Beyond Acceptable Advertisement: Better Understanding Blocking Extensions

Benjamin VanderSloot
benvds@umich.edu
University of Michigan

Steven Sprecher
swsprec@umich.edu
University of Michigan

J. Alex Halderman
jhalderm@eecs.umich.edu
University of Michigan

ABSTRACT

Users have an uncertain relationship with online analytics and advertisement. This is understandable with the lack of public understanding about the practice and its model of consent. Given this, it is important to understand the defenses a user can employ to protect their privacy. In this work we find that blocking extensions vary significantly in their impact on user privacy, reducing average exposure to advertising and analytic domains by 9–23% in the worst case and 95–96% in the best case. Similar impact is observed in other extensions focused on empowering users to control requests from their browser, and can not be explained by Acceptable Advertisement programs or completeness of advertising and analytic domains blocked as was previously believed. Our findings on blocking list composition and efficacy in protecting privacy improve our recommendations to users and blocking tool developers.

ACM Reference Format:
Benjamin VanderSloot, Steven Sprecher, and J. Alex Halderman. 2019. Beyond Acceptable Advertisement: Better Understanding Blocking Extensions. In *Under Submission '20.* ACM, New York, NY, USA, 13 pages. https://doi.org/XX.XXXX/XXXXXXX.XXXXXXX

1 INTRODUCTION

Advertising is an important piece of the current web experience and economy. Most users' search engine is Google and its parent corporation, Alphabet, drives its 800 billion USD market cap by deriving 83% of its revenue from advertising [5]. Users have a complicated relationship with ads, and for good reason: despite being intrusive on their experience, ads can provide benefit if they are relevant [58]. It is in search of relevance by advertisers for being profitable, that many online advertisers target ads. However, underlying the targeted advertisements are user profiles built to enable the targeted ads.

The underlying profile building and tracking that supports advertisements is where the users can come to harm. Anecdotally, pregnancy status [36] and sexual orientation [35] have been disclosed accidentally without direct access to the user's profile. In one study, a majority of users were uncomfortable with any targeted advertising based on common interests inferred by advertisers. These same users were made even more uncomfortable by interests pertaining to their health or religion [21]. This says nothing of inferences that combine credit card history, voter registration, and browsing behavior used by journalists specifically to demonstrate the power of targeted advertising [56].

In lieu of a strong model of consent, some users have turned to using browser extensions to prevent being tracked. This is not to be confused with the separate but related phenomenon of ad blockers [44]. It is easy to confuse the categories of blocking extensions: ad blockers advertise their utility in preserving privacy [3, 26] and use some of the same rules to block requests.

Researchers have studied several aspects of these browser extensions and the block lists that direct them how to behave. This even includes some comparative evaluation of these extensions to inform recommendations to users. However, these comparisons do not do enough to look into the factors that contribute to the extensions' effectiveness, leaving us asking the central question of our work: what are the roles of manual list creation, organizational mission, and the connections explicitly white-listed from blocking in the limits of ad blocking extensions as deployed today?

In this paper we look to understand blocking extensions better and unpack these factors as they impact extension efficacy. To do this we visit a list of the top million sites with Firefox. Our model of the blocking extensions behavior is more accurate than prior work. We also compare browser extensions' blocking behavior over all rules defined in each of their rule-sets, reducing reliance on crawls limited in scope and subject to transient advertiser practices.

We infer how many page visits would be disclosed to tracking companies over our crawl, and use this information to evaluate the browser extension efficacy. Browser extensions are evaluated both with and without their Acceptable Ads programs enabled, and find results that contradict those of prior work: Acceptable Ads programs are not the primary reason for their extensions' poor performance.

We use the inclusion graph of our crawl to identify advertisers not in the hand-curated lists used by blocking extensions. The size of these newly discovered advertiser lists help us understand the completeness of the manual curation of the lists over their target communities. When combined with our understanding of Acceptable Ads programs and the stated missions of the browser extension authors, this lets us better understand the factors that contribute to extension efficacy.

Our work answers 3 questions that had not been previously answered. *1)* How effectively do the full population of A&A companies track users across the web? *2)* What tools are most effective in stopping this tracking? *3)* What factors of the effective tools most contribute?

Under Submission, 2020,

ACM ISBN XXX-X-XXXX-XXXX-X/XX/XX.
https://doi.org/XX.XXXX/XXXXXXX.XXXXXXX

2 RELATED WORK

In this section we describe the research that comprises the background of and is related to our own work.

One subarea of web tracking research is the human interaction with web trackers and defenses. Some work has focused on user understanding and perceptions of targeted ads [16, 21, 45, 58, 63], finding significant gaps in knowledge and inconsistent user mental models. More closely related to our work, research has focused on user adoption of both ad blockers [51] and blocking extensions more generally [44] finding diverse reasons for adopting these extensions and little issue with website breakage. Finally, the justification for inclusion in the Acceptable Ads program by AdBlock Plus was refuted [62].

Some studies have focused on cookies, one of the most conspicuous means of identifying users uniquely. Much of the research has looked at identifying unique identifiers in cookies and how third parties share these identifiers to synchronize identities across multiple tracking networks [12, 28, 30, 49, 53, 65]. Some work in this area has focused on policy, specifically considering implications of cookies on governmental mass surveillance [24] and EU privacy law [18, 57, 59].

Much research has been performed to understand the behaviors, weaknesses, and extent of the advertising ecosystem [8, 10, 13–15, 31, 42, 50, 52, 61, 64]. In particular, some work focused on out-of-band exchanges of information between advertisers [9] and challenges in attributing advertisements [40]. These show limitations in fully automated, but much larger measurement techniques such as our own. Also, some research has analyzed these same questions for defensive countermeasures used by ad networks to circumvent blockers [37].

We draw much inspiration from the work of Bashir et al. [11]. In particular, the view of the entire advertising ecosystem as a graph, similar to that of Kalavri et al. [38] and Gomer et al. [32], shaped our thinking. Kalavari et al., however, focused on automatic identification of trackers using machine learning to replace current blocking extensions, supporting this by analyzing the referrer graph structure. Gomer et al. studied a different graph representation of the tracking ecosystem, creating a bipartite graph of first-parties and the third-parties they load, rather than the structure of the third-party ecosystem itself. Bashir et al. study the same graph that we do and compare some blocking extensions. However, by focusing our study on the blocking extensions and by studying a different, more diverse, and larger population of web pages with a more accurate blocking extension model we were able to come to deeper conclusions about blocking extensions and refute proposed explanations for findings of that work.

The closest research to our own is in the evaluation of existing solutions to online tracking. Early work identified the futility of the DNT header [7]. Other studies focus on the interaction of ad-blocking tools and advertisers in an arms race of detection [46, 48]. However, these fail to account for out-of-band communication between A&A companies and did not explore causes of ineffective blocking as deeply as in this work, also identifying Acceptable Ads as the reason for AdBlock Plus's poor performance. The work of Malloy et al. [43] focuses on the market penetration of ad blockers, and focuses on impact of advertisements displayed, not tracking. The work of Karaj et al. [39] and the work of Yu et al. [64] focus on the users of a single browser extension which, while very useful, is biased in uncharacterized ways and did not try to generalize beyond their primary study population. Gervais et al. [29] and Bashir et al. [11] perform less focused studies over a very small number of popular websites. We demonstrate in this work that our population of trackers discovered is larger than that measured in these papers and that tracking is quantitatively different from the top thousand to the top million (which are equally weighted under typical Zipf-ian assumptions). Merzdovnik et al. [46] provides a focused analyses of a large-scale top domain list. However, they do not study the A&A ecosystem outside of 30 actors and additionally fail to analyze how information is shared in an RTB system. This work and that of Bashir et al. [11] show that out-of-band RTB sharing has significant impact on user privacy. However, they do not explore deeply enough which tools are more effective at stopping tracking and why, only analyzing one ad blocker and one privacy tool, and also concluding incorrectly that Acceptable Ads are the reason for AdBlock Plus's poor performance. Bashir et al. [11] also used a less accurate representation of the acceptable ads list and a data set focused on popular online shopping sites, which led to a limited scope of A&A websites. Walls et al. [62] studied the impact of acceptable advertisement lists, but focused on the composition, evolution, and user impressions of the list; this omits analysis of the impact of the list on user privacy. Vastel et al. [60] studies filter list composition, but pursues questions on how to trim filter lists of rarely used rules. This did not consider the impacts to user privacy due to out-of-band communication by A&A companies or the most effective trackers and did not seek to explain why different tools were more effective. This insights of our work are complementary to those of Vastel et al.for these reasons.

3 METHODS

To support the analyses in the remainder of this paper we employ methods we develop ourselves as well as methods taken from prior work. We present the methods in this section, referencing them in later sections as they are used.

Throughout we use the acronyms FQDN and eFLD. These are the Fully Qualified Domain Name (e.g. `www.foo.co.uk`) and the effective First Level Domain (e.g. `foo.co.uk`).

3.1 Web Crawl

Prior work has studied many websites' tracking behavior using OpenWPM, an instrumented browser measurement platform [23]. We re-use this tool and its standard data formats, allowing re-use of our analyses over future and past crawls. To facilitate this, we release our source code and data at [redacted for anonymous submission].

Also like prior studies, we crawl the homepages of a large top domain list. We use the Tranco list [1] [41] created on 07 July 2019. This improves reproducibility while maintaining the scale of the largest prior work and reducing bias from domain fluctuations known to be in the more commonly used Alexa top list [4, 54]. We compare the top crawl targets and advertising and analytic domains

[1] Available at https://tranco-list.eu/list/J96Y.

Table 1: Comparison of Top Domain List to Prior Work. We compare both the top ten in each list, accounting for 20% of Zipf-weighted visits, and the top ten eFLDs blocked by EasyList and EasyPrivacy as trackers.

	Alexa (Jan 2015)		Tranco (Jul 2019)	
	Targets	Trackers	Targets	Trackers
1	google.com	google-analytics.com	google.com	doubleclick.net
2	facebook.com	doubleclick.net	netflix.com	google.com
3	youtube.com	facebook.com	facebook.com	facebook.com
4	baidu.com	googlesyndication.com	youtube.com	googlesyndication.com
5	yahoo.com	googleadservices.com	twitter.com	google-analytics.com
6	amazon.com	google.com	microsoft.com	youtube.com
7	wikipedia.com	twitter.com	wikipedia.org	facebook.net
8	qq.com	adnxs.com	baidu.com	pubmatic.com
9	google.co.in	blekai.com	linkedin.com	adnxs.com
10	twitter.com	mathtag.com	instagram.com	wp.com

contacted in our study to the most recent publicly available Alexa crawl [23] in Table 1.

We attempt to load the HTTP version of all one million domains with OpenWPM in a stateless crawl using Firefox, storing HTTP request, response, and redirect data. We performed these measurements from a dedicated research scanning subnet from 12 July—27 July 2019. Of the 1M input list, OpenWPM received responses for 895,856 sites, missing 10.5% due to timeouts, unavailable destinations, and various other errors; this error rate is similar to prior measurements of top million sites. This results in 89,136,479 total requests during the crawl to 62,970,931 unique urls.

3.2 Blocking List Testing

We compare six different blocking extensions, selected for their popularity: AdBlock Plus [26], AdGuard [3], Disconnect.me [20] (whose rules are used by Firefox in its tracking protection [47]), DuckDuckGo [22], Ghostery [17], and uBlock Origin [34]. We evaluate each tools' default configuration; however it is possible for users to improve the efficacy of each extension and even import filter lists from other extensions in most cases. We also use EasyList and EasyPrivacy [55] throughout our evaluation to identify previously known A&A domains. EasyList and EasyPrivacy are community created lists for blocking advertisements and tracking respectively, and are imported by some tools to verify We define *previously known A&A domains* heuristically as in prior work [11]: any domain that is blocked more than 10% of the time by EasyList and EasyPrivacy. We use blocking lists retrieved on April 23rd, 2019 for all extentions.

Our method of evaluating blocking lists is an improvement over prior work, as we show in Section 6.4. We more closely reflect the actual behavior of the blocking extensions while still testing over 29M requests. We enabled this in two pre-computation stages: normalizing each tool's behavior to a single syntax and reducing sets of regular expressions into single prefix-structured regular expressions.

3.2.1 Normalizing Blocking Syntax. Before we take any steps to speed up our request testing process, we first standardize our blocking list syntax. Fortunately, three extensions (AdBlock Plus [27], AdGuard [1], and uBlock Origin [33]) of the six we study and EasyList and EasyPrivacy use a very similar syntax. While these syntaxes vary, the commonly used features are shared.

The syntax for blocking rules we normalize to is a shared subset of these powerful syntaxes. Each rule is a newline-delimited string with two parts, separated by a $ character. The first part is a simplified regular expression syntax and the second part is a series of flags and corresponding values. The regular expression syntax has only five special strings that are not interpreted literally: "`||`", "`|`", "`*`", "`^`", and "`@@`". Each of these have the same meaning among all of the tools and either have a direct translation to standard regular expression syntax or indicate that the rule should be inverted, specifying requests to allow rather than to block ("`@@`").

The flag portion of these powerful syntaxes is often overlooked; it is not mentioned at all in the application of rules in prior work. In fact, the syntax for AdBlock Plus, AdGuard, and uBlock Origin vary primarily in what flags they support. We support the following flags: `important`, `domain`, and `third-party`. Respectively, these rules bypass whitelisting caused by other rules, enforce the rule only on requests originating from specific domains, and enforce the rule only on third-party requests. The only rules we ignore that could be used to aid privacy and are used in more than 1% of rules are those that block requests for types of resources, i.e. `script`, `image`, etc.. This leads to over-estimation of the power of the blocking tools, however the tools are treated equally.

Our implemented subset of these rules also elides an entire class of specification: cosmetic rules. These rules mark elements on the page as `hidden` and do not prevent requests from being sent. Therefore, we find it acceptable to omit them entirely for measurement of web tracking.

Tools that do not use this standard syntax are converted into the standard syntax. We take each of the specified rules and convert them into an identical rule that blocks exactly the same resources and subdomains, using only the limited regex syntax and `domain` flag. These rules can then be treated as the other canonicalized rule lists are in the subsequent steps.

3.2.2 Efficiently Evaluating Regular Expressions. During initial experiments it became clear that evaluation of all rules over all requests of our crawl performed naively would be impractical. Evaluating 89M requests against over 130K regular expressions proved to be a significant bottleneck, even with early-exit optimizations. To speed up our evaluation, we reduce the number of comparisons required in a few ways that are layered hierarchically.

First we note that some rules are only in use on very few domains. To take advantage of this, for each tool's rule list we create rule lists for each domain explicitly mentioned in the constraints. Those rules we do not have any constraint on, or only whitelist explicit domains from having the rule applied, are added to the default collection which is applied to all domains. In effect, we have created a blocking list for each referring origin that would change the behavior of the blocking tool. These lists being much smaller allows much less comparison in an initial check that may return a result quickly. If a result is not obtained from the first check, we then compare on the larger list that contains rules that apply to all referring origins.

Next we note that there is a priority order of rules that goes from a small subset to a large subset as the priority decreases: another opportunity for early exit. First, we can search in the set

of "important" rules that bypass whitelisting, and are the smallest set. Then we can look at the whitelist rules, those that begin with "@@". Finally we compare the remaining rules. To allow this ordered evaluation, at this stage we further partition the rules along these lines, within each domain constraint's list.

Finally, and perhaps most importantly for our speed, we reduce all of the regular expressions in each partition into a single regular expression. Doing this naively, simply listing each expression, and separating them by '|' in standard regex syntax, would provide little speedup over comparing each rule to the URL. This is the exact underlying behavior the regex engine takes in its evaluation. Instead, we convert the rules into a prefix tree of their equivalent regular expressions, then convert this prefix tree into a reduced regular expression. For example, rather than converting a search for either "food" or "fold" into "(food|fold)", this is converted into the regex "fo(ld|od)", removing duplicated comparison to the common prefix "fo'. This takes ideas from prior work in optimizing regular expression validation [6], without having to write a custom engine.

So, in review, to compare a single URL requested from a given origin, we take the following steps: first find the longest FQDN with domain constraints for which our referring origin is a subdomain and check if it is blocked or explicitly whitelisted by that structure. If not, or if there is no such subdomain, check if it is blocked or explicitly whitelisted by the rules with no domain constraints. Both of these checks are done by evaluating the next layer of the hierarchy: comparing against the important rules, whitelist rules, and other rules of that structure in that order. Again, to do this we proceed down the structure, comparing against the regular expressions built for this constraint and priority, and returning whether or not a match was found. Finally, if the result is that we found no matching rules, we interpret this as an unblocked request. In this way, we have traversed only the relevant partitions of the rules we created for this blocking tool.

While this is more complicated, it does capture the entirety of our normalized subset of the blocking syntax. Additionally, it sped our evaluation of requests from our crawl by *over 1000x*, letting us finish evaluation of all six tools in under six hours on a single 8-core Intel Xeon E5–2690. This made the comparison of multiple tools over a very large crawl practical.

3.3 Blocking List Probing

Initial analysis of how each tool performed on our crawl revealed that most rules blocking entire domains were not used at all during our crawl. We want to present an alternative perspective when comparing tools that is not biased by the particular crawl. To do this we created probing lists from each tool's blocking rules.

Each blocking rule that includes a specified domain name is mechanically converted to a set of rules that capture some of the cases a rule may apply to. Multiple transformation may occur for each special character, causing exponential growth in the number of special characters. In practice each probe list does not grow impractically large.

Rules that begin with the special character "|" simply have the character removed. Rules that begin with the special sequence "||" have the sequence replaced by both "http://" and "https://".

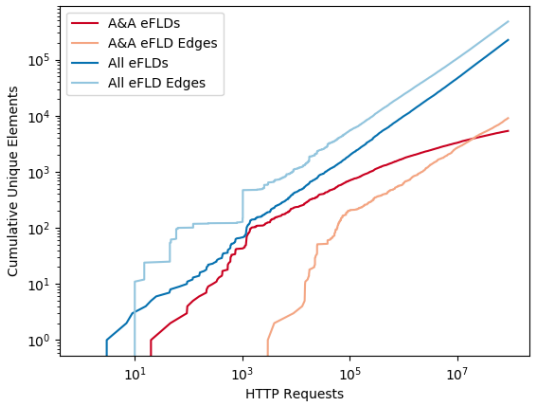

Figure 1: Cumulative Unique Inclusion Graph Nodes and Edges. We graph the cumulative number of unique eFLDs (dark lines) and inclusions between eFLDs (light lines) as we conduct our crawl. Blue lines represent the graph of all third-parties and red lines represent the subgraph of only third-parties that are known A&A domains.

The special character "^" is replaced by a "/" and optionally removed if at the end of the rule. Finally, the "*" is replaced by an empty string, "1", "/", and the random string "b75de6c2/2ce0bae5" which occurs nowhere in the block list definitions. These simple transformations capture much of the diversity of behavior of the limited syntax provided by the blocking extensions. For example, the rule "||foo.com^" is expanded to the following probing links: "http://foo.com/", "http://foo.com", "https://foo.com/", "https://foo.com".

Rules are tested with referring origins taken from the domain restrictions and an unrestricted referring origin. We group the probes by eFLD of the request and compare two test lists by checking if all test probes in the same eFLD have the same outcome.

3.4 Inclusion Graph

Prior work has demonstrated use of the inclusion graph to demonstrate interrelationship among known A&A domains [11]. The inclusion graph is a representation of a web crawl where parties (in prior work canonicalized domain name) are node and directed edges indicate that the source party initiated a request to the destination party in some page visit of that crawl. This is in contrast to the referrer graph where an edge indicates the source appeared as the referrer in the destinations' request. JavaScript requests are a source of disagreement between these two graphs and are used often in the A&A ecosystem. We construct the inclusion graph similarly, however we our graph is over all third-parties, rather than only known advertisers. We are also able to analyze the subgraph of only known advertisers for direct comparison.

The graph we produce represents only the sample of domains we perform; there may be some number of nodes or edges not included in our sample that exist in the population. We can tell if the number

of nodes and edges continue to increase as we near the end of our crawl. We show the growth of these through our crawl in Figure 1.

Our crawl does not saturate the number of eFLDs nor edges between them from our population. This is despite finding a larger number of them than prior work that found saturation in its inclusion graph. This is because we study a much larger population that is not restricted by type of domain or to only the top few thousand domains in a rank list.

This does not preclude use of the graph for the uses in this paper, however we do not compute graph statistics to compare to prior work as they would not be representative.

3.5 Privacy Metrics

Finally, we analyze the impact of blocking strategies on our crawl. Prior work has defined disclosure of a page visit, sometimes termed impression, and we use these definitions. In particular, an A&A party learns that a particular user has requested the main page of a particular visit if in the course of that visit it is sent a request. Additionally, an A&A party may learn of a request if it is shared with them through other means. We consider two such means in this work[2]: **Cookie Matching** and **Real-Time Bidding**.

In the **Cookie Matching** model, we assume information sharing along all directed edges known to perform ad-redirection [9].

In the **Real-Time Bidding** model, we attempt to capture information disclosure that occurs when ad exchanges solicit bids from Demand Side Platforms and advertisers for a given impression. To do this, we designate some eFLDs as belonging to ad exchanges, using heuristics described in Section 5, and have these exchanges disclose all visits they are aware of to all eFLDs they include during our crawl. We investigate the sensitivity to the number of ad exchanges included in Section 6.1.

The metric we use to evaluate any A&A domain in our crawl is how many page visits cause the domain to receive a request or be further notified by another third party. To measure if a domain would receive a request, we build the inclusion graph for just that page visit and compute the reachable origins if only non-blocking requests and out-of-band disclosures proceed.

Performing this inclusion graph construction and evaluation of sharing is also expensive to do at scale. To cope with this we only perform these computations over a random sample of 54,000 visits, weighted with a Zipf-ian distribution over their top site rank to approximate user traffic patterns. We graph several cumulative distributions of this metric over previously known A&A domains. By the Dvoretzky–Kiefer–Wolfowitz inequality this places the entire CDF within the line width on our figures with 95% likelihood, even with Holm-Bonferroni correction. As such we present the distributions without explicit error bars.

4 COMPARING BLOCKING LISTS

In this section we lay the foundation of our new understanding of blocking extensions by comparing the blocking lists that define the behavior of these extensions. We evaluate the extensions' behavior both over our web crawl and our constructed probing lists.

[2]Both of these are taken from Bashir et al. [11], corresponding to Cookie Matching and Real-Time Bidding – Constrained respectively

4.1 Default Extension Behavior

Blocking extensions rule lists have been compared to each other in prior work in their rate of blocking and the privacy that affords. However, some aspects of our comparison are novel. Our comparisons are used to help disambiguate the impacts of different factors on the privacy provided by these tools. Also, our evaluation uses rule lists to generate test lists providing a larger perspective and one that lets us ignore issues of the coverage of a single crawl. Finally, our data set enables insight into a larger test list than prior work focused on blocking extensions, and the first top-million study of web tracking in four years.

An important distinction must be made when evaluating the size of lists is whether the measure used is one that is evaluating the list itself or is grounded in a sample of requests. We compare list sizes in Table 2 using both.

The largest lists by size and by domain explicitly mentioned in their rule set are AdBlock Plus and AdGuard, however it is not the case that these most active filtering tools based on our crawl. AdGuard blocks the second fewest requests to a relatively diverse set of domains while one of the smallest lists, Disconnect.me, provides the most requests blocked. This disagreement highlights a difference in strategy between growing a large test list that does not block entire domains as third parties and entirely blocking a few major actors.

While comparing the list sizes by measure does tell us some differences between the tools, it does not show the intersections of each tools' behavior. To do this, we must compare this tools pairwise. Understanding the relative intersections of tools exposes which tools may source their rules from one another and how independent other tools are in their behavior.

There are subtleties in comparing different tools under metrics that can be described other than "blocked" or "not blocked". For example, comparing the domains that two tools block can be done by taking the domains mentioned explicitly in each tools' test list, however test lists can hold exceptions for many domains and may have more generic rules that do not explicitly mention all of the same domains. To counter this, we can compare the domains that are blocked under at least the same conditions as the other test list. This can be done using our probing lists described in Section 3.2. Note that this comparison is asymmetric, e.g. only Tool 1 may block `example.com/foo` and only Tool 2 may block `example.com/bar`.

In Figure 2a we compare tools by which domains where tools' test lists explicitly block any resource from a given domain and in Figure 2b we compare the relative coverage of those lists using our generated probe sets from Section 3.3. We include EasyList and EasyPrivacy as a baseline as it is used in prior work as a baseline.

In both figures the vertical bar corresponding to uBlock Origin's performance on their peers' test lists show the most consistent coverage. It has the strongest definitions when evaluated on other tools' probe lists, with the exception of Disconnect and DuckDuckGo's similar definitions. uBlock Origin shows similar behavior to AdGuard and AdBlock Plus: all three appear to be derived from EasyList with varying degrees of difference. Notably, uBlock Origin is the only that gets stronger when diverging from the EasyList definition, observed by comparing squares reflected over the diagonal line that is all black.

Table 2: Tool Block List Sizes. We compare six browser extensions' block list sizes and coverage in our crawl data set. Our counts are by rules after canonicalization, explicitly mentioned eFLDs in rules, requests blocked, and the number of eFLDs with at least 10% of requests blocked.

	AdBlock Plus	AdGuard	Disconnect.me	DuckDuckGo	Ghostery	uBlock Origin
Canonicalized Rules	50,844	51,552	2,568	1,789	4,667	65,772
Explicit eFLDs	28,529	31,427	2,117	1,736	3,658	37,912
Blocked Requests	9,580,925	12,276,925	30,840,615	20,073,773	24,019,804	25,692,118
A&A eFLDs	2,545	2,981	1,340	943	2,190	6,115

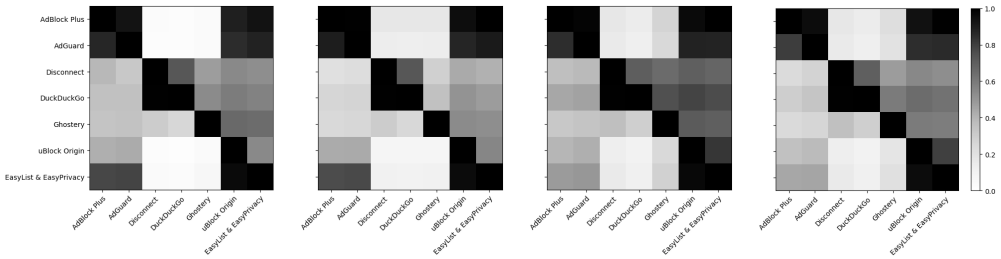

(a) Definition - Any **(b) Definition - Covered** **(c) Measured - Any** **(d) Measured - Covered**

Figure 2: Comparing Extensions Pairwise. We compare each extension to all others, and EasyList and EasyPrivacy in four different ways. The comparisons are made over both the definition of the rules and how they behave in our crawl (Definition vs Measured) as well as if the tool blocks the domain at all versus if the tool blocks the domain in at least the same instances as the tool it is compared to (Any vs Covered). For example, the entirely black square in the upper right corner of *(c)* means that for every eFLD we observed a request blocked by AdBlock Plus, we also observed a request to that same eFLD blocked by EasyList and EasyPrivacy.

Additionally, we observe that Disconnect and DuckDuckGo have similar list definitions, with Disconnect's definitions being stronger. It is not explicitly stated that one derives from the other, however this is one simple explanation.

While these comparisons describe the broadest possible coverage of each list, we are not able to infer how these lists compare on resources from real websites. To do this, we must utilize our crawl data set. We perform similar comparisons with this practical data: Figure 2c shows how many domains have at least one domain blocked by both tools and Figure 2d shows the number of domains that are blocked in at least the same cases for a pair of tools. Both values are taken as a fraction of the number of domains blocked ever by the baseline tool.

In our crawl, the strongest tool, uBlock Origin, performs even better, as does Ghostery. Overall our observations from the comparison of list definitions are borne out in our crawl data.

The smaller test lists by rule count perform poorly against the domains blocked by larger test lists in these figures. We count by eFLDs, so this does not reflect total requests blocked, nor does it reflect the privacy impact of blocking each domain. In Section 6, we will analyze the privacy benefits of these tools.

4.2 Explicit A&A Exceptions

Some blocking extensions are straightforward in their effort to reclaim user experience and privacy without consideration for the advertising and analytics companies. Others explicitly exempt some requests as acceptable, despite their known analytic or tracking nature. While these exceptions are made for a variety of reasons, we compare them directly to determine their impact. To clarify, we do not include all rules that white-list requests; their exclusion may be because they are not related to tracking despite being caught in another rule, to circumvent anti-ad-blocker tools, or to allow an analytics company to perform analytics on its own domains. These are all examples of what is not exempt A&A behavior.

In order to understand the impacts of the acceptable advertisement lists, we repeat some of our analysis in the previous section, removing the explicit exceptions. Only AdBlock Plus [25], AdGuard [2], and Disconnect [19] claim to have these explicit exceptions.

Table 3: Tool Block List Size, Explicit Exception Impact. We show the changes to the scope of each blocking tool when their explicit exceptions are removed. Counting is performed identically to that of Table 2.

	AdBlock Plus	AdGuard	Disconnect.me
Canonicalized Rules	-8,468	-30	+512
Explicit Domains	-1,892	-10	+4
Blocked Requests	+2,365,442	0	+55,248
A&A Domains	+26	0	+3

We show differences in list size in measures affected by the removal of explicit exceptions. Table 3 shows that only AdBlock Plus is affected significantly by their "Acceptable Advertising" list. AdGuard showed only 30 rules that reference 10 domains removed, and none were utilized during our crawl.

Disconnect.me showed effectively no change in its behavior with and without its exception list removed. The default behavior of Disconnect.me's browser extension and the content of its publicly available block list include those sites it claims it does not block in its list. This means that when we introduce the domains claimed to be excluded from their default list, we gain 4 new rules and 508 duplicate rules. The number of rules went up as an artifact of how we constructed the test list; our default test list includes only those that ship with the browser extension, with no care taken to exclude the alleged exceptions. This resulted in 3 new domains being blocked in our crawl and a 0.2% increase in blocked requests.

If AdBlock Plus were to block its "Acceptable Advertising" it would increase the requests blocked by 20% and while only increasing the domains that have 10% of requests blocked by 1%. These eFLDs are all blocked in over half of requests by EasyList and EasyPrivacy. Included in these are eFLDs with more than 50,000 requests: `yahoo.com`, `media.net`, `linkedin.com`, and `bttrack.com`.

We do not find agreement with prior work on the scope of the Acceptable Advertising list in our more precise model. Prior work [11] included major ad networks in the domains excluded by Acceptable Advertising, however we find that this is not the case in practice. This is due to our more faithful representation of the block lists, in particular implementing domain restrictions; the ad networks do show up in our Acceptable Advertising list, however the vast majority are restricted to being whitelisted on a small number of websites. Two notable exceptions are Google AdSense[3] and subresources of some advertisers used to monitor advertisement clicks[4], which are permitted on all but a few sites. Despite these broad exceptions and their potential for misuse by advertisers, ad networks are still blocked by AdBlock Plus because of the broad restrictions on most exceptions. This refutes a previous belief that AdBlock Plus performs poorly *because* of the Acceptable Ads program.

4.3 EasyList Variants

Three tools we test are at least in part based upon EasyList or EasyPrivacy: uBlock Origin, AdBlock Plus, and AdGuard. Notably, uBlock Origin is the only of the three that adds blocking behavior to EasyList and EasyPrivacy combined. It adds separate resources for malware risks, privacy, and resource abuse above its bulk import of EasyList and EasyPrivacy.

Both ad blockers are based upon EasyList, but do not import EasyPrivacy. However, EasyList blocks several ad exchanges that are allowed by the ad blockers. AdBlock Plus directly imports EasyList, then adds exception rules for some ad exchanges and websites. AdGuard, on the other hand, imports EasyList, then applies changes to form a single modified list. Their behavior blocking ad exchanges compared to EasyList and EasyPrivacy is shown in Figure 4.

[3]This change was not discussed in a public forum: https://hg.adblockplus.org/exceptionrules/rev/c182d9dc5600

[4]This change was discussed in a public forum, but at a high level and without consideration of potential abuse: https://adblockplus.org/forum/viewtopic.php?p=179030

Table 4: Exchange Comparison on Easylist Variants. Here we show the reduction of blocking the top ad exchanges from Easylist and Easyprivacy for the Ad Blockers. Both ad blockers show significantly reduced blocking of most ad exchanges, even though they are based upon Easylist. Neither ad blocker incorporates Easyprivacy. Percentages indicate the percent of requests blocked by EasyList and EasyPrivacy that were not blocked by the tested tool (0% indicating no change and 100% indicating no blocking).

Domain	AdGuard Reduction	ABP Reduction
doubleclick.net	28%	41%
googlesyndication.com	3%	37%
pubmatic.com	0%	52%
adnxs.com	0%	0%
openx.net	0%	46%
rubiconproject.com	0%	46%
2mdn.net	1%	0%
yandex.ru	100%	100%
amazon-adsystem.com	0%	21%
bidswitch.net	100%	100%

Some ad exchanges, like `adnxs.com`, show no change in behavior in ad blockers while others, like `yandex.ru`, show complete whitelisting. More interesting are those that show partial reduction of blocking. We observe through manual inspection that the vast majority of the partial reductions are to allow cookie matching behavior by loading empty content or tracking pixels. AdBlock Plus allows more of these tracking pixels to proceed from a more diverse set of sources.

5 BLOCKING COMPLETENESS

Most current blocking extensions build their rule sets of what to block through manual effort. In this section we quantify the completeness of these efforts for different blocking extensions. Naively, we could look at all known advertisers and treat this as each extensions' ideal, however this fails to account for the different subsets of the A&A ecosystem each extension wishes to block. Instead, we identify ad exchanges blocked by each extension and label the eFLDs that receive at least 10% of their connections from these extensions as advertisers for that extension.

Prior work using inclusion graphs has identified ad exchanges using heuristics on the ratio of indegree and outdegree and outdegree alone of nodes in the A&A inclusion graph. Ad exchanges must have large outdegree to other advertisers to be effective. They also have approximately equal indegree and outdegree. Too large an indegree indicates a tracker that only ingests information without selling impressions, while too small an indegree indicates an SSP that is not purchasing ads from a large number of platforms. We look at the values of outdegree and degree ratio for known ad exchanges and determine that an outdegree ≥ 50 and in/out degree ratio $\in [.5, 2.0]$ was representative. For EasyList and EasyPrivacy, this labeled 47 ad exchanges, enumerated in Appendix A. This is performed over the subgraph of our inclusion graph that only includes A&A eFLDs.

Table 5: List Coverage Quantifiaction. We show the results of quantifying missed A&A subdomains over the tested extensions using our exchange in-degree metric from Section 5. Note that this does not quantify privacy loss; this does not account for the blocking that prevents further requests from being performed.

List	Known A&A eFLDs	Discovered A&A eFLDs
AdBlock Plus	2,545	3,853 (151%)
AdGuard	2,981	3,893 (131%)
Disconnect.me	1,340	4,397 (328%)
DuckDuckGo	943	3,277 (347%)
Ghostery	2,190	5,397 (246%)
uBlock Origin	6,115	4,289 (70%)

With ad exchanges labeled for each extension, we then identify the eFLDs that are advertisers. We note that we labeled eFLDs as advertising or analytic if they were blocked more that 10% of the time. Since connections to ad exchanges would be blocked, preventing them from including any further domains, we look at the percent of connections each eFLD receives from ad exchanges. Therefore, we label all eFLDs with at least 10% of their connections coming from ad exchanges as advertisers. The results of this process are show in Table 5.

Note that while this quantifies coverage of the block lists, a lack of coverage does not inherently mean that privacy is adversely affected. We found these advertisers by identifying domains that *would have been blocked* by the tool in some cases by virtue of blocking ad exchanges. However, blocking the advertisers themselves can be useful to block their direct inclusion from publishers or otherwise unblocked members of the A&A ecosystem.

In Table 5 we see that the lowest coverage is observed by the smallest lists. These lists, on manual inspection, seem to focus on the larger players in the analytics ecosystem and are focused in their mission to provide privacy. A focused effort on the actors that share information the most seems to be an efficient use of manual effort; in Section 6 we will quantify the efficacy of this effort on privacy.

The best coverage of domains is provided by uBlock Origin. uBlock Origin curates its block list through a large community effort. This effort shows in the size and diversity of eFLDs blocked by its list as shown in Table 2.

6 PRIVACY IMPACTS

One aspect of the comparison of different blocking extensions in Section 4 is evaluating the number of blocked requests in our web crawl. While this does provide a rough proxy for privacy offered by the extension over the crawl, it is inaccurate for two main reasons. First, it does not evaluate how much of the tracking information is shared with each individual domain or out-of-band information sharing, such as that performed by ad exchanges. Second, it does not account for the impact of blocking one request that would otherwise lead to many other requests. We use an approximation of what impressions are learned by each advertiser to take these factors into consideration and more accurately reflect privacy violations.

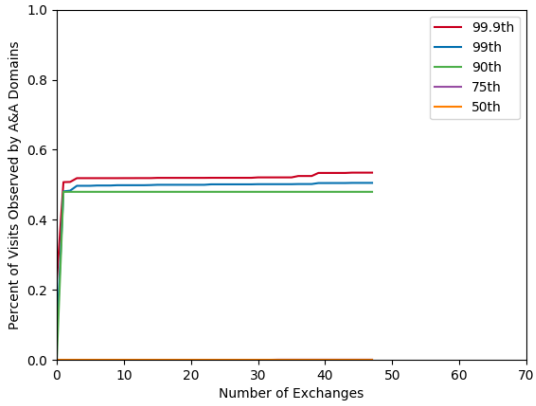

Figure 3: Ad Exchange Count Sweep. We show the degree of sensitivity to the number of ad exchanges on various percentiles of vist leakage. We add brokers in the order of the number of advertisers the initate connections to in our web crawl. Each line represents some Nth percentile of success in observing as many page visits as possible over our web crawl.

6.1 Exchange Sensitivity

While we believe we understand the A&A ecosystem well enough to identify the ad exchanges, we wish to understand the impact of this classification on our privacy metrics. Our classification depends upon the ad hoc selection of thresholds of outdegree and in/outdegree ratio, so we desire more context of the impact this makes on privacy.

To show the effect this selection has on privacy, we sweep the number of ad exchanges we include from 0 up to the 47 selected in Section 5, always adding the origin that connects to the most distinct A&A domains. We graph the Nth percentile ad exchange's visibility into site impressions as we vary the number of exchanges. The results of this are shown in Figure 3.

The most impactful ad exchanges on the privacy of end-users are generally the largest, which we add first in our sweep. We observe this by the majority of growth in each percentile we show occurring over the first few exchanges we add. Beyond this point, each exchange has very little impact on the number of impressions observed by each A&A domain. This is consistent with prior work and intuition: those agents that share information most broadly are those most impactful to privacy.

6.2 Tool Evaluation

With inclusion graphs for each site impression and the blocking behavior of six different tools captured, we can compare the privacy granted through the use of each tool. We add a seventh measurement to our six tools we analyze in previous sections: the control, in which we perform no blocking. Further details of how we perform this analysis are in Section 3.5 and both models of information sharing are described in Section 3.5. The distributions of information sharing under these models are shown in Figure 4.

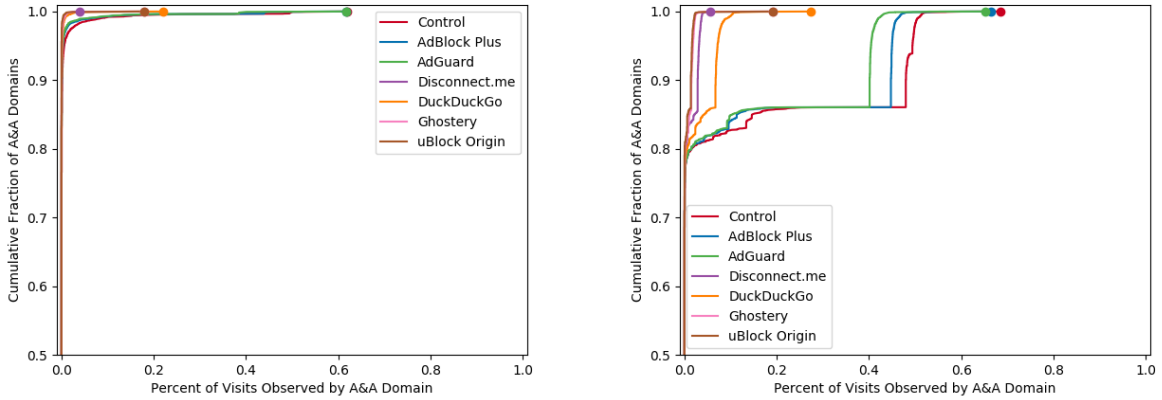

(a) Cookie Matching

(b) Real-Time Bidding

Figure 4: Privacy Evaluation of Blocking Extensions. We show the rate of visit disclosure over our known advertisers using two different disclosure models. At the left we see the Cookie Matching model and at right we see the Real-Time Bidding model.

In the Cookie Matching model there is much less impression disclosure. Any blocking performs better than nothing, however the advertising focused services (AdGuard and AdBlock Plus) provided very little benefit in blocking the worst case A&A domains. The remaining blocking tools performed similarly better with Disconnect.me having the clearly best worst-case performance. In the Real-Time Bidding model we see similar trends, but with overall more disclosure as exchanges provide more sharing.

In the Real-Time Bidding model, we observe substantial improvements in user privacy by some extensions, and similarly we see the least improvement in tools focused on advertisement blocking. However, there is no uniformly best extension for user privacy: Disconnect.me again has the best worst case performance, but lags Ghostery and uBlock Origin in the 85th to 99th percentile of A&A domains. Depending on threat model, concern with the most privileged A&A domain, specific A&A domains, or average performance may be used to select the defense among extensions focused on privacy. This figure the used to generate it are needed to answer those questions quantitatively.

Tools have similar classes of performance grouped based upon their mission and similar order of worst case performance in both cases. It is unclear whether there is a universal best strategy independent of threat model and user behavior. However, tools with an emphasis on privacy show large improvements to user privacy.

Overall, several aspects of these distributions are interesting in comparison to prior work. The rates of observation by publishers are lower, providing a more optimistic view of third-party tracking than other studies. There are several possible explanations for this: differences in how these rates of observation are computed, differences in the population of sites being measured, and differences in how sites are sampled from this population. For example, we do not focus our study on sites that have a higher incidence of tracking techniques, e.g. shopping sites. Also, as shown in Figure 5,

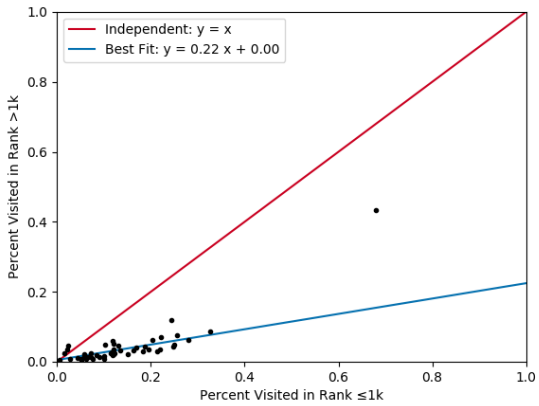

Figure 5: Exchange Connection Frequency by Rank. We demonstrate on a scatter plot with best fit line that ad exchanges generally are five times as likely to be on a website in the top thousand than in the rest of the top million domains. This is with the removal of the outlier `doubleclick.net`, at the upper right of the plot.

ad exchange prevalence is reduced in sites beyond the top thousand, where prior work has focused.

However, we see similar qualitative differences between the Cookie Matching model and Real-Time Bidding model, described in Section 3.5, as the prior work that introduced them. Under Cookie Matching, very few A&A parties have significant access to the users' browsing patterns. Under Real-Time Bidding, there are some number of A&A domains that observe almost no traffic, and a wide variety of scopes of observation for the most well positioned A&A domains.

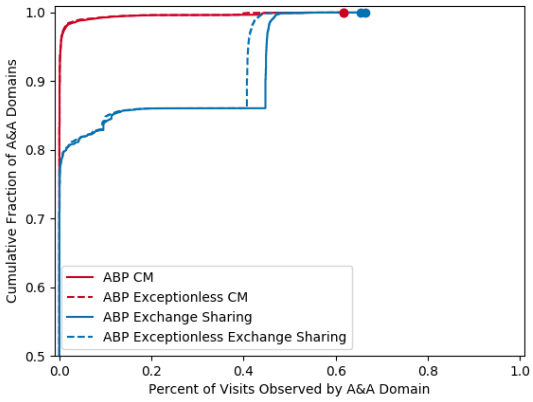

Figure 6: Evaluating Acceptable Advertisements. We show the impact of the Acceptable Advertising list on the privacy of AdBlock Plus in our crawl over both models of information sharing.

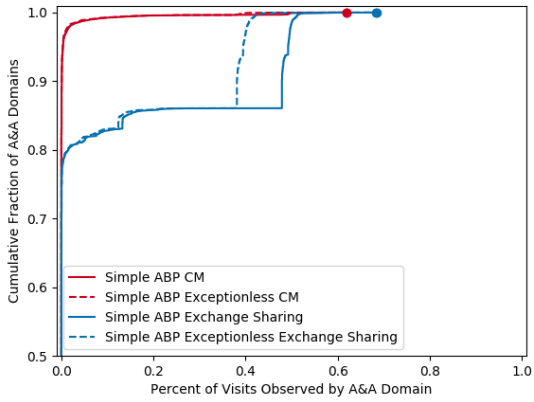

Figure 7: Evaluating Acceptable Advertisements Under Simplified Model. We show the impact of the Acceptable Advertising list on the privacy of AdBlock Plus in our crawl over both models of information sharing when the rule flags are not considered, as in prior work. The impact of the Acceptable Ads program is exaggerated.

6.3 Acceptable Advertisement

In Section 4.2 we observed that only AdBlock Plus has any change in blocking behavior when the rules explicitly allowing domains are removed. It is also the worst performing blocking extension. We know that some privacy is lost due to the Acceptable Advertising list, and prior work has identified it as a major source of reduced efficacy. However, we find this not to be the case under our more accurate representation of the filter lists. Figure 6 shows the effect of removing the Acceptable Advertising list from AdBlock Plus.

We see statistically significant difference between the two in the Two-sample Kolmogorov–Smirnov test ($p < 0.05$). However, the effect size of this difference does not in either case account for the entire difference between AdBlock Plus and the most similar tool in the composition of its rule list, uBlock Origin. Instead it only accounts for the difference to AdGuard, another advertising focused blocking extension. This indicates that most of AdBlock Plus' failures to protect its users are in its focus on advertising instead of its Acceptable Ads program.

6.4 Improved Model

In Section 3.2, we noted that our model of blocking rules is more detailed than that of prior work, and that allows us to better understand the role of acceptable advertising programs. To justify this, we repeat our experiments from Section 6.3 without consideration for rule flags, to show how results change under a simplified model like that of prior work. This is shown in Figure 7.

The effects of the acceptable advertising program is exaggerated without the rule flags. With flags removed, the tool performs even better without exceptions and even worse with exceptions. Particularly notable is that when we do not consider rule flags, we reproduce, from Bashir et al., that AdBlock Plus performs very similarly to the control. This is due to the inclusion of many ad exchanges in the exception list that have ignored domain constraints. We note that in our experiments, the difference between AdBlock Plus without rule flags and no blocking exists, but is within the confidence interval our sample size allows.

7 DISCUSSION

In this paper, we have quantified several aspects of popular blocking extensions. We found considerable differences between the extensions and a few surprising similarities. Our results lead us to a few key lessons, caveated by the limitations of this work.

7.1 Findings and Lessons

Several key lessons can be drawn from our results. First, we found that extensions focused on ad blocking (AdGuard and AdBlock Plus) provide the weakest practical protection for users. This is despite using privacy as a marketing tactic on their websites. Second, we found that despite building off of EasyList and EasyPrivacy to build their block lists, uBlock Origin, AdGuard, and AdBlock Plus all behave very differently in what they block and how much protection they provide the end user. Third, we found evidence that while AdBlock Plus has the only significant Acceptable Advertisement program, it only accounts for a 9% increase in total impression disclosures in the Real-Time Bidding model. Finally, we find that completeness of manual effort to enumerate advertisers to block does not correspond to better user privacy; Disconnect.me has the worst measured completeness and has strong privacy properties. From these findings, we can draw the following lessons.

If users seek a blocking extension that benefits their privacy, they should focus on those whose primary emphasis is privacy or user agency. In particular, Disconnect.me provides the best protection against the very strongest web trackers and uBlock Origin and Ghostery provide slightly better protection against the next thousand best web trackers. If users are comfortable configuring

their extension with custom block lists, they may not even have to make this tradeoff.

We also learn that distributed manual effort is not needed to curate an effective block list. The roughly 500 domains identified and blocked by Disconnect.me provided some of the best coverage we saw. A focus on continued identification and third-party blocking of major players in the advertising ecosystem is likely to provide the best trade-off of effort for privacy. Additionally, this may help future performance optimization of tracker blocking tools. This is likely because of the significant presence and interconnectedness a few players have in this ecosystem.

Finally, we learn that despite their marketing, and even with Acceptable Ads programs disabled, ad blockers are not nearly as effective privacy enhancing technologies as their peers that focus on user privacy and agency.

7.2 Limitations

Our work presents an improved understanding of the Advertising and Analytics ecosystems over state of the art. However it is important to understand where the potential shortcomings of our model fall so that we can both understand our results in the appropriate context and understand avenues of future work.

For example, our implementation of blocking rule application implements more features than other work that were critical to understanding the role of Acceptable Advertising programs [11]. However, we do not implement all documented features, let alone capture any possible undocumented features of each tool. While this may reduce our accuracy, we believe our improvement and the insight it provides is significant.

Additionally, our results driven by our crawl data are restricted to the sites we evaluate: home pages of the Tranco top million sites. However, we add to the literature an updated million site measurement measurement not drawn from proprietary data.

We do not evaluate the quality of the candidate advertisers we discover through our heuristics. It is a open problem to automatically differentiate between "breaking" a site such that it is no longer functional and effectively blocking A&A techniques. Current promising approaches are browser instrumentation of what sites cause users to disable their blocking extension. We consider this evaluation out of scope, and believe that our small candidate lists are currently useful for several purposes: manual evaluation for current browser extensions and a test-list with low false positive rate to reduce base-rate concerns when automatically adding them to browsers to be evaluated.

Finally, we acknowledge that our results on user privacy are based on a model of information sharing that is an approximation of one aspect of the A&A ecosystem. The model does not capture all active members of the ecosystem or information channels that do not involve interaction with the end-user machine. For example, first-parties collecting information for sale is missed in our models.

8 CONCLUSION

The insight we gain into blocking extensions from this work will help further improvement of these blocking extensions and privacy advice surrounding them. We are able to identify similarities between tool behavior that indicate list sharing. We use these lists to identify ad exchanges and candidate advertisers using simple heuristics of the inclusion graph. We discover that only AdBlock Plus has a significant explicit exception program. We find that these features are not just true of the rules exercised by our crawl, but also of the rules defined overall.

We identify privacy impacts of browser extensions by analyzing how each extension affects which parties receive requests from page visits in our crawl. This allows evidence driven recommendations to use Disconnect.me, Ghostery, or uBlock Origin to users and support for Firefox's selection of the Disconnect.me list for its content blocking deployment. The differences we observe also allow inference that organizational mission plays the largest part in determining an extension's efficacy against third party tracking under default configurations. Also, we confirm that list size is not a meaningful proxy for provided privacy.

These results comprise a significant improvement in our understanding of blocking extensions and our ability to aid in the defense of user privacy from third-party web tracking.

REFERENCES

[1] AdGuard. [n.d.]. How to create your own ad filters. Accessed 31 May 2019: https://kb.adguard.com/en/general/how-to-create-your-own-ad-filters.

[2] AdGuard. [n.d.]. Search ads and self-promotion. Accessed 31 May 2019: https://kb.adguard.com/en/general/search-ads-and-self-promotion.

[3] AdGuard. [n.d.]. The world's most advanced ad blocker! Accessed 31 May 2019: https://adguard.com/en/welcome.html.

[4] Alexa Internet, Inc. [n.d.]. Alexa Top 1,000,000 Sites. http://s3.amazonaws.com/alexa-static/top-1m.csv.zip.

[5] Alphabet. 2019. Alphabet Announces First Quarter 2019 Results. Accessed 31 May 2019: https://abc.xyz/investor/static/pdf/2019Q1_alphabet_earnings_release.pdf.

[6] Ricardo A Baeza-Yates and Gaston H Gonnet. [n.d.]. Fast text searching for regular expressions or automaton searching on tries. *Journal of the ACM (JACM)* 43, 6 ([n. d.]).

[7] Rebecca Balebako, Pedro Leon, Richard Shay, Blase Ur, Yang Wang, and L Cranor. [n.d.]. Measuring the effectiveness of privacy tools for limiting behavioral advertising. In *Proceedings of the 6th IEEE Workshop on Web 2.0 Security and Privacy (W2SP 2012)*.

[8] Paul Barford, Igor Canadi, Darja Krushevskaja, Qiang Ma, and S Muthukrishnan. [n.d.]. Adscape: Harvesting and analyzing online display ads. In *Proceedings of the 23rd International Conference on World Wide Web (WWW 2014)*.

[9] Muhammad Ahmad Bashir, Sajjad Arshad, William Robertson, and Christo Wilson. [n.d.]. Tracing information flows between ad exchanges using retargeted ads. In *Proceedings of the 25th USENIX Security Symposium (Security 2016)*.

[10] Muhammad Ahmad Bashir, Umar Farooq, Maryam Shahid, Muhammad Fareed Zaffar, and Christo Wilson. [n.d.]. Quantity vs. Quality: Evaluating User Interest Profiles Using Ad Preference Managers. In *Proceedings of the 26th Internet Society Network and Distributed System Security Symposium (NDSS 2019)*.

[11] Muhammad Ahmad Bashir and Christo Wilson. [n.d.]. Diffusion of user tracking data in the online advertising ecosystem. *Proceedings on Privacy Enhancing Technologies (PoPET)* 2018, 4 ([n. d.]).

[12] Aaron Cahn, Scott Alfeld, Paul Barford, and S Muthukrishnan. [n.d.]. An empirical study of web cookies. In *Proceedings of the 25th International Conference on World Wide Web (WWW 2016)*.

[13] Juan Miguel Carrascosa, Jakub Mikians, Ruben Cuevas, Vijay Erramilli, and Nikolaos Laoutaris. [n.d.]. I always feel like somebody's watching me: measuring online behavioural advertising. In *Proceedings of the 11th ACM Conference on Emerging Networking Experiments and Technologies (CoNEXT 2015)*.

[14] Claude Castelluccia, Mohamed-Ali Kaafar, and Minh-Dung Tran. [n.d.]. Betrayed by your ads!. In *Proceedings of the 12th Privacy Enhancing Technologies Symposium (PETS 2012)*.

[15] Claude Castelluccia, Lukasz Olejnik, and Tran Minh-Dung. [n.d.]. Selling off privacy at auction. In *Proceedings of the 21st Internet Society Network and Distributed System Security Symposium (NDSS 2014)*.

[16] Farah Chanchary and Sonia Chiasson. [n.d.]. User perceptions of sharing, advertising, and tracking. In *Proceedings of the 11th ACM Symposium on Usable Privacy and Security (SOUPS 2015)*.

[17] Cliqz. [n.d.]. Faster, safer, smarter browsing. Accessed 31 May 2019: https://www.ghostery.com/.

[18] Martin Degeling, Christine Utz, Christopher Lentzsch, Henry Hosseini, Florian Schaub, and Thorsten Holz. 2018. We Value Your Privacy... Now Take Some

Cookies: Measuring the GDPR's Impact on Web Privacy. Technical Report, arXiv preprint arXiv:1808.05096.

[19] Disconnect. [n.d.]. Tracking protection lists. Accessed 31 May 2019: https://disconnect.me/trackerprotection.

[20] Disconnect. [n.d.]. We Protect Your Privacy. Accessed 31 May 2019: https://disconnect.me/.

[21] Claire Dolin, Ben Weinshel, Shawn Shan, Chang Min Hahn, Euirim Choi, Michelle L Mazurek, and Blase Ur. [n.d.]. Unpacking Perceptions of Data-Driven Inferences Underlying Online Targeting and Personalization. In *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems (CHI 2018)*.

[22] DuckDuckGo. [n.d.]. DuckDuckGo Privacy Essentials. Accessed 31 May 2019: https://addons.mozilla.org/en-US/firefox/addon/duckduckgo-for-firefox/.

[23] Steven Englehardt and Arvind Narayanan. [n.d.]. Online tracking: A 1-million-site measurement and analysis. In *Proceedings of the 23rd ACM Conference on Computer and Communications Security (CCS 2016)*.

[24] Steven Englehardt, Dillon Reisman, Christian Eubank, Peter Zimmerman, Jonathan Mayer, Arvind Narayanan, and Edward W Felten. [n.d.]. Cookies that give you away: The surveillance implications of web tracking. In *Proceedings of the 24th International Conference on World Wide Web (WWW 2015)*.

[25] Eyeo GmbH. [n.d.]. Allowing acceptable ads in Adblock Plus. Accessed 31 May 2019: https://adblockplus.org/acceptable-ads.

[26] Eyeo GmbH. [n.d.]. Homepage. Accessed 31 May 2019: https://adblockplus.org/.

[27] Eyeo GmbH. [n.d.]. Writing AdBlock Plus Filters. Accessed 31 May 2019: https://adblockplus.org/filters.

[28] Marjan Falahrastegar, Hamed Haddadi, Steve Uhlig, and Richard Mortier. [n.d.]. Tracking personal identifiers across the web. In *Proceedings of the 17th International Conference on Passive and Active Network Measurement (PAM 2016)*.

[29] Arthur Gervais, Alexandros Filios, Vincent Lenders, and Srdjan Capkun. [n.d.]. Quantifying web adblocker privacy. In *Proceedings of the 22nd European Symposium on Research in Computer Security (ESORICS 2017)*.

[30] Arpita Ghosh, Mohammad Mahdian, R Preston McAfee, and Sergei Vassilvitskii. 2015. To match or not to match: Economics of cookie matching in online advertising. *ACM Transactions on Economics and Computation (TEAC)* 3, 2 (2015).

[31] Phillipa Gill, Vijay Erramilli, Augustin Chaintreau, Bala Krishnamurthy, Dina Papagiannaki, and Pablo Rodriguez. [n.d.]. Follow the money: Understanding economics of online aggregation and advertising. In *Proceedings of the 13th ACM Internet Measurement Conference (IMC 2013)*.

[32] Richard Gomer, Eduarda Mendes Rodrigues, Natasa Milic-Frayling, and MC Schraefel. [n.d.]. Network analysis of third party tracking: User exposure to tracking cookies through search. In *Proceedings of the IEEE/WIC/ACM International Joint Conferences on Web Intelligence (WI) and Intelligent Agent Technologies (WI 2013)*.

[33] gorhill. [n.d.]. Static filter syntax. Accessed 31 May 2019: https://github.com/gorhill/uBlock/wiki/Static-filter-syntax.

[34] gorhill. [n.d.]. uBlock. Accessed 31 May 2019: https://github.com/gorhill/uBlock.

[35] Miguel Helft. 2010. Marketers Can Glean Private Data on Facebook. New York Times. Accessed 31 May 2019: https://www.nytimes.com/2010/10/23/technology/23facebook.html.

[36] Kashmir Hill. 2012. How Target Figured Out A Teen Girl Was Pregnant Before Her Father Did. Forbes. Accessed 31 May 2019: https://www.forbes.com/sites/kashmirhill/2012/02/16/how-target-figured-out-a-teen-girl-was-pregnant-before-her-father-did.

[37] Umar Iqbal, Zubair Shafiq, and Zhiyun Qian. [n.d.]. The ad wars: retrospective measurement and analysis of anti-adblock filter lists. In *Proceedings of the 17th ACM Internet Measurement Conference (IMC 2017)*.

[38] Vasiliki Kalavri, Jeremy Blackburn, Matteo Varvello, and Konstantina Papagiannaki. [n.d.]. Like a pack of wolves: Community structure of web trackers. In *Proceedings of the 17th International Conference on Passive and Active Network Measurement (PAM 2016)*.

[39] Arjaldo Karaj, Sam Macbeth, Rémi Berson, and Josep M Pujol. 2018. WhoTracks.Me: Monitoring the online tracking landscape at scale. Technical Report, arXiv preprint arXiv:1804.08959.

[40] J. Kline, A. Cahn, and P. Barford. [n.d.]. Placement Laundering and the Complexities of Attribution in Online Advertising. *Proceedings of the 17th International Conference on WWW/Internet (ICWI 2018)* ([n. d.]).

[41] Victor Le Pochat, Tom Van Goethem, Samaneh Tajalizadehkhoob, Maciej Korczyński, and Wouter Joosen. 2019. Tranco: A Research-Oriented Top Sites Ranking Hardened Against Manipulation. In *Proceedings of the 26th Annual Network and Distributed System Security Symposium (NDSS 2019)*.

[42] Mathias Lécuyer, Guillaume Ducoffe, Francis Lan, Andrei Papancea, Theofilos Petsios, Riley Spahn, Augustin Chaintreau, and Roxana Geambasu. [n.d.]. XRay: Enhancing the web's transparency with differential correlation. In *Proceedings of the 23rd USENIX Security Symposium (Security 2014)*.

[43] Matthew Malloy, Mark McNamara, Aaron Cahn, and Paul Barford. [n.d.]. Ad blockers: Global prevalence and impact. In *Proceedings of the 16th ACM Internet Measurement Conference (IMC 2016)*.

[44] Arunesh Mathur, Jessica Vitak, Arvind Narayanan, and Marshini Chetty. [n.d.]. Characterizing the use of browser-based blocking extensions to prevent online tracking. In *Proceedings of the 14th USENIX Symposium on Usable Privacy and Security (SOUPS 2018)*.

[45] William Melicher, Mahmood Sharif, Joshua Tan, Lujo Bauer, Mihai Christodorescu, and Pedro Giovanni Leon. [n.d.]. (Do Not) Track me sometimes: users' contextual preferences for web tracking. *Proceedings on Privacy Enhancing Technologies (PoPET)* 2016, 2 ([n. d.]).

[46] Georg Merzdovnik, Markus Huber, Damjan Buhov, Nick Nikiforakis, Sebastian Neuner, Martin Schmiedecker, and Edgar Weippl. [n.d.]. Block me if you can: A large-scale study of tracker-blocking tools. In *Proceedings of the 2nd IEEE European Symposium on Security and Privacy (2017 EuroS&P)*.

[47] Mozilla. [n.d.]. Content blocking. Accessed 31 May 2019: https://support.mozilla.org/en-US/kb/content-blocking.

[48] Rishab Nithyanand, Sheharbano Khattak, Mobin Javed, Narseo Vallina-Rodriguez, Marjan Falahrastegar, Julia E Powles, Emiliano De Cristofaro, Hamed Haddadi, and Steven J Murdoch. [n.d.]. Adblocking and counter blocking: A slice of the arms race. In *Proceedings of the 6th USENIX Workshop on Free and Open Communications on the Internet (FOCI 2016)*.

[49] Panagiotis Papadopoulos, Nicolas Kourtellis, and Evangelos P Markatos. [n.d.]. Cookie synchronization: Everything you always wanted to know but were afraid to ask. In *Proceedings of the 28th International Conference on World Wide Web (WWW 2019)*.

[50] Panagiotis Papadopoulos, Nicolas Kourtellis, Pablo Rodriguez Rodriguez, and Nikolaos Laoutaris. [n.d.]. If you are not paying for it, you are the product: How much do advertisers pay to reach you?. In *Proceedings of the 17th ACM Internet Measurement Conference (IMC 2017)*.

[51] Enric Pujol, Oliver Hohlfeld, and Anja Feldmann. [n.d.]. Annoyed users: Ads and ad-block usage in the wild. In *Proceedings of the 15th ACM Internet Measurement Conference (IMC 2015)*.

[52] Abbas Razaghpanah, Rishab Nithyanand, Narseo Vallina-Rodriguez, Srikanth Sundaresan, Mark Allman, Christian Kreibich, and Phillipa Gill. [n.d.]. Apps, trackers, privacy, and regulators: A global study of the mobile tracking ecosystem. In *Proceedings of the 25th Internet Society Network and Distributed System Security Symposium (NDSS 2018)*.

[53] Franziska Roesner, Tadayoshi Kohno, and David Wetherall. [n.d.]. Detecting and defending against third-party tracking on the web. In *Proceedings of the 9th USENIX Conference on Networked Systems Design and Implementation (NSDI 2012)*.

[54] Quirin Scheitle, Oliver Hohlfeld, Julien Gamba, Jonas Jelten, Torsten Zimmermann, Stephen D Strowes, and Narseo Vallina-Rodriguez. [n.d.]. A long way to the top: significance, structure, and stability of internet top lists. In *Proceedings of the 18th ACM Internet Measurement Conference (IMC 2018)*.

[55] The EasyList authors. [n.d.]. EasyList. Accessed 31 May 2019: https://easylist.to/.

[56] Stuart A. Thompson. 2019. These Ads Think They Know You. New York Times. Accessed 31 May 2019: https://www.nytimes.com/interactive/2019/04/30/opinion/privacy-targeted-advertising.html.

[57] Martino Trevisan, Stefano Traverso, Eleonora Bassi, and Marco Mellia. [n.d.]. 4 Years of EU Cookie Law: Results and Lessons Learned. *Proceedings on Privacy Enhancing Technologies (PoPET)* 2019, 2 ([n. d.]).

[58] Blase Ur, Pedro Giovanni Leon, Lorrie Faith Cranor, Richard Shay, and Yang Wang. [n.d.]. Smart, useful, scary, creepy: perceptions of online behavioral advertising. In *Proceedings of the 8th ACM Symposium on Usable Privacy and Security (SOUPS 2012)*.

[59] Tobias Urban, Dennis Tatang, Martin Degeling, Thorsten Holz, and Norbert Pohlmann. 2018. The Unwanted Sharing Economy: An Analysis of Cookie Syncing and User Transparency under GDPR. Technical Report, arXiv preprint arXiv:1811.08660.

[60] Antoine Vastel, Peter Snyder, and Benjamin Livshits. 2018. Who Filters the Filters: Understanding the Growth, Usefulness and Efficiency of Crowdsourced Ad Blocking. Technical Report, arXiv preprint arXiv:1810.09160.

[61] Giridhari Venkatadri, Elena Lucherini, Piotr Sapiezynski, and Alan Mislove. [n.d.]. Investigating sources of PII used in Facebook's targeted advertising. *Proceedings on Privacy Enhancing Technologies (PoPET)* 2019, 1 ([n. d.]).

[62] Robert J Walls, Eric D Kilmer, Nathaniel Lageman, and Patrick D McDaniel. [n.d.]. Measuring the impact and perception of acceptable advertisements. In *Proceedings of the 15th ACM Internet Measurement Conference (IMC 2015)*.

[63] Yaxing Yao, Davide Lo Re, and Yang Wang. [n.d.]. Folk models of online behavioral advertising. In *Proceedings of the 20th ACM Conference on Computer Supported Cooperative Work and Social Computing (CSCW 2017)*.

[64] Zhonghao Yu, Sam Macbeth, Konark Modi, and Josep M Pujol. [n.d.]. Tracking the trackers. In *Proceedings of the 25th International Conference on World Wide Web (WWW 2016)*.

[65] Sebastian Zimmeck, Jie S Li, Hyungtae Kim, Steven M Bellovin, and Tony Jebara. [n.d.]. A privacy analysis of cross-device tracking. In *Proceedings of the 26th USENIX Security Symposium (Security 2017)*.

A AD EXCHANGES

Table 6: Ad Exchanges. Ad exchanges, as determined by Section 5 are shown here, with the parameters used to classify them. These parameters are the out degree and ratio of in degree to out degree in the Inclusion graph including only nodes that are advertising and analytics domains.

Domain	Out-degree	In/Out-degree Ratio
doubleclick.net	585	1.22
adnxs.com	449	0.66
googlesyndication.com	274	1.48
rubiconproject.com	254	0.57
openx.net	200	0.55
demdex.net	193	0.70
bidswitch.net	191	0.61
pubmatic.com	191	0.75
mathtag.com	189	0.58
adform.net	170	1.01
amazon-adsystem.com	163	0.67
criteo.com	153	0.64
turn.com	142	0.52
yahoo.com	140	0.78
taboola.com	137	0.75
crwdcntrl.net	133	0.70
bluekai.com	129	0.76
lijit.com	129	0.93
krxd.net	127	0.76
contextweb.com	127	0.68
smartadserver.com	122	0.65
w55c.net	105	0.52
everesttech.net	103	0.64
1rx.io	100	0.66
tapad.com	94	0.74
3lift.com	86	0.74
gumgum.com	86	0.78
spotxchange.com	80	0.51
serving-sys.com	79	1.95
wp.com	75	1.52
2mdn.net	73	1.51
teads.tv	71	0.86
imasdk.googleapis.com	68	1.40
outbrain.com	68	0.56
sharethrough.com	66	0.68
yandex.ru	66	0.76
adition.com	62	1.05
sonobi.com	62	0.77
addthis.com	62	1.15
rfihub.com	59	0.69
sharethis.com	59	0.71
eyeota.net	58	0.55
stickyadstv.com	58	0.66
yimg.com	57	0.96
33across.com	57	1.19
media.net	56	0.89
bttrack.com	51	0.65

B QUANTIFIED PRIVACY IMPACT

Table 7: Quantified Privacy Impact. We provide here the numerical values achieved for our privacy evaluation experiments in Section 6. The percent of visits disclosed averaged over all know A&A domains and to the most priveleged A&A domain are provided for both information sharing models we considered. This is a tool to clarify details that may be obscured by overlapping lines or hard to visually compare, such as area under the curve. We note that overlapping lines indicate that we did not observe a significant difference.

	Cookie Matching		**Real-Time Bidding**	
	Average	Worst	Average	Worst
Control	0.238%	36.90%	4.475%	40.88%
AdBlock Plus	0.183%	36.78%	4.063%	39.60%
AdGuard	0.162%	36.75%	3.667%	38.88%
Disconnect.me	0.015%	2.30%	0.309%	3.38%
DuckDuckGo	0.028%	13.10%	0.718%	16.38%
Ghostery	0.016%	10.70%	0.181%	11.43%
uBlock Origin	0.011%	10.65%	0.158%	11.45%
AdBlock Plus (Exceptionless)	0.166%	36.78%	3.720%	38.98%