Stop the Consent Theater

Matthias Fassl*
matthias.fassl@cispa.saarland
CISPA Helmholtz Center for Information Security
Saarland University

Lea Gröber*
lea.groeber@cispa.saarland
CISPA Helmholtz Center for Information Security
Saarland University

Katharina Krombholz
krombholz@cispa.saarland
CISPA Helmholtz Center for Information Security

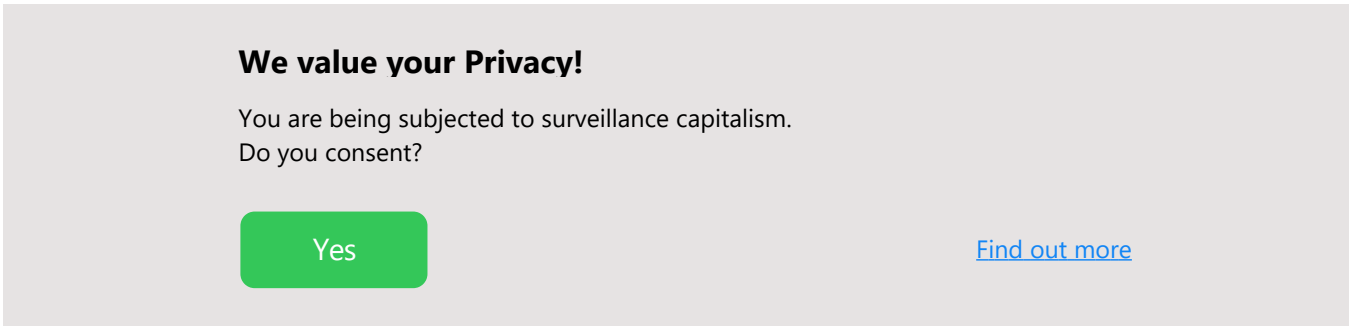

Figure 1: An artistic interpretation of an honest cookie consent notice.

ABSTRACT

The current web pesters visitors with consent notices that claim to "value" their privacy, thereby habituating them to accept all data practices. Users' lacking comprehension of these practices voids any claim of informed consent. Market forces specifically designed these consent notices in their favor to increase users' consent rates. Some sites even ignore users' decisions entirely, which results in a mere theatrical performance of consent procedures designed to appear as if it fulfills legal requirements.

Improving users' online privacy cannot rely on individuals' consent alone. We have to look for complementary approaches as well. Current online data practices are driven by powerful market forces whose interests oppose users' privacy expectations – making turnkey solutions difficult. Nevertheless, we provide a bird's-eye view on privacy-improving approaches beyond individuals' consent.

CCS CONCEPTS

• **Security and privacy** → **Social aspects of security and privacy**; **Privacy protections**; **Economics of security and privacy**; • **Human-centered computing** → *Web-based interaction.*

KEYWORDS

Cookie Consent Notices, Online Behavioral Advertising, Surveillance Capitalism

ACM Reference Format:
Matthias Fassl, Lea Gröber, and Katharina Krombholz. 2021. Stop the Consent Theater. In *CHI Conference on Human Factors in Computing Systems Extended Abstracts (CHI '21 Extended Abstracts), May 8–13, 2021, Yokohama, Japan.* ACM, New York, NY, USA, 7 pages. https://doi.org/10.1145/3411763.3451230

1 THOSE DAMN COOKIE BANNERS!!

Cookie banners (i.e., consent notices) that randomly pop up to obstruct your view of something interesting on the web are, without question, annoying. You are probably just as guilty as we have been of clicking it away without reading what it said, thereby submitting to questionable privacy practices. The sheer number of consent notices users encounter while browsing the web can be aggravating. You might ask why this is even necessary. Consent notices are the naïve answer to the complex problem of informational self-determination on the web, where sites store so-called cookies in your browser all the time. Cookies' intended use was mostly centered around e-commerce features such as virtual shopping carts and remembering logged in users. However, cookies turned out to be a useful tool to mark visitors of one site and track their behavior across other sites.

The intention behind consent notices is to grant users control over their data and thus enhance their privacy. This overall goal is becoming increasingly pressing with the rise of large-scale data gathering and accompanying surveillance. How users could effectively control if and how others profit from their personal data remains an open question. Admittedly, consent notices are doing a poor job at achieving this goal. They are supposably designed to obtain users' informed consent, yet nudge users to allow tracking

*Both authors contributed equally to this research.

CHI '21 Extended Abstracts, May 8–13, 2021, Yokohama, Japan

ACM ISBN 978-1-4503-8095-9/21/05.
https://doi.org/10.1145/3411763.3451230

behavior [1]. These nudges include default settings that allow all kinds of cookies or the deceiving visual presentation and phrasing of consent notices. In many cases, they do not comply with relevant privacy laws, such as GDPR or the ePrivacy directive [15, 17, 24, 28]. And to top it off, they often do not respect users' decisions, even if they opt-out [15]. The current situation supposedly exists for the visitors' privacy benefit but not all involved parties want to achieve that goal, resulting in a consent theater (similar to Schneier's *security theater* concept [25]). The ad-industry directs this theater, the content providers and their visitors are the performers on stage, and the regulators and policymakers are the audiences. This play aims to convince the audiences of the ongoing practices' legality while minimizing harm to the director's interest.

This paper offers a bird's-eye view on the issue of consent-notices, the accompanying data practices, and its different stakeholders. We reviewed and analyzed the relevant body of research in this area, touching on topics like tracking, online behavioral advertisements, and legal compliance. We found that in many papers, users' consent is the central point of investigation. However, this perspective is not sufficient to solve the underlying issues, as we have to deal with a complex ecosystem and conflicting interests. Building upon this, we outline future venues of research and discuss the responsibilities of different stakeholders. The goal of this work is to bring attention to and ultimately make progress towards stopping the consent theater.

2 HOW COULD IT GET SO BAD?

Internet sites have begun early on to finance their content with advertisements. Users understood and accepted that trade as consuming ads while getting content or functionality for free [29]. The nature of this accepted trade changed after the dot-com bubble burst. Google, which was under pressure to generate revenue, discovered that they could use the behavioral surplus (i.e., users' excess behavioral data that the service itself does not utilize) from their search engine to enhance the quality and value of online behavioral advertising (OBA). This triggered the discovery of surveillance capitalism's foundation: The endeavor to corner the market for big-data-based prediction products requires an arms race to collect and accumulate increasing amounts of behavioral surplus [33]. Since then, surveillance capitalists do not build new tech products for the benefit of their users but to increase their access to behavioral surplus.

Resulting from this development, website providers who use ads to monetize their content now also hand over their users' behavioral surplus (i.e., who viewed their content at which time) to advertisers. Figure 2 provides an overview of personal data flows in online behavioral advertising. In all cases, the advertisers embed code into the content providers' site. This code asks for the users' consent (optionally using consent management platforms) and only then contacts the advertisers' behavioral trackers. The trackers identify the visitors across multiple websites and use this access information to build users' profiles. Increasingly detailed user profiles make the ad-space more valuable since trackers can use them to predict the advertisements' effectiveness reliably. Commonly, trackers use regular third-party cookies to identify visitors. However, the user tracking does not depend on them since other mechanisms such as cookie-synchronization, browser fingerprinting, super-cookies, or Verizon's Unique Identifier Header (UIDH) [30] achieve the same effect. After the behavioral tracker identifies a visitor, a real-time auction (based on the visitor's profile) determines which advertisement they will see. This type of online behavioral advertising (OBA) has been mostly invisible to users since browsers just accepted any cookies by default. However, the ePrivacy regulation and GDPR made these practices more visible since users need to consent to them.

The two main stakeholders in OBA are the content providers' sites and their visitors. Visitors want to consume content or access functionality, and sites want to monetize their content. Other stakeholders, apart from these two, are crucial for the business model of OBA. The ad-industry discovered that sufficiently large datasets on users' online behavior allow behavioral prediction and behavioral modification. This extraordinarily lucrative opportunity drives the ongoing efforts to collect increasingly more behavioral data since even small improvements in predictions are worth millions [33]. Regulators and policymakers have found that these data practices negatively impact users' right to self-determination [6, 22]. This is exemplified by the Internet users who are, despite the "consent" mechanisms, unaware of how their personal data is used [18, 29]. Currently, their response consists of a policy to increase awareness and repeatedly emphasize that users have the freedom of choice of being surveilled or not (in the form of the ubiquitous consent notices). In light of the significant fines that GDPR allows, consent management platforms have become popular. They are used to ease advertisers' legal liability while simultaneously optimizing users' consent rates for profit.

The fundamental conflict of interest in OBA is that trackers base their entire business model on detailed user profiles that demand ever-increasing behavioral data. All the while visitors feel uncomfortable with that level of tracking and regulators want to reestablish reasonable ways of informational self-determination. This conflict of interest is difficult to untangle since many current internet sites' monetization depends on it.

3 OBSERVATIONS FROM PRIOR SECURITY AND PRIVACY RESEARCH

The introduction of the ePrivacy directive and GDPR reignited research interests in the topic of online consent mechanisms. Most of the research surrounds questions of regulatory effectiveness and users' perceptions of consent notices and data practices. We reviewed and analyzed prior work to see how researchers frame the continued problems, which systematic issues emerge, and which kinds of questions researchers pose. In this section, we present three key-issues we identified and take a closer look at solutions proposed in prior work.

3.1 Key Issue 1: Users largely do not understand that behavioral online ads require tracking

Research has invested considerable efforts into understanding the internet users' perceptions and behaviors concerning online tracking in general [4, 12, 18] and behavioral advertising in particular [16, 29, 31]. The findings paint a complicated picture and

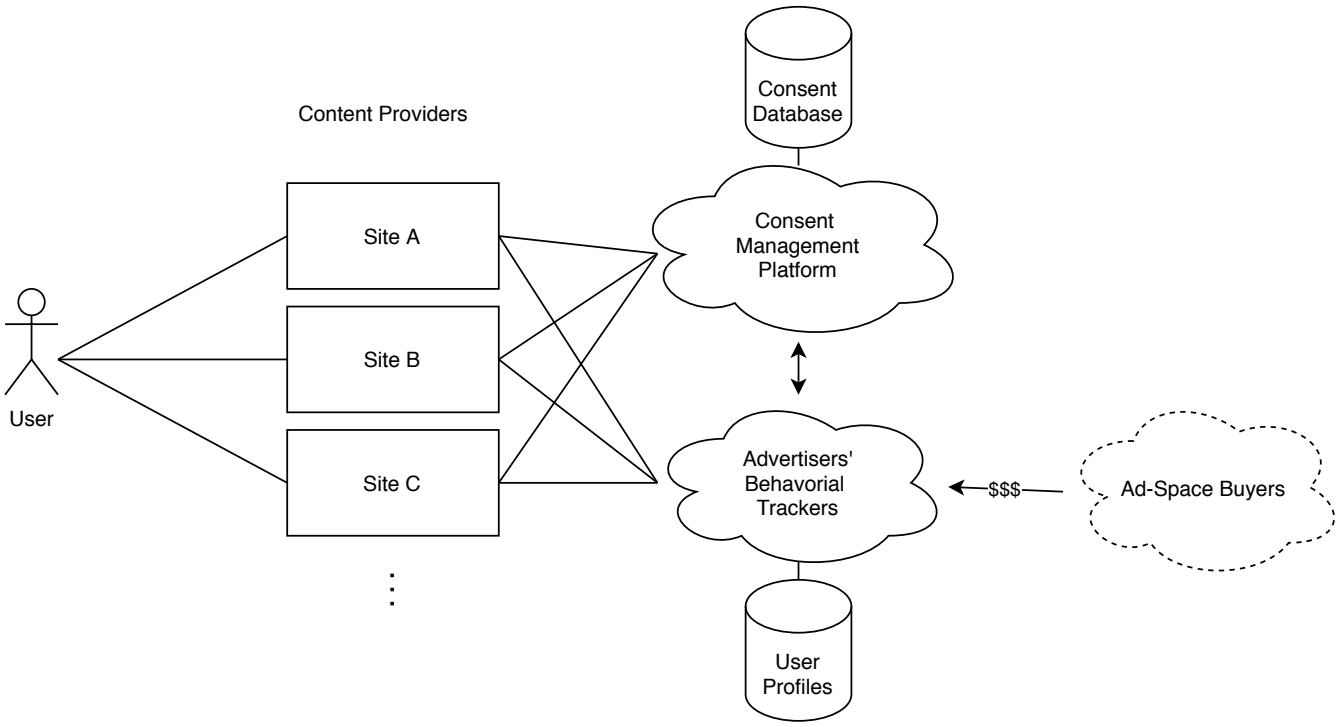

Figure 2: Online behavioral advertising (OBA) and all involved parties

conclude that user decisions are highly individual and context-dependent, while at the same time exposing little understanding for the complexity of web tracking. Melicher et al. investigated users' contextual preferences for web tracking and found that users base their decisions to accept targeted ads mostly on concerns about functionality [18]. Complementing, people saw value in user-customized targeted ads [4, 18] and behavioural advertising [29]. However, acceptance of online behavioral advertising is highly complex and context-dependent [29]. As Melicher et al. found, while 74% of their participants preferred targeted ads, 60% considered targeted ads to be harmful in at least one tracking scenario [18]. Apart from the nuanced attitudes towards targeted online advertising, people often do not connect it with the concept of tracking. That is, McDonald et al. found that while their participants in principle approved the idea of advertisement funding free online content, they did not expect to be tracked in the process [16]. Research on users' perception of tracking found a general lack of information and understanding of the technologies and concepts involved [12, 18, 31]. For example, Kulyk et al. found that people do not understand how cookies work, what that meant for their life online, and what potential countermeasures look like [12]. If asked how tracking works, people reveal concerns about their ability to make sound decisions given their limited knowledge about the topic [18]. Yao et al. found that even tech-savvy people, such as web developers, did not know how online tracking works [31]. In their study, they identified four "folk models" about how online behavioral advertisement works and found these models to be either incomplete or inaccurate. To name a few misconceptions, people thought trackers are hackers, viruses, or have access to local files on the computer. If people have mixed opinions about targeted advertisements and do not fully understand the workings and implications of tracking, which factors do influence their tracking preferences? Research identified external factors about the visited websites and the tracked data. On the one hand, trustworthiness of the webpage, familiarity to the user, and importance of the web page's contents [12]. On the other hand, the kind of data plays a role, such as personal or search information, correspondence, financial, educational, or health information [18]. Finally, Melicher et al. found that awareness and consent to tracking impacts people's comfort with such practices [18]. However, this is problematic and might promote a false sense of privacy. Also, it hints at the core of the tracking issue: you cannot consent to something you do not understand.

3.2 Key Issue 2: The ad-industry carefully crafts consent mechanisms to deceive users

In the opinion of Acquisti et al. [1], neither paternalistic (government regulation) nor strictly libertarian (self-regulatory) privacy rules guarantee desired privacy outcomes. Instead, Acquisti et al. suggest that soft-paternalistic approaches (i.e., nudges) may bridge the gap between both approaches. In 1997 advertisers fought bitterly against a proposal by the FTC that would have assigned users control over their personal information by default [33]. Instead, they formed the Network Advertising Initiative that promised to regulate the industry's behavior. Garlach et al. [8], studied the effects of NAI's self-regulatory notice-and-choice model and found that users did not see the notice and did not understand its purpose.

They found it hard to reconcile that the industry has expertise in creating noticeable messages and chose not to create noticeable and informative privacy notices.

Suspiciously, this is a re-occurring pattern in related work. User-interface elements that would give users more options to opt-out of tracking use confusing and and inconspicuous design. Acquisti et al. [1] described different modes of soft-paternalistic nudges: *information*, *presentation*, *defaults*, *incentives*, *reversability*, and *timing*. Privacy-related research has found evidence of most of these nudging modes. Information: Machuletz et al. [14] recognized a proliferation of choice in consent notices and found that users disengage if more than a few data collection purposes are available. Hence, they warn that businesses use the flexibility in designing consent dialogues for their interest by maximizing data disclosure. Presentation: The ad choices icon' design uses unobtrusive colors and sizes and hard to understand phrasing [8]. Commercial software has habituated users to accept anything that resembles an end-user license agreement [3], we have to assume that this is similar for cookie consent notices. The most popular consent notice designs are not the most effective ones [17]. Defaults: Half of the websites do not have a reject all button, and 75% of the sites that do bury this option – requiring more clicks than the alternative. Additionally, 56% of sites pre-tick optional vendors [20]. Timing: Consent-management platforms sometimes ask users excessively for their consent decision (which they already have), bordering on "consent harassment" [15].

The ad-industry seems to use the promise of self-regulation and user consent to avoid stricter legal regulation while simultaneously using nudges and biases to continue their business model. Consequently, we cannot trust the advertisement industry's mechanisms that ask users for consent. On the contrary, we have to assume all parts of these mechanisms use meticulous design to confuse users and increase "consent" rates.

3.3 Key Issue 3: The advertisement industry tries to avoid litigation by providing consent mechanisms on the surface, but often disregards the choice in the implementation

The introduction of GDPR did not substantially impact web tracking. Although its intention was strengthening users' right to privacy, researchers found no decrease in web tracking in the European Union [5]. Degeling et al. argue that the supposed increase in transparency may even lead to a false sense of privacy and security [5]. On top of that, they found that few websites offer meaningful cookie choices to control tracking. Additionally, Mehrnezhad found inconsistencies in the presentation of privacy notices and banners within and across platforms [17]. In desktop and mobile browsers, as well as mobile applications, consent notices often do not comply with GDPR since they tie tracking to page load, not user consent. Several works found non-compliance to GDPR and violations of user consent [15, 24, 28]. Millett et al. identified four types of potential GDPR and the ePrivacy directive violations of European websites which are specific to consent notices: (1) storing consent before the users' choice, (2) not providing users a way to opt-out, (3) pre-selected cookie choices, and (4) not respecting the users' choice [15]. Similarly, Trevisan et al. conducted a large-scale measurement of 35.000 European websites to check the implementation of the EU cookie directive [28]. They found that half of the websites install tracking cookies before users provide their consent. On a global scale, Sanchez et al. investigated users' ability to opt-out of tracking [24]. According to them, opt-out mechanisms are often not correctly implemented and confront users with deceiving information. Additionally, long-lasting cookies are widely used, despite users having opted-out of tracking.

Research demonstrates that users have to face a heterogeneous ecosystem of consent notices that often does not respect their choices, even though they make one.

3.4 Proposed solutions to the flood of cookie decisions

Prior work has tried to find solutions for the flood of cookie decisions for 20 years now. Most of them have tried to make the individual users' cookie decisions easier or more understandable. However, a few also tried to tackle the problem not from an individualistic perspective but rather from collective enforcement of existing laws.

The Netscape browser implemented cookies first in 1994 to enable virtual shopping carts for e-commerce applications – resulting in an increasing burden on users to make cookie-related decisions. In the early 2000s, it was already clear that this burden is growing too large for them since they had little information about these cookies' purposes, and the sheer number of cookie decisions created habituation effects [19]. Friedman et al. [7] applied value-sensitive design to create a browser-integrated cookie information interface to help users to keep track and manage their cookies. Similarly, Kulyk et al. [13] found that cookie-related browser settings confuse users. They designed an alternative settings interface to bridge the gap between the users' privacy preferences and the rather technical jargon of cookie settings. An alternative approach to the problem of increasing cookie decisions is automation. Shankar et al. [26] and Yue et al. [32] identified the problem that users face too many cookie decisions and suggested a tool that only accepts cookies required for a site's functionality. These kinds of solutions can be efficient by applying a pre-defined policy to all decisions. The authors' focus on functionality does not necessarily improve users' privacy, which opens up the question of suitable cookie decision policies. Such automation solutions provide an efficient way of dealing with large number of cookie decisions by moving from individual consent to generally applicable policies. This approach leaves room for a discussion about the types of policies that fit the users' requirements in their different contexts of use.

One thing all of these approaches have in common is that they come with an additional burden to users. They assume that it is the users' responsibility to understand the purposes of cookies and set them in a privacy-aware manner. Even automatic approaches to cookie management require users to handle edge-cases where pre-defined policies fail. However, other stakeholders such as the browser vendors, ISPs, websites, and law enforcement can combat the flood of cookies – by focusing less on individual users' behavior and more on the behavior of other stakeholders has the potential to improve the current status of privacy on the web substantially. An

example of this kind of approach to improving cookie privacy is Trevisan et al.'s large-scale measurement [28]. They measured sites' cookie handling and found many configurations that do not respect the ePrivacy directive. Their audit tool, CookieCheck, automatically verifies legal violations and simplifies the enforcement of existing regulations on a large-scale, independent from individual users' behavior.

4 QUESTIONING THE UBIQUITOUS PARADIGM OF CONSENT

In 1983 the German Federal Constitutional Court identified the fundamental right to informational self-determination in the context of modern data processing. The decision enshrined the users' right to decide if and how their personal data may be used. Since then, policymakers, privacy advocates, and privacy researchers alike focus on the issue of informed consent to data practices. Achieving informed consent requires that users (1) understand how their data will be used, and (2) agree explicitly and voluntarily to the use of their data. Both of these criteria have become increasingly difficult to achieve since the complexity of data practices has increased significantly, and privacy-infringing data practices have become so ubiquitous.

4.1 Informed consent to online data practices is difficult to achieve

Users have nuanced and context-dependent views on their acceptance of online behavioral advertising, understanding sites' need for monetization, and additional functionality enabled by tracking [4, 18]. However, they do not seem to fully understand that personalized ads rely on their tracked online behavior, how widespread these trackers are, which information these trackers can access, and what types of personal information advertisers can derive from their behavior [29, 31]. Users can, by definition, not give their informed consent to data practices they do not understand.

According to GDPR rules, user consent legitimizes most forms of privacy-infringing data practices. Hence, achieving a high rate of consent is a goal with direct monetary value for advertisers. Advertisers achieve high rates of "consent" by employing several different kinds of dark patters in the user interfaces of online consent notices [3, 8, 14, 17, 20]. Optimizing consent in this manner protects advertisers' business model while keeping up the facade of lawful and effective informational self-determination. Efficiently optimizing consent rates opened up a new business opportunity for so-called consent management platforms (CMP). "OneTrust, on its webpage presenting its CMP solution, proposes publishers to 'maximize user opt-in with customizable publisher-specific cookie banners [...] to optimize consent collection' " [15].

Aside from these fundamental issues with consent to online data practices, there is also the issue of the sheer number of consent decisions. Since these habituate and disengage users from consent requests in general, it will be necessary to reduce their number significantly.

4.2 Reduce overall number of consent decisions

The current state of the web requires a high number of consent decisions from users. Aside from an increased cognitive load, they also habituate users to accept more cookies than they feel comfortable with. To avoid user disengagement, these consent decisions need to be reserved for impactful, non-trivial, and context-dependent situations as Boehme et al. noted: "a last resort to prevent habituation is economizing consent decisions and thus reserving users' scarce decision capacity for the really important choices." [3] Many consent notices unnecessarily ask users to accept data practices that are essential for the site's functionality (such as virtual shopping carts or remembering the login status). We interpret this as a "proliferation of choice" tactic to overwhelm users since these cases do not even require consent according to GDPR. Additionally, there the consent decisions regarding tracking practices that users almost universally decline if they do not see a clear benefit to them.

Some tools that minimize consent notices already exist [11, 21, 23, 26, 32]. These apply predefined policies and thereby move away from individual consent decisions. However, these predefined policies do not necessarily care for privacy concerns and instead focus on the sites' functionality, e.g., by agreeing to all data practices to hide the consent notice.

Even if predefined decision policies handle most of the consent notices, there remains a small set of consent decisions that are non-trivial and can be very personal.

For these context-dependent decisions, users might want to weigh the site's perceived trustworthiness and the benefits of accepting behavior tracking (maybe additional functionality or monetary value) with the privacy-infringing drawbacks. In case the behavior tracking has a direct effect on the website's functionality (e.g., in the case of personalized search), the website should not be allowed to deny service, and instead, provide a gracefully degraded level of service to users who reject tracking. In any case, these remaining consent notices need a unified presentation so that users can understand and compare important points quickly [9].

4.3 Alternative approaches to improving privacy on the web

Achieving informed consent to data practices is not the turnkey solution to online privacy that we desire. The individualistic right to informational self-determination does not fight privacy-infringing behavior itself. Instead, it provides privileged, informed, and motivated individuals a seemingly neutral decision if and how they want to be under surveillance. As long as a sufficient number of users "consent" to surveillance, the underlying privacy-infringing business model continues. More regulation on how companies are not allowed to track users and ask users for consent could turn into an elaborate cat-and-mouse game. Hence, improving privacy on the web as a whole can not rely on informed consent alone.

One of the reasons for ubiquitous consent decisions on the web is that users' consent legalizes data practices that would otherwise undergo much more scrutiny. We can not rely on informed consent to improve privacy on the web. Hence, we need to make these privacy-infringing data practices more transparent, discuss the potential consequences, and use government regulation to keep

them in check – independent of users' "consent" decisions. A core issue for regulators should be the market for predictive futures since it drives the urge to collect ever-increasing mounds of behavioral data.

5 HOW FUTURE RESEARCH COULD COUNTER SURVEILLANCE CAPITALISM

As we discussed in detail in Section 3, prior research has invested significant work into describing, measuring, and improving issues revolving around tracking, targeted advertisements, and cookie banners. These works built the foundation for further discussions and raised awareness for the complexity of the advertisement industry and its stakeholders' conflicting interests. While identifying the oppressive mechanisms of the status quo is difficult already, changing is even more so. As discussed by Keyes [10] and Asad [2], prefigurative politics is a viable method of engaging with power. It is a pragmatic and applied process-driven approach to producing counter-power and counter-structures in the here and now. This section concerns itself with smaller-scale future research directions that could scale up to engage with the power behind current privacy issues to realize online privacy eventually.

5.1 Shift Focus Away From Consent

Users' consent has been a focal point of prior research, not least because it is a central aspect of privacy laws. However, as Trevisan et al. pointed out, the concept of consent alone might not be sufficient to solve the issues around behavior tracking [28]. Though consent is a necessary aspect, research can benefit from looking at the problem from a different angle. For example, future work can focus more on restricting the privacy-infringing data practices themselves without depending on the users' consent.

5.2 Examine Understudied Parties

So far, research has focused on end-users, CMPs, and the technical implementation of consent notices on websites. This focus has avoided some of the stakeholders involved in online tracking and advertisement. Consequently, the perspective of the advertisement industry and the content providers remain understudied. Of these two, the advertisement industry seems to use insights from research to maximize their profits by optimizing users' consent rates. Hence, we do not deem it worthwhile to invest further resources into understanding their point of view. However, understanding the perspective of content providers might provide valuable information to improving users' privacy on the web. One potential area of research is a detailed understanding of the content providers' motivation behind incorporating tracking mechanisms. Additionally, we consider it urgent to examine privacy-friendly forms of monetization and provide incentives to content providers to shift to them.

5.3 Individual and Collective Solutions

In our opinion, resolving the issues around tracking and cookie banners requires a combination of individual and collective solutions that address different aspects of the problem. Individual solutions target end-users. These can be plugins that help users detect tracking or make monitoring visible. Additionally, they can remove the burden of trivial consent decisions from the user and only call them to action when it is needed. Nowadays, we already have such tools, usually third-party plugins or tools that proxy the users' decisions. From the perspective of the websites, consent management platforms (CMPs), which handle the consent management for website owners, are on the rise. The rise of CMPs is somewhat counterintuitive, as the sole idea behind GDPR and similar laws is to make the process of tracking and data collection more transparent. However, it perhaps opened up business opportunities for even more parties to get involved with user data. Before privacy regulations came into effect, advertisers gathered people's data silently without their awareness. While ubiquitous consent notices made the data collection very annoying nowadays, they did not seem to have curbed it significantly.

We are in desperate need of collective solutions that target the issue on a large scale. For example, legal authorities need to be equipped with suitable measures to detect and report violations. Trevisan et al.'s CookieCheck tool is a step in the right direction [28]. Additionally, we are convinced that the different technological parties need to get involved in solving the issue, such as browsers, search engines, and social networks. One of the biggest players, Google, has a special kind of responsibility since they have one of the most popular browsers and are also heavily invested in the ad industry. Past efforts of Google in cooperation with other tech giants have demonstrated a strong ability to transform the Internet's ecosystem. For example, joined efforts resulted in a switch from HTTP to HTTPS, which at the time was not widely adopted [27]. They could also significantly contribute to an improvement of the current tracking situation if they wanted to. However, calling them into action will undoubtedly be difficult as such endeavors counteract their business model.

5.4 About Science Communication and Asking the Right Questions

We have been conducting the same research year after year – and while it is undoubtedly important to measure the status quo, we have not been able to counteract these privacy-invasive tendencies. There were also early efforts to propose design guidelines, driven by value-sensitive design, which are still relevant today [7, 19]. Unfortunately, many of today's practices violate these guidelines. We see similar tendencies also in other fields of research. For example, in the field of email encryption, essentially the same research is conducted over and over again, but at the core, there has been little progress over the last 20 years. We, as a scientific community, must try to change our perspective. If research's efforts are not being heard or do not have the desired effects, we need to ask ourselves two questions: First, are we properly communicating our results to the public and relevant authorities to make changes? Especially when research involves practical applicability and can therefore be fast-moving, we need to pay special attention to science communication. And secondly, are we potentially asking the wrong questions to begin with? Implicit assumptions for which we have no evidence and which may not be correct could, for example, mislead us to ask the wrong questions. Change is possible, but it is usually an arduous process that brings together the combined efforts of research, jurisdictions, and industry.

6 ACKNOWLEDGEMENT

We appreciate the fruitful discussion we had with Gabriel Grill (University of Michigan) on current and historic consent issues. He introduced us to the term "consent theater" which inspired this paper's title.

REFERENCES

[1] Alessandro Acquisti, Idris Adjerid, Rebecca Balebako, Laura Brandimarte, Lorrie Faith Cranor, Saranga Komanduri, Pedro Giovanni Leon, Norman Sadeh, Florian Schaub, Manya Sleeper, Yang Wang, and Shomir Wilson. 2017. Nudges for Privacy and Security: Understanding and Assisting Users' Choices Online. *Comput. Surveys* 50, 3 (Oct. 2017), 1–41. https://doi.org/10.1145/3054926

[2] Mariam Asad. 2019. Prefigurative Design as a Method for Research Justice. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (Nov. 2019), 1–18. https://doi.org/10.1145/3359302

[3] Rainer Böhme and Stefan Köpsell. 2010. Trained to Accept?: A Field Experiment on Consent Dialogs. In *Proceedings of the 28th International Conference on Human Factors in Computing Systems - CHI '10 (CHI)*. ACM Press, Atlanta, Georgia, USA, 2403–2406. https://doi.org/10.1145/1753326.1753689

[4] Farah Chanchary and Sonia Chiasson. 2015. User Perceptions of Sharing, Advertising, and Tracking. In *Proceedings of the Eleventh USENIX Conference on Usable Privacy and Security (SOUPS)*. USENIX Association, Ottawa, Canada, 53–67.

[5] 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 *Proceedings 2019 Network and Distributed System Security Symposium (NDSS)*. Internet Society, San Diego, CA, USA., 1–15. https://doi.org/10.14722/ndss.2019.23378 arXiv:1808.05096

[6] European Court of Justice. 2019. Declaration of Consent by Means of a Pre-Ticked Checkbox. http://curia.europa.eu/juris/document/document.jsf?docid=218462&doclang=EN. [Accessed: 2021-02-11].

[7] B. Friedman, D.C. Howe, and E. Felten. 2002. Informed Consent in the Mozilla Browser: Implementing Value-Sensitive Design. In *Proceedings of the 35th Annual Hawaii International Conference on System Sciences (HICSS)*. IEEE Comput. Soc, Big Island, HI, USA, 10. https://doi.org/10.1109/HICSS.2002.994366

[8] Stacia Garlach and Daniel D. Suthers. 2018. 'I'm Supposed to See That?' AdChoices Usability in the Mobile Environment. In *Proceedings of the 51st Hawaii International Conference on System Sciences (HICSS)*. IEEE, Hawaii, USA, 3779–3788. https://doi.org/10.24251/HICSS.2018.476

[9] Patrick Gage Kelley, Lucian Cesca, Joanna Bresee, and Lorrie Faith Cranor. 2010. Standardizing Privacy Notices: An Online Study of the Nutrition Label Approach. In *Proceedings of the 28th International Conference on Human Factors in Computing Systems - CHI '10 (CHI)*. ACM Press, Atlanta, Georgia, USA, 1573–1582. https://doi.org/10.1145/1753326.1753561

[10] Os Keyes, Josephine Hoy, and Margaret Drouhard. 2019. Human-Computer Insurrection: Notes on an Anarchist HCI. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. ACM, Glasgow Scotland Uk, 1–13. https://doi.org/10.1145/3290605.3300569

[11] Kiko. 2021. Firefox Browser Add-On – "I Don't Care about Cookies". https://addons.mozilla.org/de/firefox/addon/i-dont-care-about-cookies/. [Accessed: 2021-02-11].

[12] Oksana Kulyk, Annika Hilt, Nina Gerber, and Melanie Volkamer. 2018. "This Website Uses Cookies": Users' Perceptions and Reactions to the Cookie Disclaimer. In *The 3rd European Workshop on Usable Security (EuroUSEC)*. Internet Society, London, UK, 1–11. https://doi.org/10.14722/eurousec.2018.23012

[13] Oksana Kulyk, Peter Mayer, Melanie Volkamer, and Oliver Kafer. 2018. A Concept and Evaluation of Usable and Fine-Grained Privacy-Friendly Cookie Settings Interface. In *2018 17th IEEE International Conference On Trust, Security And Privacy In Computing And Communications/ 12th IEEE International Conference On Big Data Science And Engineering (TrustCom/BigDataSE)*. IEEE, New York, NY, 1058—1063. https://doi.org/10.1109/trustcom/bigdatase.2018.00148

[14] Dominique Machuletz and Rainer Böhme. 2020. Multiple Purposes, Multiple Problems: A User Study of Consent Dialogs after GDPR. In *Proceedings on Privacy Enhancing Technologies (PETS, Vol. 2)*. Sciendo, Berlin, DE, 481—498. https://doi.org/10.2478/popets-2020-0037

[15] Célestin Matte, Nataliia Bielova, and Cristiana Santos. 2020. Do Cookie Banners Respect My Choice?. In *2020 IEEE Symposium on Security and Privacy (SP)*. IEEE, San Francisco, CA, USA, 791—809. https://doi.org/10.1109/sp40000.2020.00076

[16] Aleecia M. McDonald and Lorrie Faith Cranor. 2010. Beliefs and Behaviors: Internet Users' Understanding of Behavioral Advertising. In *38th Research Conference on Communication, Information, and Internet Policy (TPRC)*. 1–31.

[17] Maryam Mehrnezhad. 2020. A Cross-Platform Evaluation of Privacy Notices and Tracking Practices. In *2020 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW) (EuroS&P Workshops)*. IEEE, Genoa, Italy, 97–106. https://doi.org/10.1109/EuroSPW51379.2020.00023

[18] William Melicher, Mahmood Sharif, Joshua Tan, Lujo Bauer, Mihai Christodorescu, and Pedro Giovanni Leon. 2016. (Do Not) Track Me Sometimes: Users' Contextual Preferences for Web Tracking. In *Proceedings on Privacy Enhancing Technologies (PETS, Vol. 2)*. Sciendo, Berlin, DE, 135—154. https://doi.org/10.1515/popets-2016-0009

[19] Lynette I. Millett, Batya Friedman, and Edward Felten. 2001. Cookies and Web Browser Design: Toward Realizing Informed Consent Online. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems - CHI '01 (CHI)*. ACM Press, Seattle, Washington, United States, 46–52. https://doi.org/10.1145/365024.365034

[20] 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)*. ACM, Honolulu HI USA, 1–13. https://doi.org/10.1145/3313831.3376321

[21] @opera. 2020. Twitter Message: "You Always Have a Choice with Us When It Comes to Blocking Cookie Dialogs. It's Also Available on Opera Touch for iOS Devices.". https://twitter.com/opera/status/1325870165033152513. [Accessed: 2021-02-11].

[22] European Parliament and European Council. 2002. Directive 2002/58/EC of the European Parliament and of the Council of 12 July 2002 Concerning the Processing of Personal Data and the Protection of Privacy in the Electronic Communications Sector (Directive on Privacy and Electronic Communications). https://eur-lex.europa.eu/legal-content/EN/ALL/?uri=CELEX%3A32002L0058. [Accessed: 2021-02-11].

[23] Permisio. 2021. Permisio – Customise Your Privacy Preferences in Just a Few Clicks. https://www.permisio.com/. [Accessed: 2021-02-11].

[24] 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 (AsiaCCS)*. ACM, Auckland New Zealand, 340–351. https://doi.org/10.1145/3321705.3329806

[25] Bruce Schneier. 2003. *Beyond Fear: Thinking Sensibly about Security in an Uncertain World*. Copernicus Books, New York, NY, USA.

[26] Umesh Shankar and Chris Karlof. 2006. Doppelganger: Better Browser Privacy without the Bother. In *Proceedings of the 13th ACM Conference on Computer and Communications Security - CCS '06 (CCS)*. ACM Press, Alexandria, Virginia, USA, 154–167. https://doi.org/10.1145/1180405.1180426

[27] Parisa Tabriz. 2018. Optimistic Dissatisfaction with the Status Quo: Steps We Must Take to Improve Security in Complex Landscapes. https://www.youtube.com/watch?v=py2qmGbyhlw. [Accessed: 2021-02-11].

[28] Martino Trevisan, Stefano Traverso, Eleonora Bassi, and Marco Mellia. 2019. 4 Years of EU Cookie Law: Results and Lessons Learned. In *Proceedings on Privacy Enhancing Technologies (PETS, Vol. 2)*. Sciendo, Berlin, DE, 126—145. https://doi.org/10.2478/popets-2019-0023

[29] Blase Ur, Pedro Giovanni Leon, Lorrie Faith Cranor, Richard Shay, and Yang Wang. 2012. Smart, Useful, Scary, Creepy: Perceptions of Online Behavioral Advertising. In *Proceedings of the Eighth Symposium on Usable Privacy and Security - SOUPS '12 (SOUPS)*. ACM Press, Washington, D.C., 1–15. https://doi.org/10.1145/2335356.2335362

[30] Verizon. 2021. Verizon Wireless' Use of a Unique Identifier Header (UIDH). https://www.verizon.com/support/unique-identifier-header-faqs/. [Accessed: 2021-02-11].

[31] Yaxing Yao, Davide Lo Re, and Yang Wang. 2017. Folk Models of Online Behavioral Advertising. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing (CSCW)*. ACM, Portland Oregon USA, 1957–1969. https://doi.org/10.1145/2998181.2998316

[32] Chuan Yue, Mengjun Xie, and Haining Wang. 2010. An Automatic HTTP Cookie Management System. *Computer Networks* 54, 13 (2010), 2182—2198. https://doi.org/10.1016/j.comnet.2010.03.006

[33] Shoshana Zuboff. 2018. *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*. Profile books, London, UK.