Privacy Policies and Consent Management Platforms: Growth and Users' Interactions over Time

NIKHIL JHA, Politecnico di Torino, Torino, Italy
MARTINO TREVISAN, Università degli Studi di Trieste, Trieste, Italy
MARCO MELLIA, Politecnico di Torino, Torino, Italy
DANIEL FERNANDEZ, illow, Miami, United States
RODRIGO IRARRAZAVAL, illow.io, Miami, United States

In response to growing concerns about user privacy, legislators have introduced new regulations and laws, such as the General Data Protection Regulation (GDPR) in the European Union and the California Consumer Privacy Act (CCPA), which force websites to obtain user consent before activating any personal data collection. The cornerstone of this consent-seeking process involves the use of Privacy Banners, the technical tools to collect users' approval for data collection practices. Consent management platforms (CMPs) have emerged as practical solutions to simplify the configuration and management of such privacy banners for website administrators, allowing them to outsource the complexities of managing user consent and activating advertising features.

This article presents a detailed and longitudinal analysis of the evolution of CMPs spanning 9 years. We take a twofold perspective: firstly, thanks to the HTTP Archive dataset, we provide insights into the growth, market share, and geographical spread of CMPs. Noteworthy observations include the substantial impact of the GDPR on the proliferation of CMPs in Europe, where more than 40% of websites currently adopt a CMP. Secondly, we analyse millions of user interactions with a medium-sized CMP present in thousands of websites worldwide. We observe how even small changes in the design of Privacy Banners have a critical impact on the user's giving or denying one's consent to data collection. For instance, over 60% of users do not consent when offered a simple "one-click reject-all" option. Conversely, when opting out requires more than one click, about 90% of users prefer to simply give their consent. This hints that their main objective is to eliminate the annoying privacy banner rather than make an informed decision. Curiously, we observe that iOS users exhibit a higher tendency to accept cookies compared with Android users, possibly indicating greater confidence in the privacy offered by Apple devices. We believe that the findings of this article contribute to a deeper understanding of the multifaceted interactions between privacy regulations, technological solutions and user choices in the evolving Web ecosystem. We also show that the availability of large open datasets, although not explicitly designed and collected for our goals, is fundamental to exploring different angles

This work was partially supported by project SERICS (PE00000014) under the MUR National Recovery and Resilience Plan funded by the European Union—NextGenerationEU and the project "National Center for HPC, Big Data and Quantum Computing", CN00000013 (Bando M42C—Investimento 1.4—Avviso Centri Nazionali"—D.D. n. 3138 of 16.12.2021, funded with MUR Decree n. 1031 of 17.06.2022).
Authors' Contact Information: Nikhil Jha, Politecnico di Torino, Torino, Torino, Italy; e-mail: nikhil.jha@polito.it; Martino Trevisan, Università degli Studi di Trieste, Trieste, Trieste, Italy; e-mail: martino.trevisan@dia.units.it; Marco Mellia, Politecnico di Torino, Torino, Torino, Italy; e-mail: marco.mellia@polito.it; Daniel Fernandez, illow, Miami, Florida, United States; e-mail: daniel@illow.io; Rodrigo Irarrazaval, illow.io, Miami, Florida, United States; e-mail: rodrigo@illow.io.

ACM 1559-1131/2025/08-ART30
https://doi.org/10.1145/3725737

of the internet evolution over time. For this, we make the data and code used in this work available to the community.[1]

CCS Concepts: • **Information systems** → **Web mining**; • **Human-centered computing** → **User studies;**

Additional Key Words and Phrases: Cookies, consent management platforms, web privacy, Web measurements

ACM Reference Format:
Nikhil Jha, Martino Trevisan, Marco Mellia, Daniel Fernandez, and Rodrigo Irarrazaval. 2025. Privacy Policies and Consent Management Platforms: Growth and Users' Interactions over Time. *ACM Trans. Web* 19, 3, Article 30 (August 2025), 25 pages. https://doi.org/10.1145/3725737

1 Introduction

The landscape of Web advertising has been continuously evolving since the Web is increasingly intertwined with people's lives and many websites base their revenues on advertisements. To maximize ad effectiveness, **Interest-Based Advertisement (IBA)** stems from the fundamental assumption that users are more interested in ads that relate to their interests. In the Web ecosystem, this fueled the deployment of mechanisms that aim at building specific per-user profiles to push the IBA to a great level of detail. Third-party cookies have largely been at the centre of the IBA ecosystem [7, 36, 37]. The so-called *profiling cookies* enable a third party to collect information on the visits a user performs on those websites where the third-party is present [21, 36, 37]. This allows the third party to collect information about the users' navigation history and derive what they are interested in.

Over the years, the use of cookies raised concerns for the users' privacy, and legislators started acting to regulate the abuse of these practices across the Web to increase the users' awareness of the process. In the European Union, the first significant step in this direction was the introduction of the "Cookie Law" [13] in 2009, which states that every website using first- or third-party cookies should obtain the user's approval via a *privacy banner*. The measure was further enforced in 2018 with the **General Data Protection Regulation (GDPR)** [23], which introduced fines of up to 4% on the non-compliant website and organisation global turnover. Similarly, other legislators around the globe started regulating the collection of personal data, making the presence of privacy banners pervasive, as we all experience while surfing the Web. Notable regulation includes California's **California Consumer Privacy Act (CCPA)** [12] in 2018, the Brazilian **General Personal Data Protection Law (LGPD)** [11] and the Canadian **Personal Information Protection and Electronic Documents Act (PIPEDA)** [42] in 2000 (currently being reformed [43]).

To avoid the burden of creating, controlling, and checking the compliance of privacy banners with existing and evolving regulations, websites often outsource the banner creation and management to external companies offering ***Consent Management Platforms (CMPs)***. CMPs take care of the technical complexity of implementing privacy banners, offering their customers a simple mechanism to install and manage the user's consent and data collection at large.

In this article, we present a longitudinal measurement study that aims to understand how the CMP landscape has evolved over the years and how users interact with them. We consider two angles in our study: the system and the end-user point of view.

From the system point of view, we break down the global evolution of CMPs, exploring the following research questions:

[1]Part of this work previously appeared at the 2023 Network Traffic Measurement and Analysis Conference (TMA).

— **RQ1:** How has the adoption of CMPs evolved in time?
— **RQ2:** Is there a relation between CMP adoption and the use of profiling cookies by websites?
— **RQ3:** Do different regions show different trends in CMP adoption?
— **RQ4:** How has the CMPs' market share evolved in time?

For this, we leverage the unique opportunity offered by the HTTP Archive [3], which includes the history of millions of websites since 2015.

From the end-user perspective, we leverage the preferences expressed by users when interacting with a medium-sized CMP, formulating the following research questions:

— **RQ5:** Does the presence (or absence) of a "one-click-choice" option affect the users' rejection rates?
— **RQ6:** Do the employed browser, device, operating system, banner position, or other factors impact the user's choice?

To this end, we leverage a unique dataset of the illow.io[2] CMP, which spans more than 2 years of data since its introduction in the CMP market, collecting more than 20 million interactions with the CMP.

Our study unveils the introduction of the GDPR and similar regulations that fuelled the rapid growth of the CMPs, now present in 40% of websites in Europe. We also observe an increased number of websites that opted to directly remove third-party tracking technologies, probably balancing the benefit and complexity of managing their presence. Surprisingly, more than 50% of websites adopting a CMP still install third-party cookies of possible tracking services before the user accepts their usage. This is, in fact, due to the complexity of managing third-party content and the lack of standard technical solutions to enforce the user's choice by CMPs. Moreover, a website has to adapt its interaction with third parties depending on regulations and the country of the user. This, in turn, calls for more analysis and more extensive archival projects.

From the user's viewpoint, we observe that more than 60% of users deny the usage of third-party tracking cookies when offered a simple "one-click reject-all" action. Yet, the large majority of these choices derive from the intent of quickly removing the intrusive banner from the screen rather than an explicit and conscious decision. In fact, in those countries where the opt-out choice requires more than one click, about 90% of users accept the usage of cookies via a one-click accept-all button. Curiously, iOS users tend to accept the usage of cookies more frequently than Android users. This may be linked to a higher confidence in privacy offered by iOS and Apple devices.

Both the HTTP Archive and the illow.io datasets have not been collected or designed specifically for this work. Yet, their availability allows us to study and present interesting insights in an opportunistic manner. In particular, the HTTP Archive collects its data uniquely from locations in the United States, which is nowadays limiting, as the web is increasingly customized based on users' location to adapt to the different privacy regulations. To extend and allow the reproducibility of our work, we publish online the preprocessed HTTP Archive dataset and the data used for the CMP analysis along with the code to obtain all of the figures in this article [4].

The remainder of the article is organized as follows. Section 2 introduces the regulation history and discusses related works. Section 3 describes the datasets we use for our analyses and the methodology we use to evaluate the results — both for CMP adoption and user interaction. In Section 4, we present how the pervasiveness of the CMPs in the Web has evolved in time (**RQ1**), their relationships with profiling cookies (**RQ2**), the differences among regions (**RQ3**), and the evolution of the market share (**RQ4**). Section 5 focuses on users' interactions when facing the privacy banners of the illow.io CMP, exploring the impact of the "one-click-choice" button (**RQ5**)

[2]https://illow.io, accessed on Tuesday 18th March, 2025.

and that of several other factors (**RQ6**). Finally, Section 6 summarizes the main findings of our work.

2 Privacy Regulations and Related Works

Behavioural advertising has always been a pillar of the Web ecosystem and entailed the massive collection of personal information through web tracking and cookies [7, 21, 22, 36, 37, 41, 44, 46, 49]. The implications of web tracking on users' privacy encouraged legislators to issue privacy-related regulations.

In the European Union, the introduction of the Cookie Law [13] in 2009 and updated in 2013 led to the proliferation of privacy banners [20]. The regulation states that, when visiting a website, users have to interact with the privacy banner to explicitly opt in to the usage of tracking mechanisms. A website (and any third party embedded in the website) is allowed to install cookies and start data collection only after getting users' explicit consent. Privacy banners, however, do not fully protect users in many cases [52]. Later in 2018, the GDPR profoundly influenced the Internet user experience [15, 18, 32, 34, 47] for European Union–based users by defining severe sanctions for violators.[3]

Other countries issued similar regulations, with notable examples in the Brazilian LGPD [11] (going into effect on September 18, 2020), the PIPEDA [42] (April 13, 2000), and the Quebec 25 law [38] (September 22, 2021) in Canada, along with the California Consumer Privacy Act (CCPA) [12] (January 1, 2020) which spurred other US states to enforce similar regulations. At the high level, these regulations provide individuals control over the personal data that businesses collect about them. In a nutshell, they mandate that users have to explicitly opt in to the collection of personal data, including the usage of cookies and other tracking technologies on the Web.

A large number of websites started to rely on CMPs, i.e., external companies that provide technical solutions to manage users' consent collection by offering customizable and easy-to-deploy privacy banners. Hils et al. [28] provided a first analysis of CMPs' popularity. By actively crawling popular websites for almost 3 years, they provide a detailed picture of the growth in popularity of CMPs, observing websites switching CMPs, and manually analysing the privacy policies they publish to check the purpose for data collection indicated. Our work complements and extends this analysis by offering a longitudinal analysis over a period of 9 years thanks to the HTTP Archive, and by including the users' perspective on the analysis.

In general, it has been shown that most users go from the tendency to ignore privacy-related notices [14, 24, 55] up to getting annoyed by them. This behaviour has gone under the name of the "privacy paradox": users claim to be concerned about their privacy while at the same time taking little action to protect their data [9]. In fact, given the advertising-based business model, some websites and CMPs make efforts to increase the cookie acceptance rate. Recent works have shown that banners often nudge users to acceptance by exploiting dark patterns in the user interface, if not openly disregarding the GDPR's requirements [35], or making it difficult for users to exercise their rights [26]. Some hide the content of a website, which is visible only upon cookie acceptance. These tactics impact automated Internet measurements since the first visit to a website may not show the actual website content [31]. Nudging includes offering the user an `Accept All` default button via intrusive banners [10, 19], which is often the case [27] with websites presenting large pop-ups or wall-style banners that cover most of the webpage content. Researchers have shown that apparently minor design choices have a significant effect on inducing the user to accept the cookies [33, 39, 45, 48, 50, 54]. For instance, Habib et al. [25] compared the behaviour of 1 109 volunteers on

[3]The United Kingdom adopted the GDPR in the "Data Protection Act" in 2018. In the rest of the article, we refer to "GDPR countries" as any European country where the GDPR is in place, including the United Kingdom.

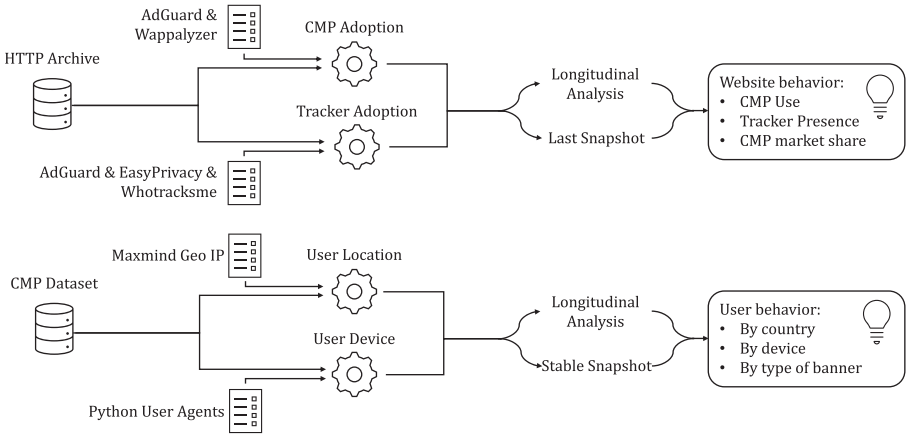

Fig. 1. Data processing and analysis workflow. The top part refers to the CMP adoption study. The bottom part refers to the user behaviour study.

12 different banners, showing that the cookie acceptance rate varies by ≈ 25% depending on banner design. Unequal paths (i.e., *refuse* action more complicated than the *accept* action) and blocking banners are the most important factors. Our work complements this by analysing the behaviour of millions of users when facing the banner of a popular CMP whose banner assumes various forms across time and countries. This allows us to compare the impact of the banner design, including many more factors than those considered by Habib et al. [25].

The part of the article focused on the study of user interaction with the CMP is an extension and refinement of a previous work [30]. In this work, we extend the timeline of our analysis to include more than 2 years. This allows us to better appreciate the impact of the banner design on the user choice and to consider several new angles, including the impact of the user's device, country, etc.

3 Methodology and Datasets

This section describes our datasets and the processing methodology we adopt to carry out our analyses. We rely on two datasets to study CMP adoption and user behaviour on CMPs that we enrich using external sources. We provide a high-level overview of our data processing and analysis methodology in Figure 1. The top part refers to the CMP adoption study, which we base on the HTTP Archive. The bottom part refers to the user behaviour study, which we base on the CMP dataset. In both cases, we start from a longitudinal dataset and integrate it with external information to observe both the evolution over time of different metrics and the detailed breakdown in the most recent period. In the appendix, we provide the exact structure of the two datasets. We release the datasets used in our analyses — both for the CMP adoption study and the user behaviour one — along with the code to replicate our results as open source [4].

3.1 CMP Adoption

3.1.1 Dataset. To study CMP adoption through the years, we use the HTTP Archive open dataset [3]. The curators of this archive visit a list of worldwide websites using an automated version of the Google Chrome browser. They perform the visits using test agents from various Google Cloud data centers in the United States. The test agents use the WebPageTest[4] to instrument

[4]https://webpagetest.org/ toolset, accessed on Tuesday 18th March, 2025.

Table 1. Breakdown of Websites Present in the HTTP Archive Dataset

	Dataset		Websites	Period
HTTP-9Years	**Total**		**47,899**	Jan 2015 - Dec 2023
	GDPR		**11,819**	
		.de	2,907	
		.es	603	
		.fr	1,060	
		.it	1,060	
		.co.uk	1,617	
	EU No-GDPR		1,412	
	Africa		336	
	Asia		3,926	
	North America		777	
	South America		1,399	
	Other		28,230	
	HTTP-Dec23		1,000,000	Dec 2023

Google Chrome and collect several metrics on the webpage visit using the HTTP Archive (HAR) format, a JSON-formatted archive format for logging a web browser's interaction with a webpage. All websites are visited every 15 days to create a longitudinal archive. The curators update the list of websites they visit at regular intervals. Until June 2018, the website list contained the top-500 k entries of the Alexa rank. In July 2018, they extended the list to include the top-1M websites taken from the Chrome UX Report, which contains a variable set of entries ($\approx$ 10 M), updated on a monthly basis.[5] The HTTP Archive data is available on Google Cloud, from which we downloaded the data into a Spark cluster for processing.

In this work, we rely on the monthly snapshots offered by the HTTP Archive starting from January 2015 until December 2023. For the longitudinal study, we focus on the subset of websites that are present the entire time. There are 47,899 websites. We refer to this dataset as HTTP-9Years. To study the impact of the countries the website belongs to, we use the **top-level domain (TLD)**, e.g., `.fr`, to associate a website with its main country.[6] We have 11,819 websites whose TLD belongs to a GDPR-regulated country and that are present for all 9 years. We focus most of our analysis on this set of websites. Table 1 details the breakdown of sites for each country.

To complement the longitudinal study and detail the most recent picture of the CMP adoption, we resort to the last available snapshot, which includes the top-1M websites as visited during December 2023. This dataset includes websites from all over the world. We refer to this dataset as HTTP-Dec23 (last line in Table 1).

For each visit, the HTTP Archive datasets report several statistics and details on each HTTP transaction carried out to fetch the webpage's objects. For each HTTP request/response pair, the dataset reports the URL, the timings, and various HTTP headers, including Cookies, Referrer, etc. This allows us to study the presence of various third-party elements in a website, including the

[5] https://developer.chrome.com/docs/crux, accessed on Tuesday 18th March, 2025.

[6] Although there exist other approaches to link a website to its country (e.g., inspecting the registrar of the DNS domain or location of the server), using the country-code TLDs represents an accurate yet conservative option even if we may erroneously exclude websites using global TLDs (`.com`, `.org`, etc.).

presence of CMPs or other third-party services. To this end, we examine all HTTP requests issued during the page loading and extract the domains from the corresponding URL. In the case in which the HTTP request refers to a third-party server (i.e., the domain is different from the webpage one), we match the domain into the lists described next to detect whether the website adopts a CMP or embeds a potential Web tracker.

3.1.2 Processing Methodology.

CMP Identification. We study the adoption of CMPs by looking for URL requests to the most popular CMPs. For this, we build a list of CMP domains by combining two sources: the AdGuard [1] list, which includes a set of CMP domain names; and the Wappalyzer [5] list, a profiler tool that is used to identify various characteristics of a website, including the presence of CMPs and analytics products. Wappalyzer explicitly lists domain names of CMPs. We merge the two sets of CMP domains and, given the small number of entries, we manually verify each one. In total, we obtain 84 CMPs, each identified by one or more domains.

Potential Trackers. To detect the presence of potential third-party trackers in a given website, we merge publicly available lists provided by Whotracksme [6] (an anti-tracking-related open-data provider), EasyPrivacy [2] (one of the lists at the core of AdBlock tracker-blocking strategy) and AdGuard [1] (a popular ad-blocking tool). Trackers are identified by their domain name. For robustness, we merge the three lists and consider as a potential tracker any third-party domain that appears in at least two lists. In total, we obtain 1,497 domains that we consider tracking services.[7]

We record the presence of a tracker during a visit if the webpage embeds an object served by a tracking domain and the latter installs a persistent cookie. The HTTP Archive dataset does not indicate the content and the lifespan of the cookie, making it impossible to verify whether the third-party cookie is actually a *profiling cookie* or not. We thus verify the *potential* presence of a tracker. In fact, such a cookie could be a technical cookie used to store, e.g., consent information, not violating the privacy policy or the regulations.

3.2 User Interaction

3.2.1 Dataset. To study the users' interaction patterns when facing a privacy banner, we use a proprietary dataset from illow.io, a medium-sized CMP company. The illow.io CMP provides web developers with the ability to include a simple privacy banner to control their preferences about data collection, via cookies or other means. The banner is shown the first time a user accesses the website as a small overlay window that can be placed in different locations on the screen. The users can express their preference across four categories of cookies: (i) necessary, (ii) statistical, (iii) preferential, and (iv) marketing cookies. The necessary cookies include those technical cookies that are needed for the site to function, including the storage of the user's preference on the usage of cookies, and the user cannot opt out from their usage. Once the users have expressed their choice, the CMP sets the necessary cookies on the users' device to store their preferences. When the user accesses the website again, no banner is shown and preferences can be updated via a dedicated page.

The privacy banner is served directly by the illow.io servers in the form of a small set of Java scripts. When users interact with the privacy banner (i.e., provide their choice, or change their consent), the illow.io servers log the event. The collection happens only when users submit their preferences (as detailed in the privacy notice). No data is collected if the user does not perform any action on the banner. The logged information includes the time of the visit, a random ID of the

[7]For simplicity, we match the domains using the *second-level domain name* — i.e., a host name truncated after the second label. We handle the case of two-label country-code TLDs, such as `.co.uk`.

entry, an anonymized ID for the website, the type and position of the banner, which cookies the user accepted, the `User Agent` header value as set by the browser, and the client IPv4 /24 subnet. We use the client's subnet to map a user to country and state via the MaxMind GeoIP[8] database.[9] This information is necessary to implement the functionalities of the platform (i.e., record user's preferences for the next visits to the website), and it is useful to customize the information provided to users (e.g., show the banner in a different language, decide what version of the banner to show to the user — more details in Section 3.2.2), and to collect statistics about the usage of the platform, including the billing of the website deploying the CMP. All of these pieces of information, including the description of the collected data and the purpose of the collection, are documented in the privacy policy the CMP offers to users.

Users submitting (or changing) their preference generate an entry in the log, which we call an *interaction* in the following. As previously stated, each entry is associated with a random ID. Thus, it is impossible to re-identify or track a user across different websites, guaranteeing users' privacy. To further protect the CMP customers, in the data used for this study, the website name is also anonymized by replacing the domain name with a random identifier.

We classify each interaction according to the combination of cookie categories accepted by the user, as follows.

— *Accepted-All*: If all cookie categories are accepted, either with a single click on the `Accept All` button, by individually accepting all of the cookies after clicking on the `Custom Permissions` button, or by clicking on the *Don't sell my data* link.
— *Rejected-All*: If only the necessary cookies are accepted, either by clicking `Reject All` button (or equivalent action) if present or by manually deactivating all cookies after clicking on `Custom Permissions` (recall that necessary cookies cannot be excluded).
— *Custom*: If at least one among the statistical, preferential, and marketing cookies is accepted through the `Custom Permissions` screen.

Summing up, the dataset contains the following attributes.

— `siteId`: An integer representing the website where the click happened
— `createdAtDate`: The day when the click has happened
— `country`: The country of the user who takes the choice
— `region`: The geographical macro-region of the user
— `os`: The operating system of the user
— `browser`: The browser employed by the user
— `position`: The position of the banner on the user's screen
— `banner`: The banner version with which the user interacts
— `blurred`: Whether the background was blurred when the banner appeared
— `status`: The choice of the user.

In the following, we compare the *Rejected-All* rate against different factors, such as banner layout and position, user's country, and type of device. This offers us precious insight into understanding which banner layout is the most effective in allowing the users to exercise their right to refusal.

The dataset contains interactions spanning from July 1, 2022 to August 22, 2023. In this period, the CMP collected $\approx 25M$ interactions from more than 6,000 websites. We report the breakdown of interactions per users' regions in Table 2. Also for this dataset, we consider the whole data when focusing on the evolution over time, while we consider the period from July 5, 2023 until the end to consider a stable period and minimize the effect of the transients.

[8] https://www.maxmind.com/, accessed on Tuesday 18th March, 2025.
[9] We do not consider IPv6, as it generates negligible traffic.

Table 2. Number of Users' Interactions Per Geographical Region - illow.io Dataset

Region	Total	July-August '23
South America	11,953,295 (48.20%)	3,415,701 (43.79%)
GDPR-regulated	9,417,253 (37.97%)	3,124,085 (40.06%)
North America	2,448,467 (9.87%)	849,008 (10.89%)
Asia	377,853 (1.52%)	184,390 (2.36%)
Rest of Europe	367,107 (1.48%)	140,969 (1.81%)
Africa	149,905 (0.60%)	45,925 (0.59%)
Oceania	87,685 (0.35%)	38,757 (0.50%)

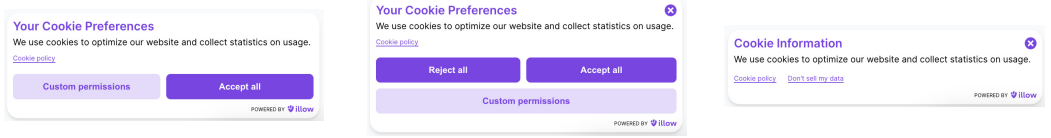

(a) General banner. (b) GDPR-compliant banner. (c) CCPA-compliant banner.

Fig. 2. The available banners introduced since July 2023 in the illow.io CMP. Figures showcase the pop-up style banners.

Banner Layouts. The different regulations specify how online companies should collect data on the visitors (from areas under their jurisdiction) and indicate how the user's consent should be obtained. A comparison of the three most widely adopted privacy regulations (GDPR, CCPA, and LGPD) can be found in recent works [8, 17]. When a new user visits a website, the CMP serves the correct type of banner according to the country and region the customer visit comes from.

Over the years, the illow.io CMP updated several times the design and actions offered by the banner to enable new functionalities or to adapt to evolving regulations. In our longitudinal study, we show the impact of such changes. The description below refers to the most recent configuration of the CMP, which has been in place since July 5, 2023. The illow.io CMP supports three different types of banners that comply with different regulations:

— The *General* banner (Figure 2(a)) presents two buttons, the `Accept All` and the `Custom Permissions`. The first button gives the permission to install any type of cookie – i.e., represent the opt-in choice. The second button brings the user to a selection window, where the user can fine-select the categories of cookies to accept. This is the default banner shown to users who connect from areas that are not covered by any particular regulation.

— The *GDPR-compliant* banner (Figure 2(b)) presents three buttons. The `Accept All` gives the permission to install any type of cookie. The `Reject All` only forbids all but necessary cookies. Finally, the `Custom Permissions` prompts to the selection window, as in the *General* banner. The banner also features a close-window button in the form of an "X". The CMP offers this banner to users connecting from GDPR-regulated countries and from countries with similar regulations – i.e., Canada (due to the PIPEDA and Quebec 25 laws) and Brazil (due to the LGPD). This banner was introduced in August 2022, whereas before, users in these countries faced the *General* banner. Moreover, the operation of the "X" close button changed from "accept-all" to "reject-all" action in December 2022. We will explore these changes in Section 5.

— The *CCPA-compliant* banner (Figure 2(c)) notifies the user that cookies are being used. Two links appear: the first brings the user to the website privacy policy. The second, named "Don't

sell my data", brings the user to the selection window, where the user can fine-select the active cookies. This banner is served by default to visitors coming from California and Utah, under the regulation of the CCPA and the **Utah Consumer Privacy Act (UCPA)**, respectively. The banner offers a close-window button in the form of an "X" that closes the banner without requiring further action from the user, which implicitly accepts the use of cookies. Note that the operation of this banner is different from the previous two: the user must explicitly opt out of cookies, which are enabled by default.

3.2.2 Processing Methodology. The dataset is securely stored on a big data cluster located in our institution. We process it using Python code to extract various metrics and statistics. Given that an imbalance exists in the website audience, we want to prevent popular websites from biasing the results. For this, we opt to show results using a website-wise macro-average of the metrics under study. In other words, we compute the desired metric average for each website separately. Then, we compute the average over the websites. This way, each website weights one, regardless of the number of interactions it generates. We consider a per-website average valid if we observe at least 10 interactions from it in at least 50 websites.

Formally, given a target metric M (e.g., *Rejected-All*), a set of websites $w \in \mathcal{W}$, each with a population of interactions $\mathcal{I}_w$, an indicator function $\mathbb{1}_M(i)$ which returns 1 if i refers to M, 0 otherwise (e.g., whether interaction i records a *Rejected-All* choice or not), we define as $\overline{M}(\mathcal{I})$ the website-wise macro-average of M computed over the samples belonging to the subset $\mathcal{I} = \bigcup_{w \in W : |I_W| \geq 10} \mathcal{I}_w$. Formally,

$$\overline{M}(\mathcal{I}_w) = \frac{1}{|\mathcal{I}_w|} \sum_{i \in \mathcal{I}_w} (\mathbb{1}_M(i)), \text{ given } |I_W| \geq 10, \tag{1}$$

$$\overline{M}(\mathcal{I}) = \frac{1}{|\mathcal{W}|} \sum_{w \in \mathcal{W}} \overline{M}(\mathcal{I}_w), \text{ given } |\mathcal{W}| \geq 50. \tag{2}$$

Equation (1) computes the per website average. Equation (2) computes the overall macro average.

We also evaluate the confidence interval of the macro average. Hence, each estimate is presented as

$$\overline{M}(\mathcal{I}) \pm c \cdot \frac{\overline{S}(\mathcal{I})}{|W|},$$

where c corresponds to the quantile of a Student's t-distribution with $|W| - 1$ degrees of freedom and $\overline{S}(\mathcal{I})$ is the sample standard deviation of each website-wise average. In this work, we consider a confidence interval of 90% and report the confidence interval as an error bar. As our main target metric, we consider the *Rejected-All* rate.

4 Temporal Evolution of CMP Adoption

In this section, we study the adoption of CMPs on websites during the last 9 years. To this end, we rely on the two HTTP Archive datasets (HTTP-9Years and HTTP-Dec23) as described in Section 3.

4.1 Overall Trend

We start our analysis with **RQ1** — i.e., the evolution of CMP adoption — with Figure 3, in which we show the fraction of websites with and without a CMP on the 11,819 European-based websites for which GDPR is in force. We further break down the statistics to consider sites that embed or do not embed a potential tracker, addressing **RQ2**. Observe how the CMP adoption (indicated by the black curve) was negligible before the GDPR (which went into effect in May 2018, indicated with a *black vertical line* in the figure): at that time, only $\approx 5\%$ of websites adopted a CMP. Recall that some form of privacy banners was already implemented to comply with the previous "Cookie

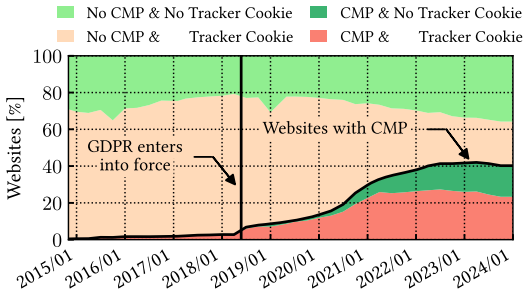

Fig. 3. Fraction of websites with/without a CMP and contacting potential trackers. HTTP-9Years dataset of 11,819 websites that are present during the whole period in GDPR-regulated countries.

Table 3. Fraction of Websites With/Without a CMP and Contacting Potential Trackers

	Tracker	No Tracker	Overall
CMP	9.9% (75.6%)	3.2% (24.4%)	13.1%
No CMP	40.2% (46.2%)	46.7% (58.8%)	86.9%
Overall	50.1%	49.9%	

HTTP-Dec23 Dataset with top-1M website worldwide.

Law" European directive. Yet, CMPs were not a common solution. Note that between 2015 and 2018, at least 3 out of 4 websites used to embed at least a potential tracker even if users did not accept their usage. This is in line with our previous finding [52]. In fact, this figure is a lower bound for the popularity of tracking services because we expect that more trackers would be contacted upon the user's consent (recall that the HTTP Archive test agents always perform a "first visit" to all websites and do not interact in any form with any banner).

Starting from May 2018, the CMP adoption quickly increases, reaching 30% at the beginning of 2021 and 42% in late 2023. As expected, the CMP adoption primarily erodes the fraction of websites with no CMP and potential tracker (*light red area*) which shrinks from ≈70% (May 2019) to ≈20% (December 2023). A possible explanation for this surge is that the introduction of the GDPR pushed those websites using ads and tracking services to adopt solid privacy management solutions (i.e., a CMP).

It is interesting to note that the fraction of websites without a CMP *and* without potential trackers (*light green area*) increases as well, growing from 20% (May 2019) to 35% of websites (December 2023). We argue that this is due to two phenomena: first, some large and popular websites started implementing proprietary consent management solutions that we cannot detect — we leave this as future work. Second, smaller Web services started balancing the benefit of including third-party ad services with the burden of correctly managing their presence, opting out from including trackers and deploying a CMP.

Focus now on the fraction of websites with CMPs (*bottom solid red and green areas*). Unexpectedly, more than 60% of websites that adopt a CMP still let the user's browser install cookies from potential trackers before collecting any user's consent (*bottom dark red area*). We argue that this is due to the complexity of managing advertisement and tracking platforms, coupled with the burden of installing and operating a CMP and the lack of a standard mechanism for CMPs to control third-party services. In a nutshell, as previously found [52], the presence of a consent banner does not guarantee the correct management of data collection policies. Recall that we cannot claim that those websites violate the GPDR: first, the HTTP Archive visits websites from the United States and a website/CMP might differentiate their behaviour when the user visit comes from a non-GDPR country. Second, a tracking service can legitimately install a technical cookie before obtaining the user's consent. As previously said, we cannot distinguish this case from the installation of an actual profiling cookie owing to limitations of the HTTP Archive data.

We complement our analysis by looking at the most recent snapshot, which contains the top 1 million ranked websites. Here, we broaden the picture to include websites from all over the world

Table 4. Potential Tracker and Third Party Prevalence on Websites Adopting Different CMPs

CMP	Websites With Trackers [%]	Average Trackers	Websites
Usercentrics	32.94	2.14	16,077
Iubenda	56.30	3.05	13,378
OneTrust	61.33	9.23	46,144
CookieScript	63.55	4.78	2,093
Cookiebot	75.85	4.99	12,415
Didomi	82.14	20.29	5,202
Funding Choices	97.85	16.30	26,921
HubSpot	99.20	9.07	20,871

HTTP-Dec23 Dataset with top-1M website worldwide.

regardless of their country and of being present in past snapshots. In Table 3, we show the CMP adoption and potential tracker share. Globally, only 13.1% of websites adopts a CMP, whereas 50.1% include potential trackers. As noted before, most websites adopting a CMP embed potential trackers (9.8% of the total, corresponding to 75% of sites with a CMP).

We investigate in more detail the surprisingly large fraction of websites with a CMP that include a potential tracker. For the top-10 most popular CMPs (discussed later), in Table 4, we report the percentage of websites that embed a potential tracker (sorted from lowest to highest), their average number, and the number of websites they are present in. Results are very heterogeneous, with some CMPs able to prevent the browser from contacting potential trackers more efficiently than others. For instance, 32.9% of websites adopting Usercentric still contact two potential trackers on average. Conversely, 58.8% of sites adopting the Sourcepoint CMP still allow the visitors to contact more than 33 potential trackers. Again, due to the limitations of the HTTP Archive, we cannot confirm that the cookies installed by potential trackers are profiling cookies. Recalling that the HTTP Archive visits come from the United States, we manually verified 10 websites adopting the Usercentric and 10 websites adopting the Sourcepoint CMP, visiting them from Italy. In all cases, we observe that both CMPs correctly permit the browser to contact trackers only after accepting the site privacy policy. This calls for further investigations and automated means to provide an extensive evaluation of the eventual violation of the privacy policies, as this would depend on multiple factors such as client country, language, device type, etc., as presented in our previous work [31]. We leave this as future work.

4.2 Country-wise and Region-wise Diversity

We now move to **RQ3**: different countries and regions of the world have adopted different regulations at different times. To observe how this impacts the evolution of CMPs, in Figure 4 we break down the CMP adoption for GDPR countries. Here, we focus on websites belonging to 5 country code top-level domains: `.de`, `.es`, `.fr`, `it`, and `.co.uk`, i.e., websites based in Germany, Spain, France, Italy, and the United Kingdom, respectively. Before GDPR (represented by the *black vertical line*), CMP adoption was negligible in all countries except for Italy, where approximately 10% of websites had already adopted a CMP starting from mid-2015. This is likely due to the implementation of the old European Directive called the Cookie Law [13], which went into effect in June 2015 in Italy. Indeed, the Italian Data Protection Authority "Garante per la protezione dei dati personali" mandated the correct implementation of a privacy banner [29] by June 2, 2015. As a consequence, Italian websites resorted to CMPs to solve the issue. The Italian CMP Iubenda grabbed a market share of 90.3% in Italy by the end of 2017. In the last part of this section, we analyse in detail the market share of popular CMPs.

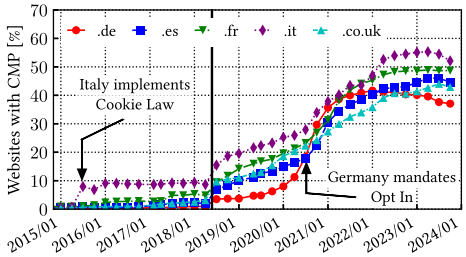

Fig. 4. Fraction of websites with a CMP for 5 European TLDs. HTTP-9Years dataset of 11,819 websites that are present during the whole period, in GDPR-regulated countries.

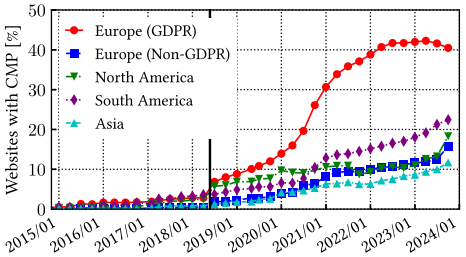

Fig. 5. Fraction of websites with a CMP for different continents. HTTP-9Years dataset of websites that are present during the whole period, in countries not regulated by the GDPR.

Following the initiation of the GPDR (May 2018), CMP adoption started increasing gradually until 2023, even if at a different pace. Italian websites (*purple line*) grow quicker, having the highest rate until the most recent snapshot. Conversely, the German websites (*red line*) have the lowest CMP adoption rate, reaching just ≈ 10% in mid-2020. From June 2020, they exhibited a very rapid increase, reaching 35% by the end of the year. This acceleration is due to the decision of the German Federal Court [16] that in May 2020 modified the GDPR in a more restrictive sense. The judgement underlined the need for *explicit* user consent before installing marketing or nonessential cookies and specified that an opt-out mechanism is invalid. We speculate that this led to the massive adoption of CMPs in Germany, causing their sudden growth to 40% by mid-2021. Spain, France, and the United Kingdom exhibit a more gradual increase. The CMP adoption was below 10% in May 2019 and increased to 40% to 50% by the end of the observation period.

Interestingly, we observe that, in all countries but the United Kingdom, CMP adoption was no longer increasing in 2023, and we even notice a moderate yet measurable decrease in Italy and Germany.

For the sake of comparison, in Figure 5 we dissect CMP adoption for different regions of the world. The *red solid line* refers to GDPR countries (i.e., it is equivalent to the back curve of Figure 3). In these countries, CMP adoption tops 42%, while we observe a moderately decreasing trend in the second quarter of 2023. In the remaining European countries (those not implementing the GDPR, *blue line*), the CMP adoption rate stands at significantly lower values and, even with an increasing trend, it has not yet reached 20%. We observe similar numbers for North America and Asia— green and cyan curves, respectively. Interestingly, in South America, the CMP adoption rate is consistently higher over time, exceeding 20% since mid-2023. This is mostly due to Brazilian (`.br`) websites, which represent 59% of the 1,399 South American websites in the HTTP-9Years dataset. In Brazil, the LGPD went into effect in September 2020, imposing a regulation similar to the GDPR. We consequently observe an average CMP adoption rate in South America of 26% by the end of 2023.

In short, we clearly observe that the initiation of and important changes to the regulations appear to stimulate the need for valid technical solutions, potentially creating a new market need that CMPs satisfy.

4.3 The CMP Market Ecosystem

We now study this new CMP market, discussing the pace at which they are born and how their market share evolves over time, to address **RQ4**. We focus on the GDPR-based websites present in the HTTP-9Years dataset. We first study when new CMPs entered this market over the years. In Figure 6, the x-axis represents time; On the y-axis, we sort the CMPs by the time of their first appearance. A dot in the figure indicates that the i-th CMP was found on at least 10 different

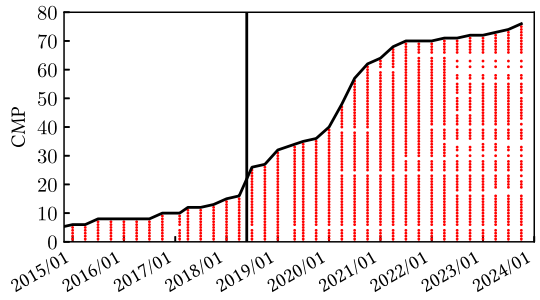

Fig. 6. Presence of the considered CMPs in the observation period. HTTP-9Years dataset of 11,819 websites that are present during the whole period, in GDPR-regulated countries.

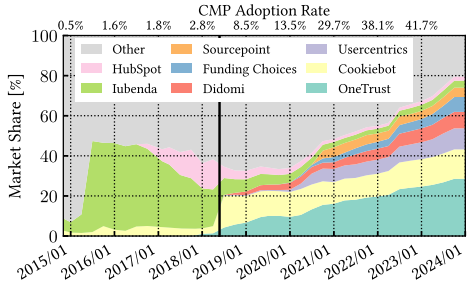

Fig. 7. Market share of the top-8 CMPs. HTTP-9Years dataset of 11,819 websites that are present during the whole period in GDPR-regulated countries.

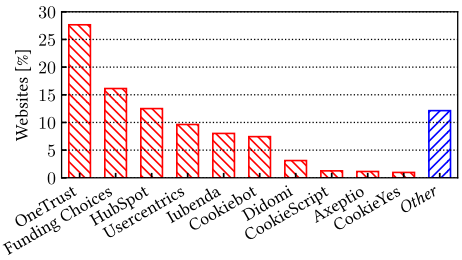

Fig. 8. Market share of the top-10 CMPs on the last snapshot. HTTP-Dec23 dataset with top-1M websites worldwide.

websites in that month. In fact, out of 84 CMPs present in our list, 76 have actually been adopted by more than 10 websites (we expect some to be active uniquely in countries other than those enforcing the GDPR). Figure 6 shows that, before the GDPR, less than 20 CMPs were active on the market. The GDPR fostered the birth of new solutions; by the end of 2019, 35 CMPs were active. In 2020, we observe a second increase in the number of CMPs, with the adoption of 30 new platforms, a symptom of a more vast, yet competitive market for privacy-management solutions. In 2021, 2022, and 2023, we observe less than 10 newcomers, hinting that the European market is rather saturated. Finally, we notice that a non-negligible fraction of CMPs (18 out of 76) has apparently shut down or has been acquired by other companies (i.e., is no longer present, as indicated by the missing dots).

We now focus on CMP market share. For each CMP, we compute its market share as the ratio between the number of websites adopting it and the number of websites using any CMP. In Figure 7, we report the market share evolution in the GDPR countries of the top-8 CMPs (as measured in the last time step of the dataset). The grey area represents the remaining CMPs. For reference, the top x-axis reports the overall CMP adoption rate (recall that it increases over time). Before the GDPR went into effect, Iubenda dominated the market: it was the choice for more than 40% of (mostly Italian) websites resorting to a CMP. The main competitor of Iubenda was Hubspot, even if with moderate penetration. After GDPR, the leading position of Iubenda was quickly eroded by new competitors: Cookiebot and OneTrust entered the market, the former reaching a market share of 16.1% by the end of 2020. Since then, Cookiebot maintained its market share of about 25% to 28%. OneTrust, a company offering various marketing and analytics solutions, exhibited a significant

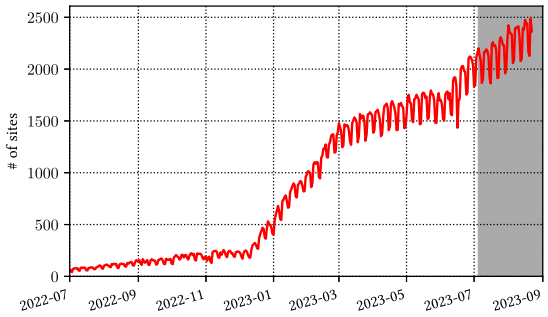

Fig. 9. Evolution of the number of sites with at least 1 daily recorded interaction in the illow.io dataset. The grey area represents the period to which the stationary analysis of Section 5.2 refers.

growth that reached 29% at the end of 2023. This is also thanks to the acquisition of the CookieLaw and CookiePro competitors. The youngest notable player is Funding Choices, a CMP operated by Google, which appeared in 2020 and currently has a market share of around 7.4%.

We complete the analysis by providing the most recent snapshot of the worldwide CMP market share. We rely on the top-ranked 1 million websites as of December 2023 in the HTTP-Dec23 dataset, in which 13.1% adopts a CMP (see Table 3). Figure 8 details the market share of the top 10 CMPs worldwide. We find most of the CMPs already presented in Figure 7, even if with a different rank and share owing to the different website bases. The leader is again OneTrust (27.6% of market share). Funding Choices by Google comes second, with 16.1% of market share among the top-1M worldwide websites. Interestingly, its market share is only 7.4% of European websites existing in all 9 years. This suggests that Funding Choices CMP is especially adopted in non-GDPR countries. Conversely, CookieBot has a market share of 14.7% on the HTTP-9Years (GDPR) and only 7.3% on HTTP-Dec23, suggesting that it is more popular among GDPR countries.

To summarize, CMPs are products born as a result of European privacy regulations. Different countries show different patterns caused by their internal dynamics and specific implementations of the European directives and privacy regulations at large. While the market was initially dominated by a few players, nowadays different companies compete. The availability of longitudinal data such as those provided by the HTTP Archive is fundamental for this type of analysis.

5 User Interaction

In this section, we focus on users' interaction patterns when facing a privacy banner. To this end, we use the proprietary dataset of the illow.io CMP described in Section 3.2.

In total, illow.io serves more than 6,000 websites all over the world, collecting hundreds of thousands of interactions daily. In Figure 9, we show the daily number of websites for which we observe at least one interaction. The weekly periodicity is owing to Sundays, which record the least amount of traffic. We first observe that the number of the CMP's customer websites grows over time, with a significant increase starting in December 2022 and a second increase in June 2023. Since then, more than 2,000 websites have been active daily. Not shown for brevity, the number of daily interactions grows proportionally, surpassing 180,000 daily interactions in the last period.

5.1 Longitudinal Analysis

We first focus on the evolution of the fraction of *Rejected-All* interactions. Figure 10 reports the daily average *Rejected-All* rate for users connecting from different areas of the world, separately for the GDPR area (Figure 10(a)) and other notable countries (Figure 10(b)). The curve profiles

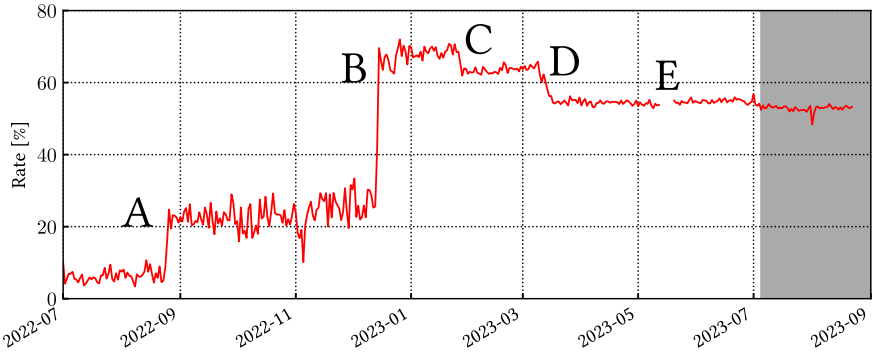

(a) GDPR-regulated countries.

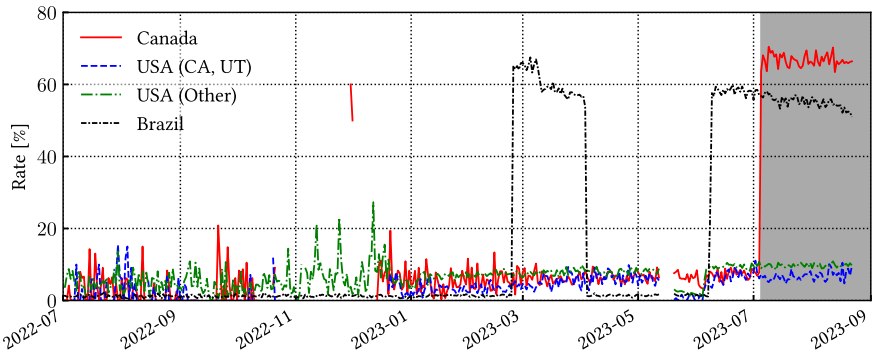

(b) Other areas.

Fig. 10. The *Rejected-All* rate in a 14-month period by macro areas. The grey areas represent the period to which the stationary analysis of Section 5.2 refers.

provide insights on **RQ5**, i.e., how the presence of a "one-click choice" affects the user's behaviour. In contrast to what we define in Section 3.2.2, we remove the 50-website constraint so that we have enough data for most of the days in our 14-month span. Figure 10 thus shows the average among all the websites with at least 10 daily interactions.

Starting from the GDPR area, observe how the curve follows a complicated shape. Each change is a consequence of changes in the banner design at illow.io. We highlight and describe the five main changes:

— **A – Add `Reject All` button**: In August 2022, the CMP changed the banner shown to the GDPR users from the *General* – two-button banner – to the *GDPR-compliant* – three-button banner. The new banner includes the `Reject All` button, allowing users to refuse the usage of non-mandatory cookies with a simple one-click action. This addition is a consequence of the fine imposed by **the Commission Nationale de l'Informatique et des Libertés (CNIL),** the French data protection authority, on Google and Facebook in January 2022.[10] This suddenly increases the *Rejected-All* rate to ≈ 20% (we already observed this behaviour in our previous work [30]).

[10]https://www.cnil.fr/sites/cnil/files/atoms/files/deliberation_of_the_restricted_committee_no._san-2021-023_of_31_december_2021_concerning_google_llc_and_google_ireland_limited.pdf, accessed on Tuesday 18th March, 2025.

— **B – Add "X" button, activate blurring, top position**: In December 2022, the CMP introduced the "Close banner" "X" button. Clicks on this button are equivalent to a *Rejected-All* choice. Following again the suggestions of European privacy-regulation bodies, the CMP also changed the `Reject All` button to have the same colour as the `Accept All` button. The banner is shown by blurring the background and placed in a central position on the page. This forces the users to interact with the banner to access the website content. Users quickly click on the "X" or `Reject All` buttons. In a nutshell, they are more interested in discarding the banner rather than making an informed decision. The reject rate tops 70% (and 30% explicitly accepting cookies).
— **C – Remove blurring; D – central position**: In January and March 2023, the CMP took some minor measures to reduce the *Rejected-All* rate: they disabled the blurring effect (C), and moved the banner to the bottom position by default (D). These changes allow the user to see the website content without getting rid of the banner, which causes a reduction in the *Rejected-All* rate. The *Rejected-All* rate drops to 55%. We link this drop to users ignoring the banner (for which we have no records) rather than intentionally selecting to accept cookies.
— **E – lack of data**: In May 2023, the data collection was suspended for a short amount of time for technical reasons.

The same events, although differently combined and at different times, explain the shape that we observe in other regions in Figure 10(b). Focus on Canada (*red curve*) and Brazil (*black curve*): the CMP introduced a GDPR-compliant banner for Brazilians in February 2023 with a configuration similar to A+B (i.e., with the `Reject All` and "X" buttons). The presence of the intrusive banner suddenly makes the users close the banner by clicking on the "X", which is equivalent to selecting *Rejected-All*. In April 2023, the CMP defaulted to C+D, reducing the *Rejected-All* rate by users. The same happened in Canada, where users started interacting with the GPDR-compliant banner in July 2023.

For other countries, the CMP shows the General banner with neither the "X" nor the `Reject All` button (see Figure 2(a)). The *Rejected-All* rate is far lower and barely reaches 10%. In this case, a user who would like to disable the usage of profiling and tracking cookies has to perform multiple clicks. We assume that this burden makes the user prefer to either ignore the banner or quickly select the `Accept All` button to dismiss the banner.

In a nutshell, users aim to quickly close the banner. Less than 10% of them go through the burden of manually disabling the data collection when more than one click is needed. When the `Reject All` button is present, users select it in 20% of cases. When the "X" button is present (and equivalent to the `Reject All` action), users just click on it to quickly close the privacy banner.

For completeness, the curves in Figure 10 include the per-region average *Rejected-All* rate if at least 10 interactions are collected for the given samples (i.e., at least 10 interactions on a website by users from the given region). This causes the curves to be noisy and possibly lack some points, especially in the early months when the number of interactions is low (see Figure 9).

5.2 Stationary Analysis

We now focus on the last 2 months of illow.io data (from July 5 to the end of August 2023, highlighted in grey in previous plots). By doing so, we ensure that the CMP operation is not changing, allowing us to drill down the analysis.

Another insight regarding **RQ5** is provided in Figure 11, in which we break down the *Rejected-All* rate based on users' country. To show only the most represented countries, we include countries with at least 50 websites recording 50 visits in the observation period rather than only 10 visits. The main differences are ascribable to the different versions of the banner the users are shown. The *Rejected-All* rate varies between 45% and 65% for countries with a GDPR-compliant banner (which

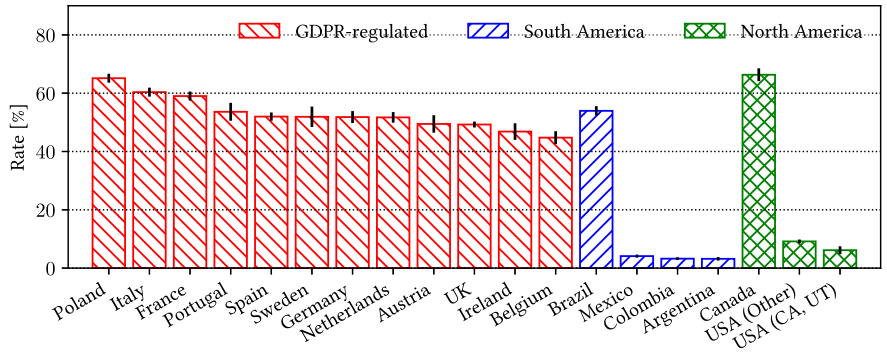

Fig. 11. The *Rejected-All* rate per country. illow.io dataset from July 3, 2023.

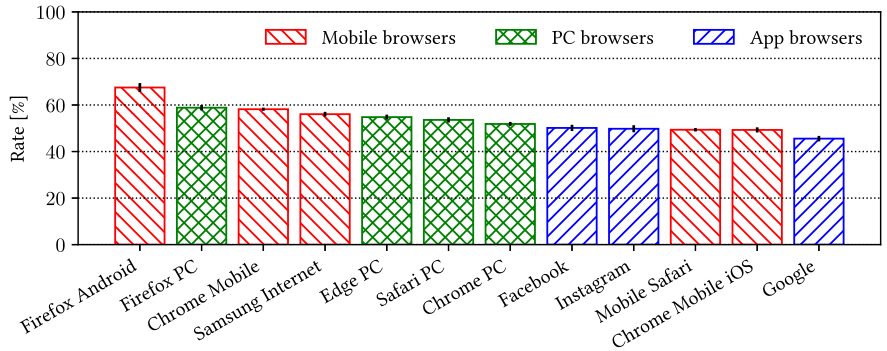

Fig. 12. Average *Rejected-All* rate by browsers. illow.io dataset from July 1, 2023.

includes Brazil and Canada). Conversely, countries with default banners do not exceed the 10% *Rejected-All* rate. We attribute the differences among countries with the same banner to different cultural habits, as differences are consistent across all of our dataset dimensions. As an example, we compare the *Rejected-All* rate for Italy and the United Kingdom in Appendix B.

Peculiar is the case of California and Utah, where users face the CCPA-compliant banner (as in Figure 2(c)). The *Rejected-All* rate for these two states (6.6%) is only moderately lower than in the rest of the United States (8.8%), where users are offered the General banner (as in Figure 2(a)). These numbers are interesting if we focus on how the two types of banners operate. The CCPA-compliant banner is based on the *opt-out* mechanism, i.e., if the user ignores or closes the banner, the consent is presumed. Conversely, with the General banner, the user must explicitly *opt in* to cookies, although opting in is faster and easier than opting out (which requires accessing the selection windows). The opt-in mechanism is considered more respectful of users' privacy, but our results show that when the banner design encourages users to opt in, the practical effect is almost comparable to an opt-out banner.

We now focus on other aspects and study how the users' browser and device impact their behaviour with privacy banners. We explore different dimensions, addressing the different subquestions of **RQ6**. In Figure 12, we compare the *Rejected-All* rate for users using different browsers to surf the Internet (**RQ6a**). We focus on users facing a GDPR-compliant banner, and we group the browsers into three categories for ease of visualization.

— **Mobile browsers – red pattern**: Full browsers operating on a mobile device, such as a smartphone or a tablet. Among them, we include Chrome Mobile, Samsung Internet, and

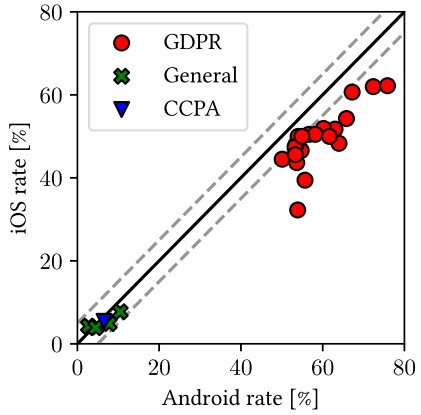

Fig. 13. Android versus iOS: average *Rejected-All* rate for users in each country since July 3, 2023. The dotted lines represent a ±5% difference.

Firefox Mobile on Android devices, and the Mobile Safari and Chrome Mobile for iOS on iOS and iPadOS devices.

— **PC browsers – green pattern**: Full browsers operating on a laptop or desktop computer. We include Google Chrome, Safari, Edge, and Firefox.

— **App browsers – blue pattern**: In-app browsers offered by popular applications on smartphones and tablets. These apps let the user browse websites in the app when they click a link. We observe a significant usage of Facebook, Instagram, and Google Search apps directly requesting and showing webpages to users.

Interestingly, Firefox stands first among both Android and PC browsers. Although we cannot prove any hypothesis on this outcome, we suppose that Firefox users are careful about their privacy and tend to refuse the use of cookies more often than other users. At the opposite end, users navigating app browsers tend to refuse less. This behaviour might have multiple causes: for example, users inside a third-party app's environment might be less encouraged to make an informed choice about cookies given the unusual surfing context they are in. Or the more occasional browsing makes them quickly return to the app without spending time on the webpage content.

We now draw attention to mobile, comparing the behaviour of Android and iOS users, for which we find notable differences (**RQ6b**). In general, we find that Android users have high *Rejected-All* rates (see, for example, Firefox Android in Figure 12). This may be caused by the large relative banner area, which covers a large fraction of the screen. Users quickly dismiss the banner by clicking on the "X" more frequently on a mobile than on a desktop browser (we will investigate this later). However, for iOS users, the picture is different. For instance, Mobile Safari and Chrome Mobile iOS are in the last positions in Figure 12. We explore the remarkable differences between Android and iOS users more in depth in Figure 13. For every country, we plot the *Rejected-All* rate for users using the Android versus iOS operating system. We consider countries with at least 10 websites with 50 interactions each. The *diagonal dashed lines* represent a ±5% difference. Interestingly enough, in every country we observe that iOS users present a significantly smaller *Rejected-All* rate than Android users. Our data cannot offer us any insight into the reason for this difference. We hypothesize that this could be owing to iOS users being more confident about their privacy than their Android counterparts because of their faith in privacy-friendly solutions offered by the Apple ecosystem. A confirmation of this hypothesis is a potential line of future research.

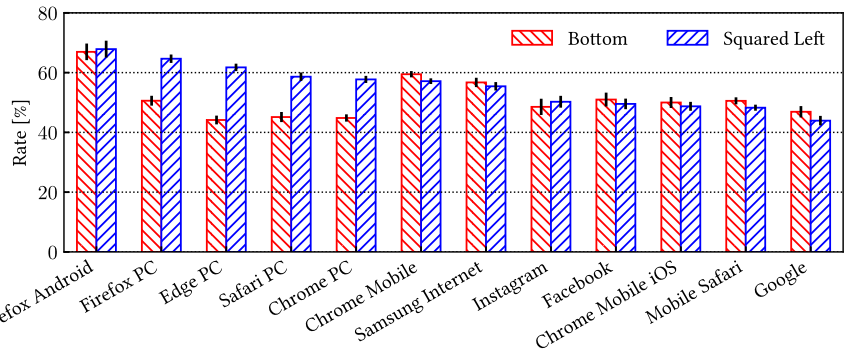

Fig. 14. Average *Rejected-All* rate by browsers, divided by banner position. illow.io dataset since July 3, 2023.

Finally, we differentiate the outcome with respect to the privacy banner position on the screen (**RQ6c**). We consider the two most popular banner options: *Bottom*, i.e., the banner appears on the lower part of the screen, covering a portion of the bottom area with an overlay window; *Squared Left*, i.e., the banner appears as a pop-up in the bottom-left corner of the screen (see Figure 2), resulting in more visibility. In Figure 14, we observe that the position of the banner makes a significant difference only on Desktop browsers. In fact, the banners overlayed on the bottom of the page register a lower *Rejected-All* rate than the Squared Left pop-up style banner. This hints at users being annoyed by the pop-up banner window, nudging more people to dismiss it via the "X" button to freely access the page. In Mobile and App browsers on smartphones and tablets, there is no significant difference in the *Rejected-All* rate between the two types of banners. This is because there is no difference between the layout of the Bottom and the Squared Left banners: both cover a large fraction of the mobile device screen.

5.3 Other Findings

Here, we report some additional findings regarding users' interaction with the illow.io privacy banner (**RQ6d**).

— Access to Cookie Policy and the Privacy Policy: From August 2022 to April 2023, the illow.io dataset reported clicks on the Cookie Policy and the Privacy Policy links offered by the banner. Only 0.22% of users explored this possibility and clicked on at least one of the two policies. It is clear that common website visitors are not interested in spending time reading long and detailed policy text, confirming the findings of related work [40, 51, 53].

— Fine-grained consent: Few users choose to fine-tune the categories of cookies they are willing to accept. In the whole illow.io dataset, we count only 8,297 interactions that offered a partial acceptance over almost 25 million interactions. The statistics category is the most accepted one (62.77 ± 1.04%), followed by preferences (52.95 ± 1.07%). The marketing category (arguably, the one most associated with the feeling of lack of privacy) is the least accepted one (19.63 ± 0.85%). All confidence intervals are computed with a 95% probability on a per-website grouping.

6 Conclusion

In this article, we discussed different aspects related to CMPs, considering both their penetration into the Web ecosystem and the way users interact with them.

Thanks to the long-lasting effort of the HTTP Archive, we showed the significant growth of CMP adoption on the Web to ease the management of different and evolving privacy regulations.

We witnessed that CMP adoption significantly varies according to region and when the privacy regulations took effect. Thanks to the GDPR, Europe leads CMP adoption, with other regions of the world growing quickly when regulation takes effect, generating new opportunities for the CMP market.

Our results also suggest that some websites and CMPs still allow the browser to contact tracking systems and let them install persistent cookies before the user accepts their use. However, the limitations of the HTTP Archive call for further investigations to verify whether this corresponds to an actual violation of privacy legislation. At large, this shows the need to extend the availability of archiving platforms such as HTTP Archive owing to the intertwining of the regulations, user location, device type, language, etc. In fact, the customisation of content that websites serve to users includes the way CMPs show the privacy banners and their actions.

Considering users' interactions with CMPs, our results showed that, when offered a viable option to refuse data collection, users tend to do so. Even more clearly, we showed that a large percentage of users tend to dismiss the privacy banner without actively interacting with it but rather simply clicking on the "close" button if present without making any explicit decision about their privacy. Such behaviour challenges the concept that offering users fine-grained options about their privacy empowers them to make the best decision. For instance, when annoyed with more intrusive banners, the reject rate increases, inflated by the need to close the banner. Conversely, when refusing data collection is cumbersome and requires more than one click, users tend not to take the opportunity. Interestingly, with an opt-out banner, where, by default, consent is presumed, the rejection rate is similar.

Finally, our results showed that users operating on mobile devices are apparently the most keen to refuse data collection—in fact, they just seem to counterbalance the intrusiveness of the banner on mobile device screens. Curiously, iOS users tend to refuse at a lower rate than Android users. Although we formulate some hypotheses about it, this observation lays the groundwork for interesting future research.

References

[1] 2024. AdGuard. Retrieved March 18, 2025 from https://adguard.com/

[2] 2024. EasyPrivacy. Retrieved March 18, 2025 from https://easylist.to/easylist/easyprivacy.txt

[3] 2024. HTTP Archive. Retrieved March 18, 2025 from https://httparchive.org/

[4] 2024. HTTP Archive Open Dataset used for CMP Analysis. Retrieved March 18, 2025 from https://smartdata.polito.it/consent-management-platforms-growth-and-users-interactions-over-time/

[5] 2024. Wappalyzer. Retrieved March 18, 2025 from https://www.wappalyzer.com/

[6] 2024. WhoTracks.me. Retrieved March 18, 2025 from https://whotracks.me/

[7] Gunes Acar, Christian Eubank, Steven Englehardt, Marc Juarez, Arvind Narayanan, and Claudia Diaz. 2014. The web never forgets: persistent tracking mechanisms in the wild. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*. 674–689.

[8] Atheer Aljeraisy, Masoud Barati, Omer Rana, and Charith Perera. 2022. Exploring the relationships between privacy by design schemes and privacy laws: A comparative analysis. *arXiv preprint arXiv:2210.03520* (2022).

[9] Susanne Barth and Menno D. T. de Jong. 2017. The privacy paradox—Investigating discrepancies between expressed privacy concerns and actual online behavior—A systematic literature review. *Telematics and Informatics* 34, 7 (2017), 1038–1058. https://doi.org/10.1016/j.tele.2017.04.013

[10] Jan M. Bauer, Regitze Bergstrøm, and Rune Foss-Madsen. 2021. Are you sure, you want a cookie?—The effects of choice architecture on users' decisions about sharing private online data. *Computers in Human Behavior* 120 (2021), 106729.

[11] Brazilian President of the Republic. 2018. Lei Geral de Proteção de Dados Pessoais. Retrieved March 18, 2025 from http://www.planalto.gov.br/ccivil_03/_ato2015-2018/2018/lei/L13709compilado.htm

[12] California State Legislature. 2018. California Consumer Privacy Act of 2018. Retrieved March 18, 2025 from https://leginfo.legislature.ca.gov/faces/billTextClient.xhtml?bill_id=201720180AB375

[13] Council of European Union. 2009. Directive 2009/136/EC Amending Directive 2002/22/EC on Universal Service and Users' Rights Relating to Electronic Communications Networks and Services, Directive 2002/58/EC Concerning the

Processing of Personal Data and the Protection of Privacy in the Electronic Communications Sector and Regulation (EC) No 2006/2004 on Cooperation between National Authorities Responsible for the Enforcement of Consumer Protection Laws. Retrieved March 18, 2025 from http://eur-lex.europa.eu/legal-content/en/TXT/?uri=CELEX:32009L0136

[14] Lynne M. Coventry, Debora Jeske, John M. Blythe, James Turland, and Pam Briggs. 2016. Personality and social framing in privacy decision-making: A study on cookie acceptance. *Frontiers in Psychology* 7 (2016), 1341.

[15] Adrian Dabrowski, Georg Merzdovnik, Johanna Ullrich, Gerald Sendera, and Edgar Weippl. 2019. Measuring cookies and web privacy in a post-GDPR world. In *Passive and Active Measurement*, David Choffnes and Marinho Barcellos (Eds.). Springer International Publishing, Cham, 258–270.

[16] Data Protection Authorities of Germany. 2024. Orientierungshilfe der Konferenz der unabhängigen Datenschutzaufsichtsbehörden des Bundes und der Länder vom 20. Dezember 2021. Retrieved March 18, 2025 from https://www.datenschutzkonferenz-online.de/media/oh/20211220_oh_telemedien.pdf

[17] Marcus Abreu de Magalhães. 2021. Data protection regulation: A comparative law approach: Proteção de dados: Estudo comparado de normas nacionais. *International Journal of Digital Law* 2, 2 (2021), 33–53.

[18] Martin Degeling, Christine Utz, Christopher Lentzsch, Henry Hosseini, Florian Schaub, and Thorsten Holz. 2019. We value your privacy... now take some cookies. *Informatik Spektrum* 42, 5 (2019), 345–346.

[19] Deloitte. 2020. Cookie Benchmark Study. Retrieved March 18, 2025 from https://www2.deloitte.com/content/dam/Deloitte/nl/Documents/risk/deloitte-nl-risk-cookie-benchmark-study.pdf (2020).

[20] Rob van Eijk, Hadi Asghari, Philipp Winter, and Arvind Narayanan. 2019. The impact of user location on cookie notices (inside and outside of the European Union). In *Workshop on Technology and Consumer Protection (ConPro'19)*.

[21] Steven Englehardt and Arvind Narayanan. 2016. Online tracking: A 1-million-site measurement and analysis. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*. 1388–1401.

[22] José Estrada-Jiménez, Javier Parra-Arnau, Ana Rodríguez-Hoyos, and Jordi Forné. 2017. Online advertising: Analysis of privacy threats and protection approaches. *Computer Communications* 100 (2017), 32–51.

[23] European Parliament and Council of European Union. 2016. Directive 95/46/EC. General Data Protection Regulation. Retrieved March 18, 2025 from http://data.consilium.europa.eu/doc/document/ST-5419-2016-INIT/en/pdf

[24] Jens Grossklags and Nathan Good. 2007. Empirical studies on software notices to inform policy makers and usability designers. In *International Conference on Financial Cryptography and Data Security*. Springer, 341–355.

[25] Hana Habib, Megan Li, Ellie Young, and Lorrie Cranor. 2022. "Okay, whatever"': An evaluation of cookie consent interfaces. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. 1–27.

[26] Hana Habib, Sarah Pearman, Jiamin Wang, Yixin Zou, Alessandro Acquisti, Lorrie Faith Cranor, Norman Sadeh, and Florian Schaub. 2020. "It's a scavenger hunt": Usability of websites' opt-out and data deletion choices. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (CHI'20)*. ACM, New York, NY, 1–12. https://doi.org/10.1145/3313831.3376511

[27] Philip Hausner and Michael Gertz. 2021. Dark patterns in the interaction with cookie banners. *arXiv preprint arXiv:2103.14956* (2021).

[28] Maximilian Hils, Daniel W. Woods, and Rainer Böhme. 2020. Measuring the emergence of consent management on the web. In *Proceedings of the ACM Internet Measurement Conference (IMC'20)*. ACM, New York, NY, 317–332. https://doi.org/10.1145/3419394.3423647

[29] Italian Data Protection Authority. 2024. Chiarimenti in merito all'attuazione della normativa in materia di cookie. Retrieved March 18, 2025 from https://www.garanteprivacy.it/web/guest/home/docweb/-/docweb-display/docweb/4006878

[30] Nikhil Jha, Martino Trevisan, Marco Mellia, Rodrigo Irarrazaval, and Daniel Fernandez. 2023. I refuse if you let me: Studying user behavior with privacy banners at scale. In *2023 7th Network Traffic Measurement and Analysis Conference (TMA'23)*. 1–9. https://doi.org/10.23919/TMA58422.2023.10198936

[31] Nikhil Jha, Martino Trevisan, Luca Vassio, and Marco Mellia. 2022. The Internet with privacy policies: Measuring the Web upon consent. *ACM Transactions on the Web (TWEB)* 16, 3 (2022), 1–24.

[32] Michael Kretschmer, Jan Pennekamp, and Klaus Wehrle. 2021. Cookie banners and privacy policies: Measuring the impact of the GDPR on the web. *ACM Trans. Web* 15, 4, Article 20 (Jul 2021), 42 pages. https://doi.org/10.1145/3466722

[33] Oksana Kulyk, Willard Rafnsson, Ida Marie Borberg, and Rene Hougard Pedersen. 2022. "So I Sold My Soul": Effects of dark patterns in cookie notices on end-user behavior and perceptions. In *Proceedings of 2022 Symposium on Usable Security and Privacy*. Internet Society.

[34] Thomas Linden, Rishabh Khandelwal, Hamza Harkous, and Kassem Fawaz. 2020. The privacy policy landscape after the GDPR. *Proceedings on Privacy Enhancing Technologies* 2020, 1 (2020), 47–64.

[35] 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'20)*. 791–809. https://doi.org/10.1109/SP40000.2020.00076

[36] Jonathan R. Mayer and John C. Mitchell. 2012. Third-party web tracking: Policy and technology. In *2012 IEEE Symposium on Security and Privacy*. IEEE, 413–427.

[37] Hassan Metwalley, Stefano Traverso, Marco Mellia, Stanislav Miskovic, and Mario Baldi. 2015. The online tracking horde: A view from passive measurements. In *International Workshop on Traffic Monitoring and Analysis*. Springer, 111–125.

[38] National Assembly of Québec. 2021. An Act to Modernize Legislative Provisions as Regards the Protection of Personal Information. Retrieved March 18, 2025 from https://www.publicationsduquebec.gouv.qc.ca/fileadmin/Fichiers_client/lois_et_reglements/LoisAnnuelles/en/2021/2021C25A.PDF

[39] Midas Nouwens, Ilaria Liccardi, Michael Veale, David Karger, and Lalana Kagal. 2020. Dark patterns after the GDPR: Scraping consent pop-ups and demonstrating their influence. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (CHI'20)*. ACM, New York, NY, 1–13. https://doi.org/10.1145/3313831.3376321

[40] Jonathan A. Obar and Anne Oeldorf-Hirsch. 2020. The biggest lie on the internet: Ignoring the privacy policies and terms of service policies of social networking services. *Information, Communication & Society* 23, 1 (2020), 128–147.

[41] 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. ACM, New York, NY, 2130–2141.

[42] Parliament of Canada. 2000. Personal Information Protection and Electronic Documents Act. An Act to Modernize Legislative Provisions as Regards the Protection of Personal Information. Retrieved March 18, 2025 from https://laws-lois.justice.gc.ca/eng/acts/p-8.6/

[43] Parliament of Canada. 2023. An Act to Enact the Consumer Privacy Protection Act, the Personal Information and Data Protection Tribunal Act and the Artificial Intelligence and Data Act and to Make Consequential and Related Amendments to other Acts. Retrieved March 18, 2025 from https://www.parl.ca/legisinfo/en/bill/44-1/c-27

[44] Enric Pujol, Oliver Hohlfeld, and Anja Feldmann. 2015. Annoyed users: Ads and ad-block usage in the wild. In *Proceedings of the 2015 Internet Measurement Conference*. 93–106.

[45] European Commission, Joint Research Centre, N. Rodríguez-Priego, and R. v. Bavel. 2016. Testing the effect of the cookie banners on behaviour. Publications Office. DOI: https://doi.org/doi/10.2791/22197

[46] Valentino Rizzo, Stefano Traverso, and Marco Mellia. 2021. Unveiling web fingerprinting in the wild via code mining and machine learning. *Proceedings on Privacy Enhancing Technologies* 2021, 1 (2021), 43–63.

[47] 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 *Proceedings of the 2019 ACM Asia Conference on Computer and Communications Security (Asia CCS'19)*. ACM, New York, NY, 340–351. https://doi.org/10.1145/3321705.3329806

[48] Ashutosh Kumar Singh, Nisarg Upadhyaya, Arka Seth, Xuehui Hu, Nishanth Sastry, and Mainack Mondal. 2022. What cookie consent notices do users prefer: A study in the wild. In *Proceedings of the 2022 European Symposium on Usable Security*. 28–39.

[49] Janice C. Sipior, Burke T. Ward, and Ruben A. Mendoza. 2011. Online privacy concerns associated with cookies, flash cookies, and web beacons. *Journal of Internet Commerce* 10, 1 (2011), 1–16.

[50] Than Htut Soe, Oda Elise Nordberg, Frode Guribye, and Marija Slavkovik. 2020. Circumvention by design—dark patterns in cookie consent for online news outlets. In *Proceedings of the 11th Nordic Conference on Human-Computer Interaction: Shaping Experiences, Shaping Society (NordiCHI'20)*. ACM, New York, NY, Article 19, 12 pages. https://doi.org/10.1145/3419249.3420132

[51] Nili Steinfeld. 2016. "'I agree to the terms and conditions"': (How) do users read privacy policies online? An eye-tracking experiment. *Computers in Human Behavior* 55 (2016), 992–1000. https://doi.org/10.1016/j.chb.2015.09.038

[52] Martino Trevisan, Stefano Traverso, Eleonora Bassi, and Marco Mellia. 2019. 4 years of EU cookie law: Results and lessons learned. *Proc. Priv. Enhancing Technol.* 2019, 2 (2019), 126–145.

[53] Janice Y. Tsai, Serge Egelman, Lorrie Cranor, and Alessandro Acquisti. 2011. The effect of online privacy information on purchasing behavior: An experimental study. *Information Systems Research* 22, 2 (2011), 254–268.

[54] Christine Utz, Martin Degeling, Sascha Fahl, Florian Schaub, and Thorsten Holz. 2019. (Un)Informed consent: Studying GDPR consent notices in the field. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security (CCS'19)*. ACM, New York, NY, 973–990. https://doi.org/10.1145/3319535.3354212

[55] Tony Vila, Rachel Greenstadt, and David Molnar. 2003. Why we can't be bothered to read privacy policies models of privacy economics as a lemons market. In *Proceedings of the 5th International Conference on Electronic Commerce*. 403–407.

Appendices

A Dataset Structures

In this Appendix, we provide the exact structure of our two datasets. Both datasets are in tabular format. In what follows, we report the columns relevant to our analysis, omitting the columns not useful for our purposes.

A.1 HTTP Archive

We use the public data offered by the HTTP Archive to build the HTTP-9Years and HTTP-Dec23 datasets. The preprocessed data and the code to obtain the figures of this article are available online [4]. In particular, we downloaded via Google Cloud Storage the tables called *Summary Pages* and *Summary Requests*, containing details on the visited pages and all issued HTTP requests, respectively. We downloaded the tables for the 9 years of interest. Relevant columns are the following.

- **Summary Pages**
  — Numerical identifier of the visit
  — URL of the visited website
  — Time of the visit
- **Summary Requests**
  — Identifier of the corresponding visit
  — URL requested by the browser in the HTTP request
  — Presence of the `Set-Cookie` header. The dataset does not include the content of the cookie.

A.2 illow.io CMP Dataset

This dataset is contained in a single data table, in which each entry corresponds to an *interaction*. The columns relevant to our analysis are as follows.

- Time of the interaction
- Anonymous identifier of the interaction
- Anonymous identifier of the website
- Client's /24 IPv4 subnet
- Client's `User-Agent` as set by the browser in the HTTP requests
- Consents provided among the following cookie categories
  — Necessary (always present)
  — Statistical
  — Preferential
  — Marketing

B Italy versus United Kingdom

Here, we compare the *Rejected-All* rate of Italy and the United Kingdom to show that consistent differences exist across all of the dimensions of our dataset. We presume that the differences in rates between the two countries (as shown in Figure 11) only depend on users' cultural habits.

In Figure 15, we compare the *Rejected-All* rate by Italian and British users. The distance between the two countries has been sizeable and stable since the beginning of 2023. We can thus rule out that any variation in the CMP operation has biased the measurement.

In Figure 16, we partition the dataset across the most popular browsers. The *Rejected-All* rate for Italy is always higher than for the United Kingdom regardless of the considered browser (although for some of them, the distance is minimal – e.g., Firefox PC). We can also rule out that the difference between the two countries is due to the different browser popularity. With the information at our

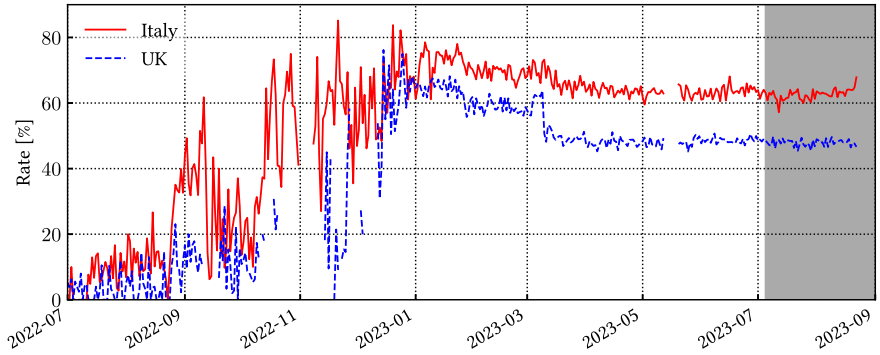

Fig. 15. The *Rejected-All* rate of Italy and UK.

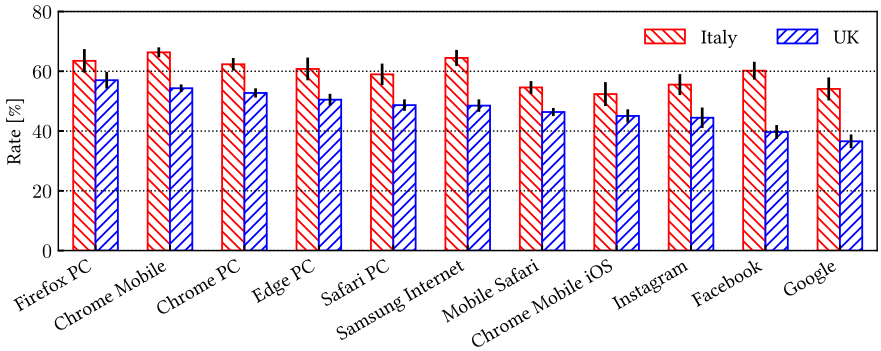

Fig. 16. The *Rejected-All* rate of Italy and UK.

disposal, we thus assume that the difference is attributable to different cultural habits among the populations.

Received 29 January 2024; revised 23 December 2024; accepted 22 September 2024