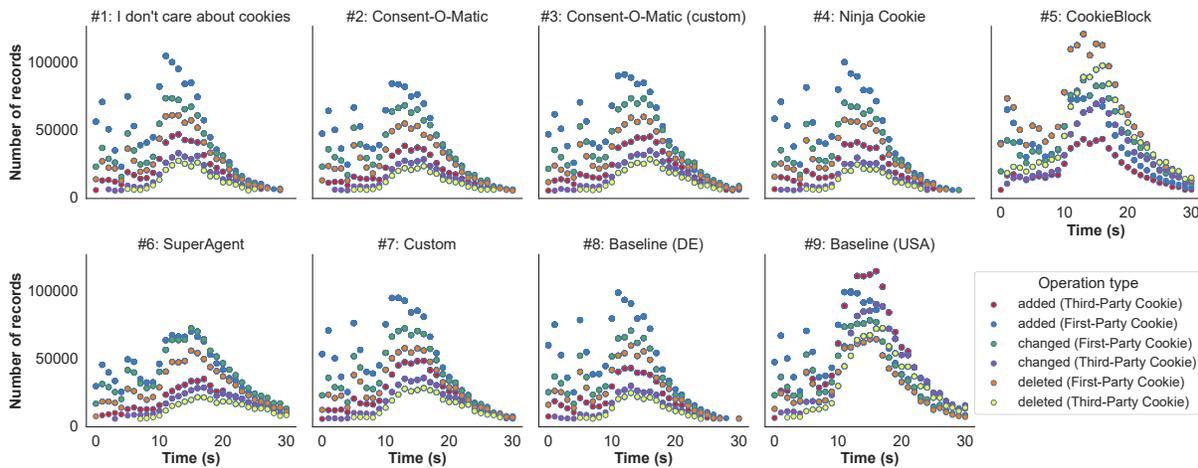


Figure 5: Cookie usage patterns of the analyzed webpages for the different profiles over time.

more cookies, which aligns with previous work, but also that the analyzed tools are more effective on these websites. Furthermore, the differences in the effectiveness of the analyzed extensions are more considerable for top sites. Thus, less popular sites seem to have a more diverse set of cookie banners, making it more challenging for extensions to interact with them than popular sites. The extension *CookieBlock* (profile #5) outperforms all other extensions across all cookie categories and ranks of sites. For all cookie types, we see a steady decrease in the number of used cookies, and the results show that the delta between the most effective extension (in terms of cookies) and the least effective extension also shrinks. For example, for targeting cookies on popular websites (rank 1–5k), the absolute difference in the number of set cookies is 4 (SD: 1.4; min: 11; max: 15), but for the least popular sites (rank 500,001–1M) the delta is only .8 cookies (SD: .28; min: 5.6; max: 6.4). These figures exclude profile #5 due to its meaningfully better performance.

Lessons learned. Our results indicate that the performance of cookie-blocking extensions may vary depending on the popularity of the visited website. Popular websites tend to use more cookies, and the differences in the effectiveness of the analyzed extensions are more pronounced for these sites. The best-performing extension in terms of blocking cookies, *CookieBlock*, demonstrates a clear advantage across all cookie categories and website ranks. However, it should be noted that the difference in the number of cookies set between the best and least effective extensions decreases for less popular websites. This observation suggests that users may experience varying levels of cookie blocking depending on the websites they visit, with potentially more diverse effects on popular websites.

5.6 Cookies in the Local Storage

The *HTML Web Storage API* (“local storage”) is a way to store data on the client. Compared to classic HTTP cookies, the storage can hold large amounts of data without negatively impacting a page’s performance. It works similarly to HTTP cookies as objects in the local storage are key-value pairs that can be accessed via

JavaScript [32]. In the following, we use the term *object* rather than *cookie* for objects stored in the local storage to make it easier to distinguish between HTTP cookies and data in the local storage. On average, we identified 122k objects in the local storage of the profiles. A site uses 12 (SD: 31; min: 1; max: 2,720) keys on average, which is 60% less than the average number of cookies.

Here, we analyze the differences the profiles cause in the elements stored in the local storage. Again, we use profile #8 as our baseline. On average, a site in profile #8 uses 11.6 objects in the local storage (cookies). The permutation test shows significant (p -value < 0.001) differences between the profiles that accept all cookies (#3 and #7) and the baseline. All other profiles do not significantly impact the number of local storage objects per site (p -value \geq 0.62). We record that the profiles #3 and #7 increase the number of keys per site by 13% (1.5 keys), on average. The increase could be because some sites may use the local storage to store more data than they could store in HTTP cookies, as the local storage allows for larger amounts of data. Furthermore, websites may use the local storage in combination with HTTP cookies to store data more persistently, as the data stored in local storage has a longer lifespan than cookies.

Profile #6 is the only profile in which fewer objects are present in the local storage (21%). We tried to use the classifications provided by *Cookiepedia* to classify the usage purpose of objects in the local storage. However, only 6% of the keys could be classified. Due to this low success rate and the low impact of extensions on this type of storage, we decided against a deeper analysis.

Lessons learned. The results show that the use of local storage objects, or key-value pairs stored on the client, is not significantly impacted by the extensions, except for profiles that accept all cookies. These profiles increased the number of objects stored in the local storage.

5.7 Potential Subsequent Effect

This section analyzes the potential impact of different extensions on tracking mechanisms. Overall, on average, interacting with consent banners increases the number of HTTP requests per page

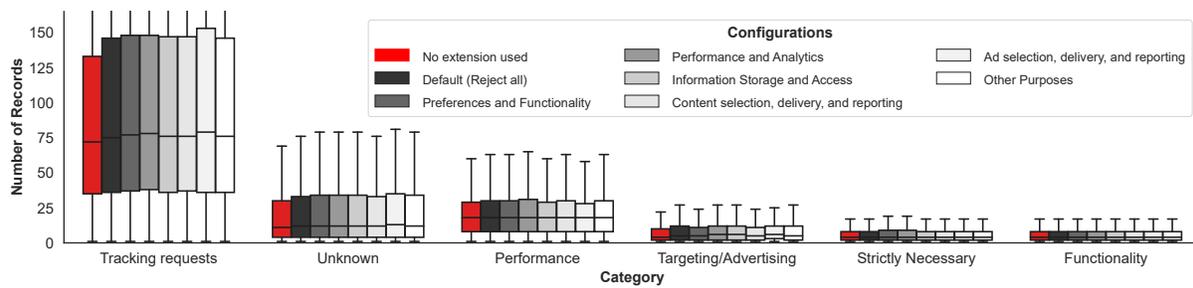


Figure 6: Distribution of number of records for tracking requests and different cookie categories grouped by different configurations (allowed cookie types) of *Consent-O-Matic*.

by 5%. The baseline profile (#8) has, on average, 89 (SD: 113; min: 1; max: 18,208) HTTP requests and 15.5 (SD: 57; min: 1; max: 18,154) tracking requests per page. Here, we dive deeper and test each extension group’s impact separately to understand their effects.

Accepting All Cookies. Overall, the results show that profiles #3 and #7 notably increase the number of tracking requests. In profile #3, we see an average increase of tracking requests per page by 12%, for profile #7 by 8%, and for profile #1 by only 3%. The permutation test found a statistical significance (p -value < 0.001) impact of these extensions on the number of tracking requests (cf. Section 4.5). Let’s only look at profile #7, for which we can measure the success rate of interaction with cookie banners. We find that the average number of tracking requests per page increases by (60%) (from 16.6 to 26) after interacting with the cookie banner, and the bootstrap confidence intervals for both profiles differ: we yield [17.12, 17.39] and [15.29, 15.90] for profile #7 and #8, respectively. The results demonstrate that accepting cookies leads to increased use of cookies and other privacy-harming techniques—in this case, user tracking. Thus, the used extensions can have a far-reaching impact on users’ privacy than ‘just’ blocking cookies.

Accepting Functionality Cookies. The observations of the group *AcceptFunc* show different results for both profiles. We see that profile #5 increases the average number of tracking requests per page by 8%, while we record a minimal increase of 3% for profile #6. The bootstrap confidence intervals for the profiles differ with [16.69, 16.96] and [15.86, 16.10] for profile #5 and #6, respectively. Note that profile #5 (*CookieBlock*) deletes cookies rather than telling websites not to use them. Thus, these finding indicates that extensions with the same purpose can have vastly different impacts on the users’ privacy, as different handling of cookies can lead to unexpected results, such as an increase in tracking requests. This result challenges the assumption that blocking cookies (especially targeting ones) will always lead to fewer tracking requests and highlights the need for further investigation of this phenomenon.

Rejecting All Cookies. We turn our analysis into the effects of the group *RejectAll*. Overall, we see that the profiles in this group do not affect the tracking requests on the pages, and the number of tracking requests per page stays relatively stable (mean of the profile #2: 15.4; #4: 15.5) in the observed profiles. The bootstrap confidence intervals for both profiles stay similar and yield [15.25, 15.50] and [15.29, 15.70] for profile #2 and #4, respectively. In this case, extensions that reject all cookies do not negatively impact the

number of tracking requests, making them a suitable option for end users interested in reducing tracking on their devices.

Lessons learned. It is important to note that each extension may affect privacy, functionality, and performance differently. For example, profile #5 substantially decreases the number of targeting cookies but increases the number of tracking requests per page. This finding highlights the challenge of balancing privacy and functionality for users. Additionally, the results show that different extensions, even with similar goals and functionalities, can have noteworthy differences in their impact on tracking mechanisms. For instance, profiles #5 and #6, in the *AcceptFunc* group, have different effects on the number of tracking requests per page, despite being designed to accept only functionality cookies. To conclude, the results demonstrate the complexities of balancing privacy and functionality for users and the potential consequences of using such extensions.

5.8 Different Extension Configurations

Some of the examined extensions offer different configuration options of which types of cookies the extension should allow or deny if the visited site provides a choice (cf. Table 2). To understand the effect on users’ privacy of these options, we measure the effectiveness of the different configurations of the most popular extension in our corpus that offers multiple configuration options: *Consent-O-Matic*. In particular, we performed separate measurements using eight different extension configurations. Note that we measure each available option individually but in combination with other options and only analyze the top 7,000 sites from our dataset (cf. Section 4.1). Appendix E describes the used setup in more detail.

Fig. 6 provides an overview of the results of the comparison between the different options (grayscale bars) and a profile with no extension installed (red bar). While the extension offers many different configuration options, each configuration’s outcome (effect) is similar. Note that the default configuration rejects all cookies. Within each cookie category, the results are comparable with the most variations for ‘*Tracking/Advertising*’ cookies. Note that the extension always accepts ‘*Functionality*’ cookies no matter which configuration is used, which explains this category’s almost identical effect size. Table 5 further provides an overview of the cross-comparison between the configurations. The experiment’s results show that the different configuration options have similar effects on

the tracking requests and number of cookie types, with minor variations. This finding suggests that the granular configuration options do not considerably enhance their privacy control in the current state of the extension. It is important to note that the analyzed extension is specialized in interacting with CMPs and, therefore, only has a limited effect on sites that do not use them. However, our findings indicate that due to missing standards on how to interact with cookie banners and banners that do not offer fine-grained options to limit the usage of cookies, offering multiple options to users does not help them, at least in our experiment. However, this might change in the future if consent mechanisms get more elaborate and automatically interacting with them gets easier.

5.9 Repeated Experiments

Measuring the Web is challenging due to its dynamic nature. Previous work has shown that repeated measurements are essential to assess the dynamic effects of a measurement [9] and that the vantage point [23] of a measurement can considerably affect an experiment’s outcome. To verify the correctness of our results, we performed three temporally close measurements from different vantage points from within the EU (cf. Appendix G) using *Amazon Web Services*. We use the same profiles as in our primary measurement (cf. Section 4.2) in all three measurements. However, we reduce the scale of each measurement for simplicity. On average, we visit 612,000 pages on 7,000 sites for each profile. We choose the top 7,000 sites from the used Tranco list and randomly sample 15 pages from each of the used buckets (cf. Section 4.1).

For this experiment, we resort to the volume and type of observed cookies and tracking requests measured in each profile and compare the three measurements based on these fundamental data points. The rationale is that if we find statistically significant overlaps in these numbers across the measurements, one can assume that our overall results are correct. We find that the analyzed pages set a similar amount of cookies in the used profiles across the three measurements: on average, we observe a delta of 2.2 (SD: 11; min: 767; max: 324) cookies on the analyzed pages. The PERMANOVA test found no statistical significance (p -value > 0.87) neither between the measurement runs (i.e., country and time) nor the number of cookies set on each page in each profile. However, the Kruskal-Wallis test shows varying results for some combinations of profiles on specific metrics, such as tracking requests and types of cookies. For instance, five profiles show significant differences in targeting cookies (p -value < 0.001). Given these varying results, we investigate the magnitude of these differences by measuring the effect size. The effect size provides additional information beyond statistical significance, measuring the degree of difference. In this study, we use the Eta-squared (η^2) measure to assess the effect size, reflecting the proportion of total variance attributable to an effect. Our results show that all computed η^2 values are less than 0.01 (SD: .0003; min: .00001; max: .001), suggesting that while some differences are statistically significant, their practical impact is minimal. We attribute these findings to the Web’s dynamic nature and argue that while some differences between measurement runs exist, they do not tamper with the reliability or correctness of the results.

6 LIMITATIONS AND ETHICS

Our study comes with several limitations. The main limitation of our work is that we cannot measure that the analyzed extensions work as they claim they do. More specifically, we cannot measure if a website ignores a user’s choice (or an extension’s choice on the user’s behalf) or if an extension is not working correctly. Our experiment can only report different tools’ effects on the cookies in the browser’s cookie jar. However, our study aims to understand the effect of these extensions on users’ privacy, which means that the execution of the user’s choice is out of the scope of this work and should be explored in future work. Furthermore, our computation of the similarity of set cookies only compares the presents of cookies based on their keys and not their functionality. A page could set a user-specific key, but this key could be used for the same purpose across all users. Finally, since our crawler does not interact with the pages as a regular user would, we only observe an excerpt of a page’s functionality. Therefore, our results can only be seen as a lower bound and incomplete. However, implementing a system that automatically generically interacts with any given page to trigger a different state is probably impossible and is out of the scope of this work. Our approach follows best practices and uses well-established tools to reduce the effect of the named limitations.

Our study uses *OpenWPM* to visit the sites and pages in our dataset. However, the way the framework instruments the browser is detectable by a website so that the website might serve different content if and when *OpenWPM* is detected [28]. We have not performed any means to circumvent such detection mechanisms as this might be ethically questionable. Furthermore, a site might detect that we use the *Google Cloud Platform* and *Amazon Web Services* for visiting the pages and serves different or no content [21]. We filter out sites and pages that we cannot reasonably compare to contain the effects of such crawling detection mechanisms. Our filter policy excludes a site if it does not serve a page for at least one crawler, meaning the crawler has been detected. If the site detects all our crawlers but acts consistently across all crawlers (e.g., showing a blank page), our filter does not take effect. Note that this does not change the overall findings of our work notably: When a blank site is served, no cookie banners are present, and no cookies are used. In this case, all extensions have the same effect.

Since our study excludes human subjects, it was exempt from the Institutional Review Board’s review. However, like all web measurement studies [9, 23, 42, 43], our study has some inevitable ethical considerations. By automatically visiting several pages, we create traffic and use up resources. We limit the effects by only visiting up to 15 website pages. Also, our crawlers may get served ads, increasing the difference in served and viewable impressions and thus potentially decreasing the revenue of the advertiser.

7 RECOMMENDATIONS

Based on the results, we provide recommendations for users, developers, and researchers of “cookie banner interaction” tools.

Recommendations for Users. Users might install “cookie banner interaction tools” for two reasons: (1) they do not want to be bothered by cookie banners, or (2) they want to automate the process

to communicate their cookie preferences. Users that use the extensions to improve their privacy online should choose an extension that limits the usage of cookies that may impact their privacy (e.g., “tracking cookies”) and should *not* use an extension or configuration that accepts all cookies. At the time of writing this paper, the extension *CookieBlock* showed the best performance regarding blocking (deleting) unwanted cookies. Finally, the results suggest that other tracking activities may increase if cookie usage is limited. Thus, privacy-aware users need to consider using additional tools (e.g., *Privacy Badger* [40], *Ghostery* [14], or *uBlock Origin* [16]) that help them to limit other forms of Web tracking.

R1: Users of cookie banner interaction tools should additionally use ad or tracking blockers to protect their privacy and should not use configurations that accept all cookies.

Additionally, users must be cautious when choosing a tool regarding its coverage. Users need to generally understand how a tool works and which limitations come from specific design choices. For example, our analysis shows that popular tools focus on interacting with popular CMPs, meaning such tools will only work on sites that utilize these CMPs. Other tools aim to hide banners, which might tamper with a site’s functionality as users cannot allow specific cookies; thus, some functions might not work correctly. Nevertheless, in our experiment, none of the analyzed extensions broke a considerable number of sites.

R2: Depending on their goals, users should check how an extension works (i.e., hiding a banner vs. interacting with them) and if the handling method meets their expectations.

Recommendations for Developers. Developers develop the extensions to (1) maximize the coverage of the tool (i.e., number of banners and websites to interact with) and (2) to implement the user’s choice as well as possible. To maximize their coverage, extensions should not only resort to interacting with popular CMPs but extend their scope. It is challenging for developers to keep up with the numerous implementations of different cookie banners. Similar to detecting trackers and ads, a community-driven approach to identify banners and choices within them could help handle this scaling issue. A promising attempt in this direction is the *EasyList Cookie List* [11], which can be used as a starting point to identify banners. Furthermore, our results suggest that the analyzed extensions perform notably worse on less popular sites, which means that such sites should be included when testing an extension.

R3: Developers can include community-driven projects helping identify banners to maximize the coverage of the tool. Also, the testing of applications should incorporate less popular sites.

According to the results, the only tool that actively checked for a cookie’s purpose (*CookieBlock*) showed the best performance regarding the number of present cookies. However, such classification of cookies is challenging. Again, community-driven projects like the named *EasyList Cookie List* can assist developers in this regard.

R4: Developers may consider actively checking for (and deleting) cookies that contradict the user’s privacy configuration to increase a tool’s effectiveness and, to some extent, its coverage.

Recommendations for Researchers. Among many other things, researchers might want to utilize “cookie banner interaction” tools if they are interested in how websites use cookies or to simulate a more ‘realistic behavior’ of a site in a Web measurement. However, our results suggest that the usage of such tools creates another, to some extent uncontrollable, bias as the effectiveness of the tools varies, and it is (without manual analysis) not possible to assess on which pages the tools could interact with a banner and on which they could not. Therefore, researchers must carefully consider whether (or to what extent) such a tool is tenable.

R5: When planning a study, researchers should evaluate on which pages the utilized tools work (e.g., by instrumenting them) to assess the impact on the outcome of the study.

Finally, researchers use the tools to reduce limitations affecting Web measurements (conducted from the EU). However, our results show that the tools would impact a study’s outcome differently, even if configured to accept the same categories of cookies, which hinders the comparability of different studies. Thus, it is inevitable to contemplate the effect a tool might have on a study.

R6: Researchers need to identify new limitations and challenges that arise when utilizing a cookie banner tool and should rigorously weigh the benefits and drawbacks for the study.

8 CONCLUSION

This paper compares six different tools that assist users in interacting with cookie banners in a large-scale web measurement study. We configured the extensions to communicate different types of cookies that the users are willing to accept or reject. The results show that the effectiveness of the different tools, in terms of cookies set by a visited page, varies significantly. We find differences in the number of used cookies, type of used cookies, and set cookies (i.e., different keys), even for similar configured tools. Our findings reveal that if users find consent banners disruptive while browsing, they can use the extensions tested in this study to eliminate them with varying degrees of success. However, it is essential to note that extensions with similar purposes lead to different sets of cookies present on a page. For example, *Consent-O-Matic* (#2 and #3) only works with so-called “Consent Management Platforms” (CMPs) and may not help block other types of consent banners.

Furthermore, the results suggest that blocking targeting cookies does not necessarily result in improving user privacy, as it can lead to an increase in tracking requests. This observation could tamper with the intention of the users to counter Web tracking attempts by sites. Therefore, end users seeking to protect their privacy—through cookie banner blocking extensions—should consider the trade-offs between functionality and privacy and select the tool that best fits their needs. However, further research must clarify the correlation between tracking requests and blocking tracking cookies.

In summary, the analyzed extensions aim to provide a user-friendly way to control the cookies users want to accept or reject. However, this study shows that the effectiveness of the extensions varies and that it can be challenging for users to assess or compare the extensions’ impact on their privacy. Our findings show that, while these tools grow in popularity, they can be helpful for users to give or withdraw consent for the usage of different cookie types.

ACKNOWLEDGMENTS

The authors gratefully acknowledge funding from the *Helmholtz Association* (HGF) within topic “46.23 Engineering Secure Systems,” the *Federal Ministry Economic Affairs and Climate Action of Germany* (grants 01MK20008E “Service-Meister” and 16KIS1629 “Ubi-Trans”), and the the *Federal Office for Information Security of Germany* (01MO23033B “5Guide”).

REFERENCES

- [1] Marti J Anderson. 2001. A new method for non-parametric multivariate analysis of variance. *Austral ecology* 26, 1 (2001).
- [2] Waqar Aqeel, Balakrishnan Chandrasekaran, Anja Feldmann, and Bruce M. Maggs. 2020. On Landing and Internal Web Pages: The Strange Case of Jekyll and Hyde in Web Performance Measurement. In *ACM SIGCOMM Internet Measurement Conference (IMC)*. <https://doi.org/10.1145/3419394.3423626>
- [3] Adam Barth. 2011. *HTTP State Management Mechanism*. RFC 6265. Internet Engineering Task Force. <https://tools.ietf.org/html/rfc6265>
- [4] Dino Bollinger, Karel Kubicek, Carlos Cotrini, and David Basin. 2022. Automating Cookie Consent and GDPR Violation Detection. In *USENIX Security Symposium (USENIX Sec.)*.
- [5] BuiltWith. 2022. Privacy Compliance Usage Distribution in the Top 1 Million Sites. <https://web.archive.org/web/20201021075918/https://trends.builtwith.com/widgets/privacy-compliance/>.
- [6] Stefano Calzavara, Tobias Urban, Dennis Tatang, Marius Steffens, and Ben Stock. 2021. Reining in the Web’s Inconsistencies with Site Policy. In *Symposium on Network and Distributed System Security (NDSS)*. <https://doi.org/10.14722/ndss.2021.23091>
- [7] Cookiepedia by OneTrust. 2022. Largest Database of Pre-Categorized Cookies. <https://cookiepedia.co.uk/>.
- [8] Martin Degeling, Christine Utz, Christopher Lentzsch, Henry Hosseini, Florian Schaub, and Thorsten Holz. 2019. We Value Your Privacy ... Now Take Some Cookies: Measuring the GDPR’s Impact on Web Privacy. In *Symposium on Network and Distributed System Security (NDSS)*. <https://doi.org/10.1007/s00287-019-01201-1>
- [9] Nurullah Demir, Matteo Große-Kampmann, Tobias Urban, Christian Wressneger, Thorsten Holz, and Pohlmann Norbert. 2022. Reproducibility and Replicability of Web Measurement Studies. In *International Conference on World Wide Web (TheWebConf)*. <https://doi.org/10.1145/3485447.3512214>
- [10] Nurullah Demir, Tobias Urban, Kevin Wittek, and Norbert Pohlmann. 2021. Our (in)Secure Web: Understanding Update Behavior of Websites and Its Impact on Security. In *Conference on Passive and Active Measurement (PAM)*. https://doi.org/10.1007/978-3-030-72582-2_5
- [11] EasyList. 2023. EasyList Cookie List. <https://secure.fanboy.co.nz/fanboy-cookiemonster.txt>.
- [12] Steven Englehardt and Arvind Narayanan. 2016. Online Tracking: A 1-Million-Site Measurement and Analysis. In *ACM Conference on Computer and Communications Security (CCS)*. <https://doi.org/10.1145/2976749.2978313>
- [13] Imane Fouad, Cristiana Santos, Feras Al Kassar, Nataliia Bielova, and Stefano Calzavara. 2020. On Compliance of Cookie Purposes with the Purpose Specification Principle. In *International Workshop on Privacy Engineering (IWPE)*. <https://doi.org/10.1109/EuroSPW51379.2020.00051>
- [14] Ghostery GmbH. 2023. Ghostery—Privacy Ad Blocker. <https://www.ghostery.com/>.
- [15] Colin M. Gray, Cristiana Santos, Nataliia Bielova, Michael Toth, and Damian Clifford. 2021. Dark Patterns and the Legal Requirements of Consent Banners: An Interaction Criticism Perspective. In *ACM SIGCHI Conference on Human Factors in Computing Systems (CHI)*. <https://doi.org/10.1145/3411764.3445779>
- [16] Raymond Hill. 2023. uBlock Origin—Free, open-source ad content blocker. <https://ublockorigin.com/>.
- [17] Maximilian Hils, Daniel W. Woods, and Rainer Böhme. 2020. Measuring the Emergence of Consent Management on the Web. In *ACM SIGCOMM Internet Measurement Conference (IMC)*. <https://doi.org/10.1145/3419394.3423647>
- [18] Xuehui Hu, Nishanth Sastry, and Mainack Mondal. 2021. CCCC: Corraling Cookies into Categories with CookieMonster. In *ACM Web Science Conference (WebSci)*. <https://doi.org/10.1145/3447535.3462509>
- [19] IAB Europe. 2022. IAB Europe Transparency & Consent Framework Policies. <https://iab europe.eu/iab-europe-transparency-consent-framework-policies/>.
- [20] International Chamber of Commerce UK. 2012. ICC UK Cookie guide. https://www.cookie law.org/wp-content/uploads/2019/12/icc_uk_cookiesguide_revnov.pdf.
- [21] Luca Invernizzi, Kurt Thomas, Alexandros Kapravelos, Oxana Comanescu, Jean-Michel Picod, and Elie Bursztein. 2016. Cloak of Visibility: Detecting When Machines Browse a Different Web. In *IEEE Symposium on Security and Privacy (S&P)*. <https://doi.org/10.1109/SP.2016.50>
- [22] Nikhil Jha, Martino Trevisan, Luca Vassio, and Marco Mellia. 2022. The Internet with Privacy Policies: Measuring The Web Upon Consent. *ACM Trans. Web* 16, 3, Article 15 (sep 2022), 24 pages. <https://doi.org/10.1145/3555352>
- [23] Jordan Jueckstock, Shaown Sarker, Peter Snyder, Aidan Beggs, Panagiotis Papadopoulos, Matteo Varvello, Ben Livshits, and Alexandros Kapravelos. 2021. Towards Realistic and Reproducible Web Crawl Measurements. In *International Conference on World Wide Web (TheWebConf)*. <https://doi.org/10.1145/3442381.3450050>
- [24] Daniel Kladnik. 2022. I don’t care about cookies 3.4.2—Get rid of cookie warnings from almost all websites! <https://www.i-dont-care-about-cookies.eu/>.
- [25] David Klein, Marius Musch, Thomas Barber, Moritz Kopmann, and Martin Johns. 2022. Accept All Exploits: Exploring the Security Impact of Cookie Banners. In *Proceedings of the 38th Annual Computer Security Applications Conference (Austin, TX, USA) (ACSAC ’22)*. <https://doi.org/10.1145/3564625.3564647>
- [26] Michael Kretschmer, Jan Pennekamp, and Klaus Wehrle. 2021. Cookie Banners and Privacy Policies: Measuring the Impact of the GDPR on the Web. *ACM Transactions on the Web* 15, 4 (2021). <https://doi.org/10.1145/3466722>
- [27] Chiara Krisam, Heike Dietmann, Melanie Volkamer, and Oksana Kulyk. 2021. Dark Patterns in the Wild: Review of Cookie Disclaimer Designs on Top 500 German Websites. In *European Symposium on Usable Security (EuroUSEC)*. <https://doi.org/10.1145/3481357.3481516>
- [28] Benjamin Krumnow, Hugo Jonker, and Stefan Karsch. 2022. How Gullible Are Web Measurement Tools? A Case Study Analysing and Strengthening OpenWPM’s Reliability. In *International Conference on Emerging Networking Experiments and Technologies (CoNEXT)*. <https://doi.org/10.1145/3555050.3569131>
- [29] Oksana Kulyk, Annika Hilt, Nina Gerber, and Melanie Volkamer. 2018. “This Website Uses Cookies”: Users’ Perceptions and Reactions to the Cookie Disclaimer. In *European Workshop on Usable Security (EuroUSEC)*. <https://doi.org/10.14722/eurosec.2018.23012>
- [30] Victor Le Pochat, Tom Van Goethem, Samaneh Tajalizadehkhoob, Maciej Korczyński, and Wouter Joosen. 2019. Tranco: A Research-Oriented Top Sites Ranking Hardened Against Manipulation. In *Symposium on Network and Distributed System Security (NDSS)*. <https://doi.org/10.14722/ndss.2019.23386>
- [31] Célestin Matte, Nataliia Bielova, and Cristiana Santos. 2020. Do Cookie Banners Respect my Choice?: Measuring Legal Compliance of Banners from IAB Europe’s Transparency and Consent Framework. In *2020 IEEE Symposium on Security and Privacy (SP)*. <https://doi.org/10.1109/SP40000.2020.00076>
- [32] MDN Web Docs. 2022. Window.localStorage. <https://developer.mozilla.org/en-US/docs/Web/API/Window/localStorage>.
- [33] Ninja Cookie. 2022. Have you had enough of cookie banners? Forget about them! Ninja Cookie will take care of these and can say “no” to them for you! <https://ninja-cookie.com/>.
- [34] Midas Nouwens, Rolf Bagge, Janus Bager Kristensen, and Clemens Nylandstedt Klokmoose. 2022. Consent-O-Matic: Automatically Answering Consent Pop-Ups Using Adversarial Interoperability. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems (CHI)*. <https://doi.org/10.1145/3491101.3519683>
- [35] Midas Nouwens, Ilaria Liccardi, Michael Veale, and Lalana Karger, David and Kagal. 2020. Dark Patterns after the GDPR: Scraping Consent Pop-Ups and Demonstrating Their Influence. In *ACM SIGCHI Conference on Human Factors in Computing Systems (CHI)*. <https://doi.org/10.1145/3313831.3376321>
- [36] 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 *International Conference on World Wide Web (WWW)*. <https://doi.org/10.1145/3442381.3450056>
- [37] Panagiotis Papadopoulos, Nicolas Kourtellis, and Evangelos Markatos. 2019. Cookie Synchronization: Everything You Always Wanted to Know But Were Afraid to Ask. In *International Conference on World Wide Web (WWW)*. <https://doi.org/10.1145/3308558.3313542>
- [38] Iskander Sanchez-Rola, Matteo Dell’Amico, Platon Kotzias, Davide Balzarotti, Leyla Bilge, Pierre-Antoine Vervier, and Igor Santos. 2019. Can I Opt Out Yet? GDPR and the Global Illusion of Cookie Control. In *ACM Asia Conference on Computer and Communications Security (AsiaCCS)*. <https://doi.org/10.1145/3321705.3329806>
- [39] Super Agent. 2022. No more Pop-Ups! Privacy can be Simple. <https://www.super-agent.com/>.
- [40] The Electronic Frontier Foundation. 2023. Privacy Badger. <https://privacybadger.org/>.
- [41] Michael Toth, Nataliia Bielova, and Vincent Roca. 2022. On Dark Patterns and Manipulation of Website Publishers by CMPs. In *Proceedings on Privacy Enhancing Technologies (PoPETs)*. <https://doi.org/10.56553/popets-2022-0082>
- [42] Tobias Urban, Martin Degeling, Thorsten Holz, and Norbert Pohlmann. 2020. Beyond the Front Page: Measuring Third Party Dynamics in the Field. In *International Conference on World Wide Web (TheWebConf)*. <https://doi.org/10.1145/3366423.3380203>
- [43] Tobias Urban, Dennis Tatang, Martin Degeling, Thorsten Holz, and Norbert Pohlmann. 2020. Measuring the Impact of the GDPR on Data Sharing. In *ACM Asia Conference on Computer and Communications Security (AsiaCCS)*. <https://doi.org/10.1145/3321705.3329806>

//doi.org/10.1145/3320269.3372194

- [44] Christine Utz, Martin Degeling, Sascha Fahl, Florian Schaub, and Thorsten Holz. 2019. (Un)informed Consent: Studying GDPR Consent Notices in the Field. In *ACM Conference on Computer and Communications Security (CCS)*. <https://doi.org/10.1145/3319535.3354212>
- [45] Web Almanac By HTTP Archive. 2022. Privacy: Compliance with privacy regulations. <https://almanac.httparchive.org/en/2022/privacy#consent-management-platforms>.
- [46] William J. Welch. 1990. Construction of Permutation Tests. *J. Amer. Statist. Assoc.* 85, 411 (1990). <https://doi.org/10.1080/01621459.1990.10474929>
- [47] Daniel W. Woods and Rainer Böhme. 2022. The Commodification of Consent. *Computers and Security* 115, C (2022). <https://doi.org/10.1016/j.cose.2022.102605>

A AVAILABILITY OF DATA & CODE

To foster future research, we release our code, queries for the entire data processing pipeline and evaluation, and other supplementary information openly online at:

<https://github.com/internet-sicherheit/A-Large-Scale-Study-of-Cookie-Banner-Interaction-Tools-and-Their-Impact-on-Users-Privacy>

Furthermore, we provide all raw data collected during our experiment:

Part 1: <https://doi.org/10.35097/1708>

Part 2: <https://doi.org/10.35097/1717>

B SUCCESS RATES OF THE ANALYZED EXTENSION IN THE MANUAL TEST

Table 4 provides an overview of the respective success rates of the analyzed extensions. We assume that the ‘interaction’ was successful if a prompted cookie banner disappears automatically (or is not present) when we visit a page with an installed extension. Otherwise, we assume that the extension does not work on the visited page. Note that the extension *CookieBlock* does not interact with cookie banners but actively deletes cookies of specific categories. Overall, we see that the analyzed extension have different success rates in interaction with the banner, but no extension broke a page in our manual experiment.

Table 4: Overview of the success rates of the extension to interact with the 19 cookie banners in our manual test.

Name	Successful interaction		Unsuccessful interaction	
I don't care about cookies	18	95%	1	5%
Consent-O-Matic	10	53%	9	47%
Ninja Cookie	12	63%	2	37%
SuperAgent	9	47%	10	53%
CookieBlock	—	—	—	—

C FAILURE RATES OF THE CRAWLER

We only include sites in our experiment if at least eight profiles successfully crawled them to ensure a fair and meaningful comparison of all tools. This filtering resulted in the exclusion of 36% of all sites. It is worth noting that this (high) rate is solely attributed to

the combination of the profiles—each profile has a failure rate of <15%. More precisely, profile #9 has the highest (15%) and profile #6 the lowest (13%) failure rate; the mean failure rate is 14%. These numbers are typical for large-scale Web measurements [6, 9].

D COOKIES IN THE DIFFERENT PROFILES

Fig. 7 shows the number of cookies present in each profile. It is evident from the figure that each profile has a distinct distribution of cookies, with profile #5 having the lowest number of cookies and profile #9 having the highest. These results highlight that the different browser extensions can have a considerable effect on the type and quantity of cookies set by a website.

E CONFIGURATIONS OF CONSENT-O-MATIC

The *Consent-O-Matic* extension offers, in addition to the standard configuration, six configuration options to allow or deny different types of cookies: (1) Preferences and Functionality; (2) Performance and Analytics; (3) Information Storage and Access; (4) Content selection, delivery, and reporting; (5) Ad selection, delivery, and reporting; and (6) Other purposes. Thus, we analyzed seven different configurations of the extension. One should note that the default configuration of *Consent-O-Matic* aims to reject all cookies that are not necessary for the page to work. To experiment with a reasonable time frame, we analyzed each option individually and not a combination of them, which would result in analyzing 63 configurations ($\binom{6}{6} + \binom{6}{5} + \binom{6}{4} + \binom{6}{3} + \binom{6}{2} + \binom{6}{1} = 63$). Furthermore, we choose the top 7,000 sites from the utilized Tranco list and randomly sampled 15 pages (cf. Section 4.1), visited them with our measurement framework and filtered the pages to analyze according to our filtering rules (cf. Section 4). We conducted the experiments between 05/08/2023 and 05/12/2023 from a European IP address (Germany) using *Amazon Web Services AWS*.

F STATISTICAL EFFECTS OF DIFFERENT CONSENT-O-MATIC CONFIGURATIONS

Table 5 shows the p -values when computing the statistical significance, using the Kruskal-Wallis test, of the different *Consent-O-Matic* profiles regarding the presence of tracking requests and different cookie types. The lighter-gray fields highlight the profile combinations for which we found statistical significance. Especially the profile that limits the use of “Performance” cookies (Perfm.) shows a notably different behavior than most other profiles. For the configurations that allow functional cookies (Func.) and “Information Storage and Access” cookies (Infor.), we get mixed results meaning that they show statistical difference to roughly half of the configurations. Most profiles (i.e., the baseline, the default configuration, reports, ads, and other) show a very similar (i.e., almost no statistical significance) behavior compared to the other configurations. These comparisons concerning the statistically significant impact on the metrics of interest show that most configuration options have little impact on them.

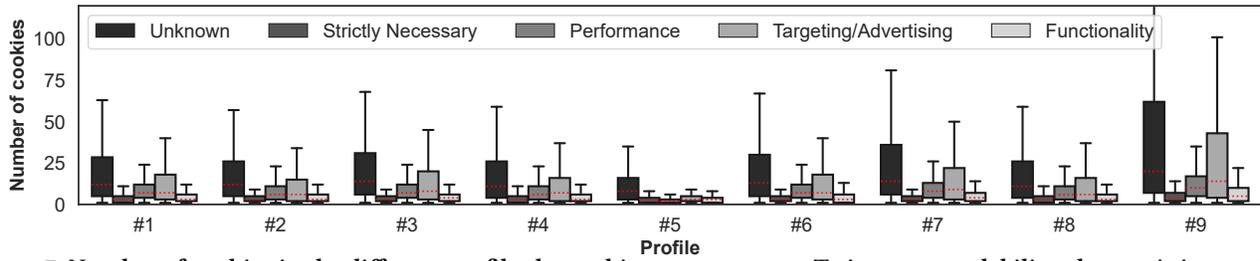


Figure 7: Number of cookies in the different profiles by cookie usage category. To increase readability, the y-axis is cut at 100. The upper whisker of the category “Unknown” in profile #9 is around 150.

Table 5: Kruskal-Wallis test results for tracking request numbers and cookie types across different *Consent-O-Matic* configurations.

	Baseline	Default	Func.	Perfm.	Infor.	Report	Ads	Other
Baseline	—	1.00	0.06	0.00	1.00	1.00	1.00	1.00
Default	1.00	—	0.02	0.00	1.00	1.00	1.00	1.00
Func.	0.06	0.02	—	1.00	0.04	0.46	0.00	0.01
Perfm.	0.00	0.00	1.00	—	0.00	0.05	0.00	0.00
Infor.	1.00	1.00	0.04	0.00	—	1.00	1.00	1.00
Report	1.00	1.00	0.46	0.05	1.00	—	1.00	1.00
Ads	1.00	1.00	0.00	0.00	1.00	1.00	—	1.00
Other	1.00	1.00	0.01	0.00	1.00	1.00	1.00	—

G LOCATIONS OF THE REPEATED MEASUREMENTS

Table 6 shows the three locations that we used for our repeated control experiments. We performed all experiments successive using different *Amazon Web Services* instances during April 2023.

Table 6: Locations and times of the repeated measurements.

#	City	Country	Start Date	End Date
1	Frankfurt	DEU	04/21/23	04/24/23
2	Paris	FRA	04/24/23	04/28/23
3	Stockholm	SWE	04/28/23	05/01/23

H ANALYSIS ON COOKIES THAT COULD NOT BE CLASSIFIED

In the following, we analyze the cookies that could not be identified by *Cookeipedia* [7], and test their possible impacts on our analysis. Overall, we could not classify 43% of the cookies, which means that

Cookeipedia could not identify their purpose (cf. Section 4.5). To better understand if a manual classification of some of the observed unclassified cookies is feasible and would notably enhance our analysis, we analyze their distribution, overall occurrence, and further properties. We first test the characteristics of these cookies and find that 65% of the unclassified cookies are third-party cookies and 86% of them are session cookies. A deeper analysis of these cookies shows that almost a third (37%) of the unclassified cookies have unique names. Table 7 provides an overview of the most common unclassified cookies. Overall, a manual classification of the top 5 cookies would increase the number of *classified cookies* by 2,7%, and a manual classification of the top 10 cookies would increase the number by 4%, and a classification of the top 100 cookies would increase the number by 15%. It is worth noting that 2 of the top 5 cookies are probably “functional” cookies, meaning all extensions would accept them. Therefore, we did not conduct any manual classification on these cookies as it would have a minimal effect on the overall results.

Table 7: Overview of the top unclassified cookies, their overall occurrence in our dataset, an example value, and their expected functionality based on an Internet search.

Name	Occurrence	Ex. value	Functionality
__cf_bm	1.08%	fwhRuDEX7J...	Functional cookie to detect bots.
c	0.59%	1662508350	Timestamp of unknown purpose
CMTS	0.36%	1150	Unknown
A3	0.35%	d=AQABBIK...	Unknown
li_gc	0.32%	MTswOzE2Nj...	Functional cookie to store consent preferences.