Addressing Anonymous Abuses: Measuring the Effects of Technical Mechanisms on Reported User Behaviors

Wajeeha Ahmad
MIT CSAIL
Cambridge, MA, USA
wajeeha@csail.mit.edu

Ilaria Liccardi
MIT CSAIL
Cambridge, MA, USA
ilaria@csail.mit.edu

ABSTRACT
Anonymous networks intended to enhance privacy and evade censorship are also being exploited for abusive activities. Technical schemes have been proposed to selectively revoke the anonymity of abusive users, or simply limit them from anonymously accessing online service providers. We designed an empirical survey study to assess the effects of deploying these schemes on 75 users of the Tor anonymous network. We evaluated proposed schemes based on examples of the intended or abusive use cases they may address, their technical implementation and the types of entities responsible for enforcing them. Our results show that revocable anonymity schemes would particularly deter the intended uses of anonymous networks. We found a lower reported decrease in usage for schemes addressing spam than those directly compromising free expression. However, participants were concerned that all technical mechanisms for addressing anonymous abuses could be exploited beyond their intended goals (51.7%) to harm users (43.8%). Participants were distrustful of the enforcing entities involved (43.8%) and concerned about being unable to verify (49.3%) how particular mechanisms were applied.

Author Keywords
Anonymous networks; Trust; Abuse; Empirical study; Tor.

CCS Concepts
•**Security and privacy** → **Social aspects of security and privacy;** •**Human-centered computing** → *User studies;* •**Social and professional topics** → Censorship; Surveillance;

INTRODUCTION
In an era of mass surveillance by governments and corporations alike, online anonymity is often considered indispensable to free expression and individual privacy. People seek anonymity online for various important reasons such as to gain protection from governments and repressive regimes [67, 74], evade commercial surveillance, better manage boundaries in personal and professional relationships, and avoid harassment from online, offline and unspecified entities [38, 44]. Other uses include anonymously accessing information or censored materials [39, 76], gathering intelligence or tips [59], and discussing stigmatized topics [30].

Yet anonymity makes it difficult to trace or exclude abusive users. Some exploit the veil of anonymity to engage in illegal drug exchanges [9], harassment [46] and terrorist plots [13, 75]. Moreover, the Tor anonymous network suffers from botnet attacks [36, 53, 66] among other abuses. There also exist botnet constructions that researchers claim could be nearly impossible to subvert without blocking all access to anonymous networks [61]. Because some use Tor to attack services, spam forums and scan for vulnerabilities, many service providers and content delivery networks treat all users connecting from known anonymous networks as "second-class" web citizens [45], forcing them to solve multiple CAPTCHAs or blocking them.

Can we simultaneously promote the legitimate uses of anonymous networks while mitigating their abuses? In 2007, Tor's original developers remarked: *"Simple technical mechanisms can remove the ability to abuse anonymously without undermining the ability to communicate anonymously"* [35]. But do users perceive technical mechanisms as effectively curtailing anonymous abuses without reducing their own legitimate uses? What additional factors need to be considered in making such decisions? Using both quantitative and qualitative approaches, we study the desirability of different mechanisms to deter abuse among users of anonymous networks. We show how and why three main factors associated with proposals for countering abuses affect the intended uses of anonymous networks. We illustrate how users' awareness of different activities conducted via anonymous networks could reflect their responses to various technical mechanisms. Finally, we describe how users' responses inform policies for the design and implementation of measures for addressing anonymous abuses.

RELATED RESEARCH
Anonymous networks were designed to prevent online tracking in order to protect free expression and enhance privacy [20, 34, 60] as well as resist censorship [33]. Many studies detail anonymity as allowing for more disclosure [64] across all intimacy levels [54], encouraging both beneficial and harmful behaviors in collaborative learning [23] and other social [25] settings. Several others explore peoples' motivations for seeking anonymity ranging from gaining protection against various actors [38, 39, 44] to general usage and exploration [39, 76]. People attain online anonymity by using different tools [72], incorporating behavioral changes such as creating several accounts [21] or altering personal profiles [71]. Anonymous networks such as Tor, I2P and Freenet aim to hide users' network identity (i.e. IP address) from

CHI '20, April 25–30, 2020, Honolulu, HI, USA.

ACM ISBN 978-1-4503-6708-0/20/04 ...$15.00.
http://dx.doi.org/10.1145/3313831.3376690

unwanted observations. Of these, Tor is considered the largest network with millions of daily users.[1] To prevent tracking of users' communication, Tor reroutes traffic through three randomly chosen and globally distributed volunteer-run servers called "nodes" or "relays" [34]. Tor also offers onion services, which are websites that protect both their own and users' anonymity.

Debate about advancing or banning online anonymity has been ongoing among security researchers [28], policy experts [2, 22, 43, 49], and community designers [48] among others. While users' opinions range from viewing Tor as a force for freedom to a tool for cybercriminals and terrorists, many believe that the balance between individual privacy and national security should be closer to privacy [39]. Some users have also complained about insufficient protection specifically from authorities or big companies with a few raising concerns about the criminal content of onion services [76]. From the perspective of some open collaboration service providers, anonymous users make valuable contributions and do not violate community norms more frequently than other users [56]. According to one study, Tor users contribute similar proportions of damaging and good faith edits on Wikipedia as non-Tor users with no substantial differences in quality [68].

However, online anonymity is also associated with toxic behaviors that are hard to control [51]. Given threats from users of anonymous networks [24, 78], websites such as Wikipedia and Slash-dot have had to ban their contributions [41]. While onion services have been found to offer both illegal and other content (about human rights, free speech, security, etc.) [14], those serving criminal and unethical uses including botnets and adult content are among the most popular services [14, 40, 57, 79]. Some claim that the inability to deter the abuse of anonymous networks hinders their widespread adoption [31, 73] and leads to service providers blocking all anonymous users [32, 41, 52, 70]. To address these concerns, researchers have proposed several cryptographic schemes, which fall into two main groups based on their goals: revocable anonymity and access-limiting schemes.

Revocable anonymity schemes aim to provide anonymity for ordinary users, while simultaneously guaranteeing traceability of abusive users. Such schemes are meant to deter abuse by allowing potential investigators to find the identity of suspected users. Some of these schemes use trusted third parties (TTPs) to register all users and revoke the anonymity of certain users [26, 31, 32, 47, 73, 77]: registration entities aim to offer unique credentials such as new pseudonyms to enable users to access anonymous networks whereas revocation entities may cooperate with registration entities to revoke a user's anonymity in case of a legal investigation. These TTPs may be centralized or implemented distributedly [19, 31, 73, 77] via secret sharing that allows a set of parties to reconstruct a secret key only when a sufficient number of them all consent and collaborate to do so [62]. Revocable anonymity schemes without TTPs [7, 17] allow investigators to trace back the source of an anonymous communication stream by requiring all nodes of the anonymous network to reveal their predecessors.

Contrasting this approach, researchers have proposed using existing software vulnerabilities for lawful access to communications in case of legal investigations since there will always be urgent situations involving national security or major crimes, where built-in intercept mechanisms are not available [10, 11]. This approach has been used by law enforcement agencies and is expected to increase in utility as anonymous tools become more widespread [55].

Secondly, access-limiting schemes aim to enable service providers (e.g. websites) to selectively limit the access of certain users without revealing their identities. Some access-limiting schemes incorporate TTPs: in Nymble, service providers require a TTP (the "nymble manager") to provide a token linking the user's identity to their actions in order to temporarily block the user [70]. TorPolice aims to allow service providers to rate-limit only those anonymous users engaging in botnet-enabled abuses (e.g. spamming forums, scraping content, etc.) using CAPTCHAs or computational puzzles [52]. In access-limiting schemes without trusted third parties, users present zero-knowledge proofs to a service provider to demonstrate that they are not part of the service provider's blacklist before accessing its services [4, 5, 6, 16, 69].

Our work extends prior research in three ways. First, we test the desirability and impact of proposed technical mechanisms on actual users of anonymous networks. Second, we investigate why users respond differently to various technical mechanisms depending on the case, scheme and decision-making entity involved in the mechanism. Finally, we gather data on users' understanding of abuses of anonymous networks to glean insights on the debate of how such issues may be addressed without negatively impacting the intended uses of anonymous networks. This is the first study that explores the tensions between protecting anonymity and addressing its potential abuses from the perspective of anonymous network users.

USER STUDY

Our study was designed to test the effects of proposed anti-abuse technical mechanisms on current users of anonymous networks. Specifically, we sought to understand whether users would alter their usage of anonymous networks depending on the type of technical scheme (e.g. revocable anonymity or access limiting scheme) implemented. To capture the diverse social contexts in which these schemes may be deployed, we also tested the effects of five popular use cases of anonymous networks that may be addressed in different circumstances (spam, phishing, illegal drug exchange, communication and reporting[2]), and five types of decision-making entities responsible for addressing potential abuses (anonymous network administrators, non-government organizations, anonymous nodes, government agencies and commercial services). We also wished to understand whether users' decisions might be influenced by their own prior knowledge of encountered or known abuses associated with anonymous networks. We aimed to identify the circumstances, if any, under which users may view technical schemes as useful and not impacting or deterring their own usage. In particular, we are interested in investigating:

- Which factors i.e. case, scheme and/or entity affect users' self-reported usage of anonymous networks?
- Does knowledge of abuses or security vulnerabilities associated with anonymous networks affect users' responses to technical mechanisms for addressing anonymous abuses?

[1] https://metrics.torproject.org/userstats-relay-country.html

[2] In scenarios presented to participants, communication and reporting were depicted as being illegal in places where they were undertaken.

Study Design

We designed our study as an online survey consisting of six sections: 1) primary use of the anonymous network; 2) measuring the effects of specific cases, technical schemes and entities on users' reported behaviors; 3) motivations for using anonymous networks; prior knowledge of 4) abusive activities and 5) investigators' existing de-anonymization practices; and 6) demographics. Only the first two sections were compulsory. Section 2 was the only section designed as between-subjects. Section 2 was aimed at investigating the effects of three independent variables: cases (5), schemes (5) and entities (5). The combinations of these three variables yielded 125 scenarios, which we divided between 5 user groups. We ensured that each participant encountered one scenario only once in the study to minimize learning and confounding effects.

Procedure

Participants willing to take our study were directed to a Qualtrics link. They were first asked about the anonymous services they used, their "most important or needed" i.e. primary use and frequency of usage of anonymous services (Section 1). Participants were then randomly assigned to 1 of the 5 groups.[3] Each group contained 25 unique scenarios (Section 2). The order of the scenarios presented in each group was randomized across participants. For each scenario, participants were asked to assume that the anonymous network they used had introduced the described functionality to address the type of case presented, and then asked two questions. First, how would their own anonymous network usage change for their primary activity? Participants could select from options presented on a Likert scale, ranging from decrease to no change in usage (Figure 1). Second, what reason(s) applied to their change in usage or lack thereof? Participants could select from a randomized list of options and write their own reasons.

We asked participants to select their motivations for using anonymous networks from 20 randomized options synthesized from prior work (Section 3).[4] Participants were asked if they had prior knowledge about any "malicious, criminal or unethical" uses of anonymous networks (Section 4). Those aware were then asked to identify any relevant activities they knew about. We also asked if participants were aware of existing practices by investigators such as law enforcement to exploit software vulnerabilities for de-anonymizing certain anonymous users, and how knowledge of this practice affected their usage of anonymous networks (Section 5). Finally, we inquired about age, gender, education level, technical skills, employment status, residence and nationality (Section 6).

Translating Technical Schemes into Testable Scenarios

We analyzed proposed technical schemes (i.e. both revocable anonymity schemes that aim to selectively trace certain users, and access-limiting schemes that only seek to block or limit the rate of access of some users) for deterring anonymous abuses along three dimensions using a systematized framework [37]. First, we analyzed the *goals* of each scheme to examine how it addresses the prevention, detection, evidence, judgement and punishment aspects of countering abuse. Second, we examined how *information* about potential abuses was identified and disclosed in each scheme. Third, we analyzed how each scheme addressed potential abuses with automated or mediated *actions* implemented by centralized or decentralized entities. After this initial analysis, three researchers abstracted the technical details and implications of the proposed schemes to derive their similarities and differences over six sessions between February 1 and March 7, 2019. We then refined the abstract descriptions to be comprehensible to users of different technical backgrounds while still reflecting the overall functionalities and aims of the original proposals.

We derived five types of schemes. Two involved anonymity revocation by trusted third parties: "Anonymity revocation by 1" i.e. one entity can revoke a user's anonymity [26, 32, 47], and "Anonymity revocation by 3" i.e. three entities can revoke anonymity only by consensus among themselves as done in distributed revocation schemes [19, 31, 73, 77]. One involved blocking with the consent of a trusted third party: "Blocking with TTP" [70]. Two involved access limitations by service providers: "Blocking" [4, 5, 6, 16, 69] and "Rate-limiting" [52].

We chose five commonly reported use cases of anonymous networks that various entities may deem worth addressing. Particularly, we were interested in finding out whether there is any distinction in the way users regard computer attacks ("spam" and "phishing"), which are regarded as illegitimate uses of anonymous networks as opposed to those concerning free expression ("illegal communication" and "illegal reporting" on censored topics), which are deemed legitimate in democratic societies, but criminalized by some authoritarian regimes. The remainder case involved the illegal exchange of drugs ("illegal drugs").

To allow for a diverse set of potential enforcing entities, we included government agencies in the user's country of residence (e.g. appropriate judicial bodies), commercial services (among Google, Comcast or Cloudflare), international non-profit organizations or NGOs,[5] anonymous nodes of the network (e.g. volunteer-run Tor relays), and organizations administrating the anonymous network (e.g. the Tor Project). For schemes involving decision-making by third parties, the third party was one of these five entities. In schemes involving decision-making by service providers alone (i.e. Blocking and Rate-limiting), the entity was in charge of deciding to limit a user's access to the anonymous network altogether if sufficient service providers set access limitations for that user. For anonymity revocation by three entities via consensus, the entities involved were all of the same type, e.g. three anonymous nodes.

Finally, to understand why users might change their usage of anonymous services in response to various anti-abuse mechanisms, two researchers analyzed how each scheme could be exploited beyond its intended goals over 6 sessions. By evaluating how the different cases, schemes and entities involved may deter usage, we derived a list of 10 potential reasons to present to participants.

Participant Recruitment

We launched our survey after receiving ethical approval from the Institutional Review Board at MIT. We primarily targeted Tor

[3]We used the *Randomizer* element in Qualtrics' *Survey Flow* both to randomly and evenly assign participants within the 5 groups. We used Qualtrics' *Quotas* to ensure equivalent participant numbers in all 5 groups.

[4]We examined prior research on why people seek anonymity [38, 39, 44, 76] and extracted a list of reported motivations from each paper. We then compared and consolidated all reported motivations into 20 options.

[5]We varied the NGOs presented for each case type, e.g. "The SpamHaus Project" (spam), "a member of the Anti-Phishing Working Group" (phishing), "Reporters without Borders" (reporting), "Access Now" (communication) and "The World Federation Against Drugs" (illegal drugs).

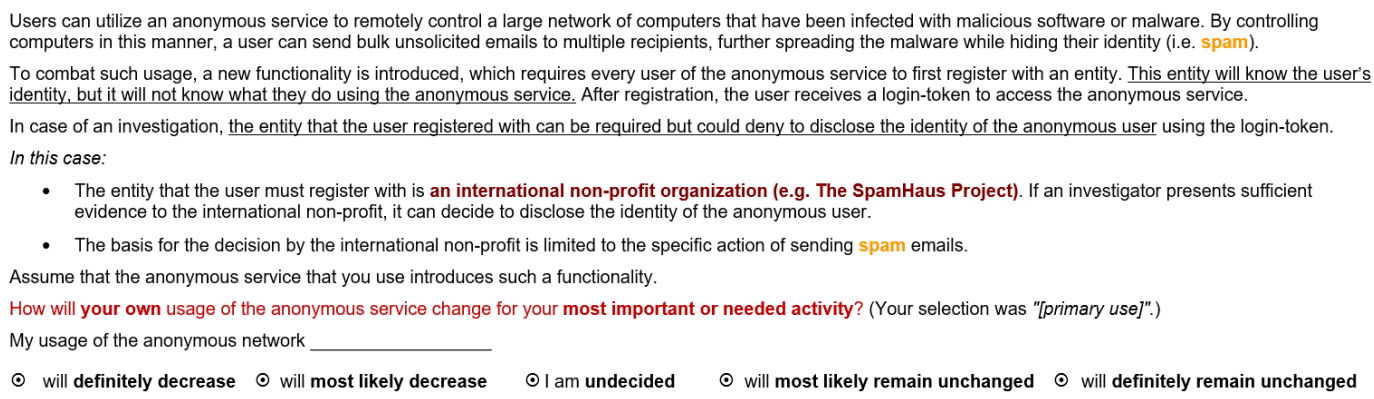

Figure 1. An example of a scenario involving spam (case), rate-limiting (scheme) and an anonymous node (enforcing entity) as shown to participants.

users via social media, online forums and a Tor-specific mailing list, including through help from the Tor Project. Participants who completed the survey were offered remuneration using a separate form to unlink their responses and respect their anonymity. Our survey ran from March 28 to May 7, 2019.

Study Validity

To ensure that participants understood our survey questions and scenarios consistently, we tested the entire study with ten people of varying ages, education levels, genders, employment statuses and technical backgrounds. After completing the survey, these participants were asked specific questions, e.g. *"What do you understand by '[survey question]'?"* and *"Could you walk me through how this scenario works?"*. These systematic probes [27] were targeted at evaluating how their interpretations matched our intended meaning. This allowed us to both simplify wording for non-technical users and include specific implementation details to allow more technical users to understand the implications of the schemes.

Data Validity

To ensure that participants did not randomly respond to our scenarios, we incorporated attention checks [12] in the form of two repeated scenarios. These were used to validate users' responses and remove participants with inconsistent answers. The attention check responses were removed from the data-set prior to analysis.

Data Analysis

Sections 1, 3, 4, 5 and 6 were analyzed by aggregating the number of responses for each answer choice. Participants' reported usage changes in Section 2 were analyzed via one-way ANOVA. This method was used to test if there was a statistically significant difference in participants' reported behaviors between scenario conditions. Participants' reasons for their reported behaviors were aggregated for all answer choices. We coded participants' open-ended reasons using an iterative process to identify recurring themes [15]. After two coders disjointly coded an agreed random sample of participants' responses, they convened to consolidate an initial set of codes. Then the two coders re-coded all qualitative data on open-ended reasons and calculated the Cohen's Kappa.

RESULTS

Participants

In total, 331 participants began the survey but only 100 completed all scenarios. Of these, 54 participants requested remuneration and all those who confirmed payment means received at least $10. Because being paid required disclosing PII such as email, some participants did not opt-in. As an added incentive, we randomly selected 7 participants for additional payments of $40 (5) and $90 (2). Among the 100 completed responses, 75 responded consistently to both attention checks (15 per group).[6]

Nine participants were female (avg. age 33.5), 48 were male (avg. age 32.6), and 7 chose "Other" (avg. age 36.6) while the remainder did not disclose their gender (11) and age (9). Education levels varied from having no diploma (4) to having completed high school (19), college or university (26), and post-graduate work (17) whereas 9 did not disclose their highest completed education level. Employment status varied from unemployed (10) to self-employed (13), part-time (9), full-time (25), other (7) and undisclosed (11). Participants described themselves as "very technical" (26), "fairly technical" (24), "somewhat technical" (19), "slightly technical" (1), "not at all technical" (1) or did not disclose their technical skills (4).

Among participants who answered, the largest number both lived (25) and were citizens (20) of the USA, followed by Germany (5) and Canada (4). Two reported multiple nationalities. Of those who disclosed both countries of residence (51) and nationality (46), all but four lived and were from the same country; four lived in the US but were from Tunisia (1), Italy (1) and India (2). Other countries represented included Bulgaria, Catalonia, China, Cyprus, France, Iran, Ireland, Mexico, Portugal, Russia, Singapore, Slovenia, Spain, Sweden, the UK and Ukraine.

Types of anonymous services used

All participants used the Tor network. Two (P24, P64) accessed Tor only via Orbot. While several participants used Tails (28) and Orbot (11),[7] anonymous networks such as I2P (8) and Freenet (6) were used less frequently. Other anonymous services used included Briar, Ricochet, Torphone, Onion Share and Whonix.

Frequency of use of anonymous networks

Twenty-seven participants used anonymous networks for $\sim$25% of their online activities. Equivalent numbers relied on anonymous

[6]Upon obtaining 15 participants who passed our attention checks in each of the five groups (which allowed us to obtain 375 responses to each type of case, scheme and entity), we terminated the study.

[7]Tails is an operating system that forces all internet connections via the Tor network, whereas Orbot is an Android application for accessing Tor.

MS	N	Description	MS	N	Description
M1	49	To keep different aspects of my identity separate from one another	M2	24	To prevent harassment
M3	47	To contribute to the anonymous community for the benefit of other users	M4	63	To avoid invasive use of my personal information
M5	51	To avoid revealing my personal information for reasons I consider inappropriate	M6	15	To avoid financial attacks
M7	57	To prevent companies from making money from my personal information	M8	17	For fear of my internet access being revoked
M9	40	To avoid discrimination based on my identity or my online activities	M10	43	To avoid unknown threats
M11	36	For fear of exposure for political associations, opinions and/or related activities	M12	26	For fear of legal sanctions, e.g. imprisonment.
M13	54	To avoid commercial tracking of my participation to online communities or projects	M14	15	To avoid accountability for my past actions or statements
M15	22	To avoid potential retaliation from a business/service after I leave an online review	M16	55	To avoid potential misuse of my personal information
M17	45	To avoid losing control of my personal data and the ability to delete my information	M18	53	For safety from unknown surveillance for unknown reasons
M19	32	To avoid repercussions for my online activities that may be perceived as unlawful or unethical	M20	24	For safety against physical harms against one-self and/or loved ones

Table 1. Motivations selected (MS) for seeking anonymity and number of participants (N) who selected each. Participants could select multiple motivations.

networks for $\sim$75% (13) and less than 5% (13) of their online activities. Twelve conducted all (i.e 100%), and 10 performed half (i.e. $\sim$50%) of their online activities via anonymous networks.

Primary uses of anonymous networks

The most frequent activities reported by participants were *"searching for information"* (30), *"communicating with others"* (14) and *"accessing censored content"* (12). Nine used anonymous networks for *"sharing personal views"* (6), *"buying/selling items"*(2) and *"participating in social networks or communities"* (1). The remainder 10 wanted to avoid online tracking for specific or general uses, including personal uses, activism, and research or testing.

Motivations for Using Anonymous Networks

All participants gave their motivations for seeking anonymity. Table 1 shows that participants most frequently selected motivations concerning the protection of personal information, and avoiding commercial or unknown tracking (**M4**: 63, **M7**: 57, **M16**: 55, **M13**: 54, **M18**: 53, **M5**: 51 and **M1**: 49). Forty-seven wanted to contribute to enhance others' anonymity (**M3**). Only 15 participants indicated their desire to avoid accountability (**M14**).

Effects of Cases, Schemes and Entities

Participants reported different changes in usage of anonymous networks according to several factors (Figure 2). The descriptive statistics associated with each such factor are shown in Table 2.

Effect of Case

A one-way ANOVA test yielded a statistically-significant effect ($F(4,1870) = 3.31, p = 0.0104$), showing that the type of case does have an effect on the degree to which participants' primary use of the anonymous network changes. A Tukey post-hoc test revealed a significantly larger decrease in anonymous network usage for cases involving illegal reporting ($\mu = 2.17$) relative to spam ($\mu = 2.45, p = 0.0463$), as shown in Figure 2(a) and Table 2. Moreover, a t-test found significant larger decreases in anonymous network usage for cases involving illegal reporting relative to spam ($p = 0.0059$), phishing ($\mu = 2.40, p = 0.0243$) and illegal drugs ($\mu = 2.38, p = 0.0413$) as well as for cases involving illegal communication ($\mu = 2.18$) relative to spam ($p = 0.0081$) and phishing ($p = 0.0318$). We observed no pairwise significant differences among other types of cases. This shows that participants were more opposed to cases countering reporting and communication of censored topics relative to spam and phishing attacks.

Effect of Scheme

We found a significant difference in the degree to which the type of scheme implemented affected participants' change in anonymous network usage for their primary activity (Figure 2(b)), as determined by one-way ANOVA ($F(4,1870) = 8.36, p < 0.0001$).

	Factor	μ	σ	95% CI
CASES	Spam	2.45	1.38	[2.31, 2.59]
	Phishing	2.40	1.39	[2.26, 2.54]
	Illegal drugs	2.38	1.39	[2.24, 2.52]
	Illegal communication	2.18	1.35	[2.05, 2.32]
	Illegal reporting	2.17	1.37	[2.03, 2.31]
SCHEMES	Anonymity revocation by 1	2.06	1.34	[1.93, 2.20]
	Anonymity revocation by 3	2.21	1.37	[2.07, 2.35]
	Blocking with TTP	2.26	1.39	[2.12, 2.40]
	Rate-limiting	2.51	1.36	[2.38, 2.65]
	Blocking	2.54	1.39	[2.40, 2.68]
ENTITIES	Administrator	2.43	1.41	[2.29, 2.58]
	International NGO	2.42	1.38	[2.28, 2.56]
	Anonymous node	2.38	1.35	[2.25, 2.52]
	Commercial service	2.21	1.40	[2.07, 2.36]
	Government agency	2.13	1.34	[2.00, 2.27]

Table 2. Means (μ), standard deviations (σ) and 95% confidence intervals for all three types of factors: cases, schemes and entities.

A Tukey post-hoc test revealed significant larger decreases in anonymous network usage for schemes involving anonymity revocation by one entity ($\mu = 2.06$) relative to blocking ($\mu = 2.54, p < 0.0001$) and rate-limiting ($\mu = 2.51, p < 0.0001$) by service providers, anonymity revocation by three entities via consensus ($\mu = 2.21$) relative to blocking ($p = 0.0085$) and rate-limiting ($p = 0.0187$), and blocking with the consent of a third party ($\mu = 2.26$) relative to blocking ($p = 0.0445$). Additionally, a t-test found significant larger decreases in anonymous network usage for schemes involving anonymity revocation by one entity relative to service providers blocking ($p < 0.0001$), rate-limiting ($p < 0.0001$) and blocking with the consent of a trusted third party ($p = 0.0486$), anonymity revocation by three entities via consensus relative to blocking ($p = 0.0010$) and rate-limiting ($p = 0.0022$), and blocking with the consent of a third party relative to blocking ($p = 0.0056$) and rate-limiting ($p = 0.0114$). We found no pairwise

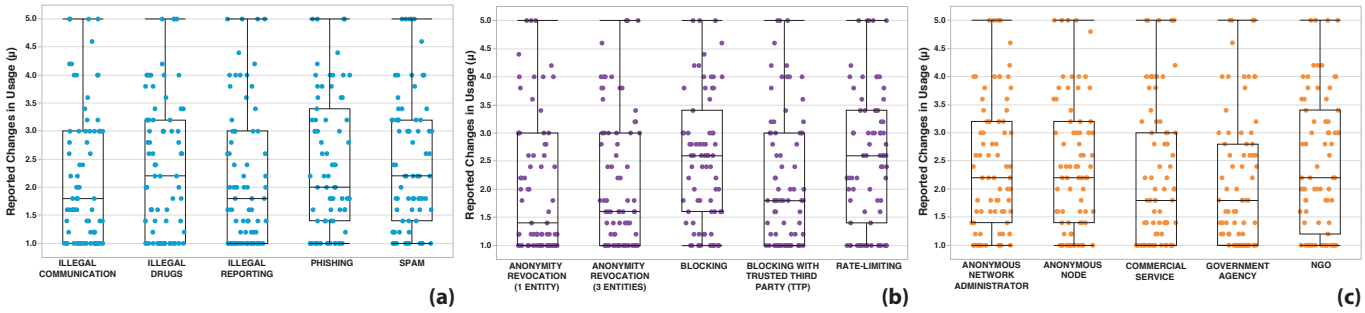

Figure 2. Mean reported changes in usage of anonymous networks for each type of (a) case, (b) scheme and (c) entity (1: *definitely decrease*, 2: *most likely decrease*, 3: *undecided*, 4: *most likely unchanged*, 5: *definitely unchanged*).

significant differences among other types of schemes. This shows that schemes involving anonymity revocation and trusted third parties schemes would deter usage more so than access-limiting schemes, which are directly implemented by service providers.

Effect of Entity

A significant difference was found between participants' self-reported changes in anonymous network usage based on the type of decision-making entity, as determined by one-way ANOVA ($F(4,1870) = 3.64, p = 0.0058$). A Tukey post-hoc test revealed a significantly larger decrease in usage when government agencies ($\mu = 2.13$) are in charge relative to anonymous networks administrators ($\mu = 2.43, p = 0.025$), and NGOs ($\mu = 2.42, p = 0.0317$), as shown in Figure 2(c) and Table 2. A t-test also found a significantly larger decrease in usage when government agencies are in charge relative to anonymous network administrators ($p = 0.0030$), NGOs ($p = 0.0039$) and anonymous nodes ($\mu = 2.38, p = 0.0127$) in addition to a significantly larger decrease in usage when commercial services ($\mu = 2.21$) are in charge relative to anonymous network administrators ($p = 0.0297$) and NGOs ($p = 0.0362$). No other pairwise significant differences were observed. This shows that participants distrusted government and commercial entities more than other enforcing entities.

Participant Profiles

What influenced participants' reported changes in anonymous network usage? We examined whether participants always reported the same change in usage (i.e. decrease, undecided, or no change) or reported variable changes (e.g. ranging from decrease to no change, etc.) for each type of factor. Our analysis revealed five distinct user profiles:

- **Anonymity-conscious users (27)** reported a **decrease** in usage regardless of the types of entities, schemes or cases presented. While 18 participants reported a decrease for all scenarios, nine had one or two exceptions for which they reported no change or were undecided, which typically involved blocking or rate-limiting cases of spam or phishing as enforced by anonymous network administrators, NGOs or anonymous nodes.

- **Anonymity-indifferent users (9):** Six users reported that their usage will **remain unchanged** regardless of the entities, schemes or cases involved. Three others also reported no change with one exception for which they were undecided; these involved anonymity revocation or blocking with the consent of a third party enforced by anonymous nodes and an NGO for phishing, illegal reporting and illegal drug sale cases.

- **Factor-specific users (31)** were affected by one or more factors, being case-conscious, scheme-driven and/or entity-based.
  - **One-factor users (15)** responded consistently for only one factor, i.e. type of case (9), scheme (4) or entity (2) while reporting variable changes in usage for the other two factors. Case-conscious users typically indicated no change in usage for phishing or illegal drug sale cases or a decrease in usage for cases countering free expression, as shown by the lowest means for illegal reporting and/or communication cases (Figure 2(a)). Scheme-driven users reported a decrease in usage for revocable anonymity schemes and/or no change in usage for one or more of the access-limiting schemes. Entity-based users typically reported a decrease in usage for government agencies.
  - **Two-factor users (12)** responded consistently for two factors, i.e. types of case and scheme (5), case and entity (4), and scheme and entity (3) while having variable responses for the remaining factor. Of these, case-conscious users indicated a decrease in usage for illegal reporting and/or communication cases, or no change in usage for spam and phishing cases. Scheme-driven users reported a decrease in usage for revocable anonymity schemes and/or were undecided about access-limiting schemes. Entity-based users reported a decrease in usage for government and/or commercial entities, or reported being unchanged or undecided for one or more of the other three entities.
  - **Three-factor users (4)** responded consistently for all 3 factors simultaneously. They reported either a decrease in usage for illegal reporting or communication cases, revocable anonymity schemes and government agencies as enforcing entities, or no change in usage for spam and phishing cases, access-limiting schemes and NGOs.

- **Undecided users (3)** were undecided regardless of the factors involved. Two had an exception for which they reported a decrease in usage; these included a government agency blocking users involved in an illegal drug case, and an anonymous network administrator blocking users for illegal communication.

- **Uncategorized users (5)** had variable reported changes in usage for all entities, schemes and cases, so their behavior cannot be explained by any factor shown in Figure 2.

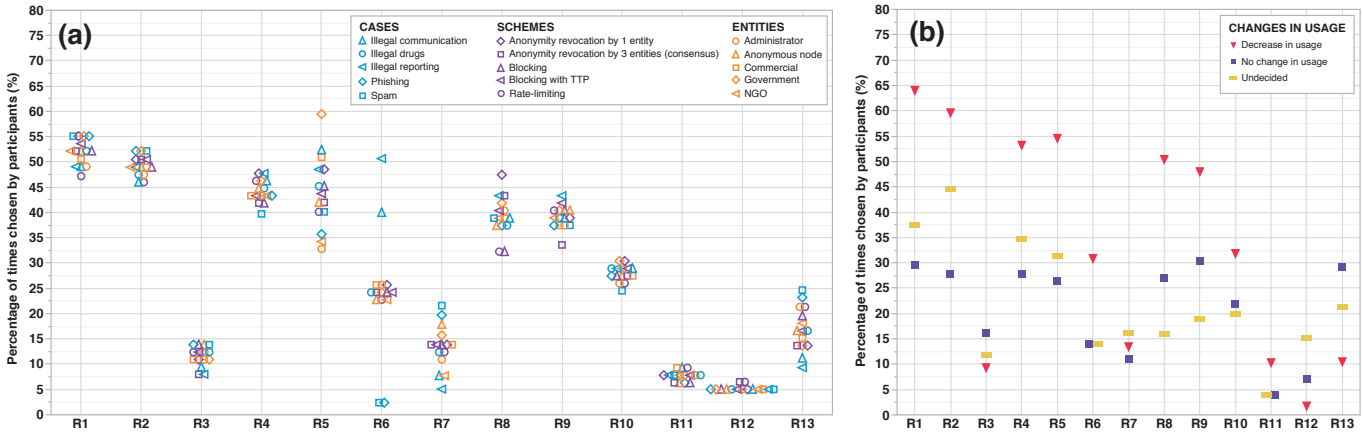

Figure 3. Reasons for participants' reported changes in usage of anonymous networks divided by (a) type of factor and (b) reported change in usage.

RS	N	%	Description	RS	N	%	Description
R1	970	51.73	This functionality can be abused and applied to other types of uses.	R2	924	49.28	I cannot verify that this functionality is used to counter only *[case]* and not other uses.
R3	214	11.41	I would be more comfortable if this functionality involved consensus by more than only *[number & type of entity]*.	R4	824	43.95	This functionality can negatively affect other anonymous users, not just me.
R5	822	43.84	I do not trust the judgement of *[entity]* about *[case]*.	R6	449	23.96	I think that *[case]* should not be countered.
R7	249	13.28	*[case]* should be countered, but not by *[entity]*.	R8	733	39.09	An anonymous user's identity should not be revealed at any cost.
R9	731	38.99	All users should have equal anonymous access.	R10	515	27.47	There is no mechanism to appeal the entity's decision while remaining anonymous.
R11	144	7.68	Other (open-ended response)	R12	97	5.17	I do not wish to disclose my reason(s).
R13	316	16.85	I understand the value of this functionality for this scenario.				

Table 3. Descriptions and overall statistics of the reasons selected (RS) by participants for their reported changes in usage in response to all scenarios. For each scenario, factors italicized in brackets contained the *[case]*, *[scheme]* and/or *[entity]* appearing in the scenario.

Reasons for Changes in Anonymous Network Usage

Participants reported several reasons for their changes in usage of anonymous networks or lack thereof. Figure 3(a) shows the percentage of times participants chose each reason depending on the types of factors involved, Figure 3(b) shows the selected reasons based on participants' reported change in usage and Table 3 shows the selected reasons' descriptions and overall statistics. Our thematic analysis of participants' open-ended reasons (**R11**) resulted in the set of codes described in Table 4 (Cohen's kappa $k = 0.731; p < 0.0001$).

Altogether, the most frequently selected reasons were about the potential for abuse of various technical mechanisms (**R1**:51.73%) and their negative impact on others (**R4**:43.75%), distrust in the judgement of different entities (**R5**:43.84%) and the ability to verify their actions (**R2**:49.28%), and the right to maintain anonymity (**R8**:39.09%, **R9**:38.99%). In general, all participants more frequently indicated a lack of trust (**R5**) in government (59.5%) or commercial (50.9%) entities relative to anonymous network administrators (32.8%) and NGOs (33.9%). Relatively more participants thought that reporting (50.7%) and communication (40%) on censored topics should not be countered if they were criminalized (**R6**) as opposed to spam (2.4%) and phishing (2.4%) attacks.[8]

[8]All participants (except undecided users) selected R6 more frequently for cases involving free expression (i.e. communication & reporting).

Code & Description	N	Code & Description	N
O1: Anonymity compromised	53	**O2:** Lack of usefulness or desirability	29
O3: Distrust	23	**O4:** Unwillingness to participate	20
O5: New security risks	17	**O6:** Ineffectiveness	13
O7: Resentment	11	**O8:** Disgust	10
O9: Censorship	8	**O10:** Incomprehensible	5
O11: Do not care	5	**O12:** External influence	4

Table 4. Thematic coding of open-ended reasons & number of mentions (N). Of the 144 total open-ended responses (R11), some had multiple codes.

Anonymity-conscious users frequently raised concerns about the abuse of various mechanisms beyond the specified cases (**R1**:65.93%) to compromise anonymity (**O1**:46) and negatively impact others (**R4**:51.41%, **O6**:5), e.g., *"I detest spammers, but don't want anything compromising my pseudonyms"* (P42), and *"There is no way to limit this system solely to the illegal drug market without tracking a user's access in general..."* (P50). They also frequently raised concerns about the right to maintain anonymity (**R8**:51.26%, **R9**:46.96%), and the new security risks (**O5**:17) associated with various schemes, e.g. *"this is an increase in the risk surface area, and opens the tor project...to political attacks..."* (P47). Some indicated an unwillingness to participate (**O4**:18)

in certain schemes, e.g. *"I don't want to register with any entity"* (P42). Others pointed out the lack of usefulness or desirability of some schemes and cases (**O2**:22), e.g. *"The random computational puzzles take time to solve."* (P69), *"I am pro-illegal drugs. People should be able to buy, sell, use and trade them..."* (P47), and in response to a scenario involving illegal communication, *"...An anonymity network that attempts to provide only conditional anonymity is like a democracy where voting for certain candidates gets you executed. Either you have anonymity, or you don't: there is no middle ground here, and trying to forcibly establish one only results in inevitable abuse, and eventual abandonment once enough users realize the betrayal."* (P50).

Anonymity-conscious users also frequently reported distrust (**R5**:50.96%, **O3**:20) and resentment (**O7**:10) towards various entities, especially government and commercial entities, in addition to concerns about being unable to verify their actions (**R2**:57.48%) or appeal their decisions while remaining anonymous (**R10**:24.3%). While some expressed distrust for specific entities, e.g. *"judicial bodies doesn't approve revolutions, but revolutions are much needed these days."* (P46), others did so for all entities, e.g. *"Allowing any entity the ability to regulate communication invariably leads to the entity blocking communications about problems or criticisms of such an entity"* (P27) and warned of external influences (**O12**:3), *"...all non-government bodies can just be forced without warrant to surrender data."* (P27).

Among factor-specific users, those influenced by all three factors simultaneously more frequently selected reasons about the potential abuse (**R1**:81%) of various technical mechanisms and their negative impact on others (**R4**:79%), distrust of entities (**R5**:64%) and inability to verify (**R2**:88%) or appeal (**R10**:83%) their decisions, and the right to maintain anonymity (**R8**:44%, **R9**:50%) than users influenced by only one or two factors. Factor-specific users also pointed out the ineffectiveness (**O6**:3) of some mechanisms, e.g. *"spam classifiers aren't very accurate"* (P15), and the lack of usefulness or desirability (**O2**:2) of others, e.g. *"Registration of every user defeats the purpose of the network anonymity"* (P36). Some indicated concerns about incomprehensibility (**O10**:4) and censorship (**O9**:5), e.g. *"What is 'illegal communication'? Sounds like censorship like China doesn't allow communication with human right activists, press or uncensored messengers/e-mail-provider."* (P37) in addition to distrust (**O3**:2), e.g. *"ANY entity, non profit or otherwise is ran by people. people are inherintly biased and cannot be expected to apply rules fairly and unanimously"* (P54).

Anonymity-indifferent users most frequently selected only reasons regarding the right to maintain anonymity (**R8**:40.88%, **R9**:42.67%). Undecided users most frequently opted to not disclose their reason(s) (**R12**:48.0%). Among uncategorized users, one participant raised concerns about anonymity being compromised (**O1**:4), including for schemes involving only blocking by service providers, *"While ZKPs [zero-knowledge proofs] are good, this feature would still partition the anonymity set of the network into blocked and non-blocked users"* (P58).

Impact of Prior Knowledge

Sixty-seven participants were aware of various abuses of anonymous networks while the remainder 8 had no such awareness. We categorized participants' free-form responses into three main types of abuses they mentioned as shown in Figure 4. Table 5 shows that participants aware of more serious abuses (i.e. physical harms and illegal exchanges) reported a greater decrease in their anonymous network usage in response to various technical anti-abuse mechanisms than those unaware or aware of non-physical harms.

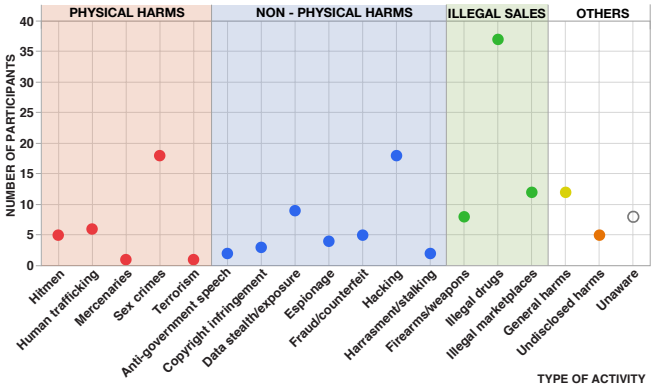

Figure 4. Types of "malicious, criminal or unethical" activities conducted via anonymous networks, as identified by participants. Some participants mentioned several different activities. Sex crimes involve materials containing illegal or child pornography/abuse, rape, etc. Hacking covers botnet attacks, spam, phishing, ransomware, money tumbling, etc. Fraud includes counterfeit documents, money laundering, etc. Espionage includes dumping government and corporate secrets. Illegal marketplaces include illicit services like organ markets and crime-for-hire. General harms are non-specific mentions of "abuse", "criminal activities", etc. Undisclosed harms include instances where users reported awareness but did not reveal any abuses.

All but one participant reported observing the activities they mentioned on forums and chat rooms accessible via anonymous networks. One participant witnessed similar uses in the physical world: *"Streets of my city have stickers with *.onion addresses promoting illegal drug retail"* (P20). Two participants added personal views, stating, *"...I don't believe online markets should be banned. They build a safe space and a community to share opinions and reviews for substances"* (P23), and, *"I am familiar with...markets such as silkroad, agora, etc. They were...typically how the media tries to portray every user of the web who likes anonymity"* (P70).

Type of abuse	n	μ	σ	95% CI
Non-physical harms	27	2.43	1.42	[2.32, 2.54]
Illegal sales/exchanges	43	2.20	1.28	[2.12, 2.27]
Physical harms	21	2.16	1.32	[2.05, 2.27]
General harms	12	1.96	1.20	[1.83, 2.10]
Undisclosed harms	5	1.66	0.92	[1.50, 1.83]
Unaware	8	2.87	1.74	[2.63, 3.11]

Table 5. Reported mean changes in usage of anonymous services for participants aware of different types of abuses.

Impact of Investigators' Existing Practices

Fifty-eight participants indicated being aware of investigators' practices to identify certain anonymous users via software vulnerabilities whereas 17 participants reported being unaware.

Participants aware of investigators' existing practices (58)
Table 6 summarizes the responses of such participants. Thirty-four participants were affected by investigators' practices in various ways, e.g. *"I try routing most...of my traffic over anonymous services. That way, metadata is much noisier to correlate against any particular internet activity"* (P39) and *"I keep it as up to date as I can. I also use it a bit less than I otherwise would"* (P63). Eighteen participants stated that investigators' use of software vulnerabilities had no effect on their anonymous network usage. Of these, five believed that they had not breached any laws or had nothing to worry about, e.g. *"I do not use Tor for anything that makes me afraid of investigators"* (P35). Others gave multiple reasons, including *"No. That privacy can be compromised does not mean I should give up entirely"* (P55). Six participants did not directly answer how investigators' practices affected their own anonymous network usage. They made comments, e.g. *"Makes me feel uneasy. I neither have faith in these agencies' intentions, nor in their competence to keep these bugs secret"* (P60), and *"...We all end up paying for those who decided to do illegal stuff"* (P36).

Participants unaware of investigators' existing practices (17)
When asked how knowledge of investigators' practices would change their own anonymous network usage for their primary activity, 11 said their usage would *"remain unchanged"*, 4 said their usage would *"decrease"*, 1 was undecided and 1 did not respond.

DISCUSSION AND IMPLICATIONS

We wanted to understand how technical schemes developed to address anonymous abuses may impact the legitimate uses of anonymous services. We found that a number of social and technical factors affect users' preferences and should considered in the design and enforcement of potential counter-abuse mechanisms.

Relation to prior work

By grounding technical anti-abuse schemes in concrete scenarios with specific cases and entities, we empirically demonstrate participants' greater opposition to cases countering communication and reporting of censored topics (which are crimes in some jurisdictions) relative to other cases. Our work supports the notion that free expression without tracking and censorship are the intended [20, 33, 34, 60] use cases of anonymous networks. In showing the relative distrust of government and commercial entities, we extend prior work on users' motivations for seeking anonymity [38, 39, 44, 76], which depicted such entities as oft-reported threat actors. An early discussion of the technical issues facing revocable anonymity schemes identified fundamental security flaws in their architecture [29]. It suggested that the potential for its abuse might lead users to place less trust in the anonymous network even when the revocation mechanism is not exercised. Our results empirically show that revocable anonymity schemes indeed deter the use of anonymous networks for several intended and legitimate purposes. We also show that this decrease in usage is driven by several factors, including the inability to limit the counter measures to specific abuses and distrust in the judgement of enforcing entities involved. Our study also corroborates prior findings on the criminal and unethical content found via scanning onion services [14, 40, 57, 79] since our participants reported a wide range of harmful activities they had observed or become aware of.

Revocable anonymity: security and trust implications

Although revocable anonymity schemes have not been implemented for the Tor network, the AN.ON communication system deployed a feature to track future connections from users in case of a valid court order. This revocable anonymity feature came in response to a 2003 legal request against a server hosting child pornography in Germany and was criticized by many users despite being made transparent via changes to the open source code [8]. While the AN.ON case highlights the precarious balance between the two needs of strong anonymity and crime prevention, our study shows that revocable anonymity mechanisms would deter several legitimate uses of anonymous networks. This is evidenced by the significantly greater decrease in anonymous network usage associated with revocable anonymity schemes and participants' more frequent concerns about anonymity being compromised for such schemes relative to access-limiting schemes.

Since schemes involving anonymity revocation and third parties alter the trust model of decentralized anonymous networks by introducing new trusted parties or giving existing entities greater power, participants' concerns about security risks and entities being susceptible to external influence are plausible. Such concerns, especially prevalent among anonymity-conscious users, are not unfounded in light of companies succumbing to pressure from foreign governments to censor specific content, as in the case of Apple removing VPN apps from its China App store to comply with Chinese censorship [63]. Even in cases where the third party enforcing revocable anonymity is well-trusted, they can make the overall system vulnerable to abuse or political meddling, as has

CHANGE IN USAGE (34)	PARTICIPANTS	NO CHANGE IN USAGE (18)	PARTICIPANTS
Being more cautious and vigilant in setting up/using anonymous networks	P1*, P6$^{\triangleleft\diamond}$, P14$^{\diamond}$, P21*, P23*, P53*, P69*	No expectation of being targeted by investigators	P20$^{\diamond\triangleright}$, P25$^{\triangleleft}$, P35$^{\diamond}$, P56*, P65‡
Keeping software updated	P4*, P13$^{\triangleleft}$, P21*, P23*, P29$^{\triangleleft\diamond\triangleright}$, P31$^{\triangleleft\diamond}$, P32$^{\triangleright}$, P37$^{\diamond}$, P53*, P59$^{\triangleleft}$, P63$^{\diamond}$	Not having many highly critical personal uses of anonymity	P28‡, P30*, P46*
Using multiple layers of security (i.e. additional tools and add-ons)	P1*, P10$^{\triangleright}$, P34$^{\triangleleft}$, P39$^{\triangleleft}$, P40*, P41*, P52*, P58‡, P67*, P70$^{\triangleleft\triangleright}$	No reason provided	P15$^{\triangleleft\triangleright}$, P26$^{\triangleleft\diamond}$, P64†, P72‡, P74$^{\triangleleft\diamond}$, P75$^{\triangleleft}$
Only using anonymous tools via public networks	P29$^{\triangleleft\diamond\triangleright}$, P31$^{\triangleleft\diamond}$, P32$^{\triangleright}$	Means of circumventing investigators exist	P16$^{\pm}$
Not connecting personal data to online persona	P18*, P29$^{\triangleleft\diamond\triangleright}$, P32$^{\triangleright}$, P59$^{\triangleleft}$	Ability to check open-source code	P51*
Increasing the use of anonymous networks	P39$^{\triangleleft}$	Unwillingness to give up privacy entirely	P55 $^{\pm}$
Decreasing the use of anonymous networks	P3$^{\triangleleft\diamond\triangleright}$, P50*, P52*, P63$^{\diamond}$	Disturbed by investigators' practices	P2$^{\triangleleft\triangleright}$
Avoiding JavaScript and vulnerable software	P9,$^{\triangleleft\diamond}$ P13$^{\triangleleft}$, P27*, P37$^{\diamond}$, P43*, P47*, P54$^{\triangleleft}$, P57$^{\pm}$, P68*	**OTHER (6)**	P17$^{\triangleleft\diamond\triangleright}$, P19*, P36$^{\triangleleft\diamond}$, P60†, P66$^{\pm}$, P73$^{\pm}$

Table 6. Impact of investigators' use of software vulnerabilities on participants' usage of anonymous networks. Participants' profiles are also shown: *Anonymity-conscious; $^{\triangleleft}$case-conscious; $^{\diamond}$scheme-driven; $^{\triangleright}$entity-based; $^{\pm}$anonymity-indifferent; † undecided; ‡uncategorized.

been the case with Interpol being politically influenced by authoritarian regimes to arrest dissidents and human rights activists [3].

Since participants aware of more serious abuses (i.e. physical harms and illegal exchanges) reported a greater decrease in their anonymous network usage, this suggests that such participants view the proposed technical mechanisms as making anonymous networks more insecure or susceptible to abuse.[9] While most participants aware of investigators' use of software vulnerabilities improved their security practices as a result, most of those unaware reported no change in their anonymous network usage upon finding about investigators' existing deanonymization methods.[10] This suggests that existing investigative practices of de-anonymization pose less risks for users than built-in lawful access mechanisms to selectively revoke anonymity, which is consistent with arguments by security researchers that engineered lawful access mechanisms would introduce new security risks into communication networks [1, 11]. In light of mounting attacks on anonymity [18, 42, 50, 58, 65], revocable anonymity schemes would exacerbate the security concerns already associated with anonymous networks.

Implications for design and policy

Our results have three main implications for addressing anonymous abuses. First, technical schemes should not be introduced to enable a third party to broadly target anonymous users for any type of abuse. Schemes should only counter specific well-defined abuses without infringing on users' human rights.

Second, anonymity revocation would be especially harmful if the revocation authority is a local government agency or a commercial service that could easily track user's communication. This could lead to unintended consequences, e.g. an authoritarian regime could seek to reveal the identity of anonymous activists reporting news critical of the government either by itself or by coercing other entities to do so. Anonymity revocation compromises the intended goal of anonymous networks, especially since several users seek anonymity predominantly to evade threats. Such schemes also introduce additional insecurities, rendering anonymous networks more susceptible to abuse. Hence, access-limiting schemes, which aim to only block or rate-limit abusive users, would be more consistent with the threat model of anonymous networks.

Finally, while some schemes allow anonymous users to check whether or not they have been blocked by specific service providers [6, 70], technical mechanisms proposed so far do not allow anonymous users to verify why particular abuses were addressed (e.g. why certain connections were blocked or rate-limited). To gain the trust of anonymous users, schemes should be adopted in a manner that enables verification of the decision-making criteria and the actions of the entities enforcing them. Incorporating the ability to appeal the decisions of the enforcing entity while remaining anonymous should also be considered.

[9]Several users who highlighted security risks (P47, P50), external influences (P27, P50) and distrust of entities (P1, P19, P27, P37, P42, P46, P47, P50, P52, P54, P58) in open-ended reasons mentioned physical harms (P1, P27, P46, P50, P54, P58), illegal exchanges (P1, P27, P37, P42, P46, P50, P52, P54, P58), and general harms (P47, P52).

[10]Eight participants reported a decrease in their anonymous network usage due to investigators current deanonymization practices. Of these, 4 reported being aware of such practices while 4 reported being unaware.

LIMITATIONS AND CHALLENGES

We used a survey methodology since we wanted to engage with a population of users that valued their anonymity. While this method ensured their anonymity, it also limited us in further probing participants to get more detailed responses. Given our targeted population and distribution method, we required only Sections 1 and 2 to be compulsory in order to retain participation. While 96.7% participants answered all questions in Sections 1-5, we missed one of two responses for two participants in Section 4 and for three participants in Section 5. In Section 6, only 46 (61.3%) participants provided all demographic information, 7 (9.3%) did not report any, and the remainder partially answered demographic questions.[11]

Participants could only report changes in anonymous network usage ranging from decrease to no change for our scenarios. This constraint may have biased their responses as some participants might have chosen to increase their usage in response to anti-abuse mechanisms. However, we believe that such users would leave their usage unchanged at most either because of their belief that the technical schemes would not substantially impact their anonymity or because of their lack of concern for the impact on their own anonymity. Additionally, our results might not have included more anonymity-conscious users, who may have decided against participating. Our survey platform, Qualtrics, required JavaScript, which is deactivated by the Tor Browser's highest security setting. This feature could have deterred some users from taking our survey.

CONCLUSION

Using a survey-based experiment that situated technical schemes for addressing anonymous abuses in the various social contexts in which they could be implemented, we show that different factors affect several legitimate uses of anonymous networks. Our 75 participants had five main types of profiles. While our participants were significantly less opposed to addressing spam and phishing attacks, they distrusted government and commercial entities more than other types of enforcing authorities. Our participants regarded schemes involving anonymity revocation and third parties as more undesirable than those only involving access limitations such as blocking or rate-limiting. We also found that participants with prior knowledge of more serious abuses reported a greater decrease in usage of anonymous networks in response to the anti-abuse schemes, which reflects concerns about the potential for abuse of such technical schemes. Knowledge of investigators' current deanonymization practices resulted in more participants adopting better security practices as opposed to decreasing their usage, which further indicates the greater security risks associated with revocable anonymity schemes. Since participants most frequently raised concerns about schemes being abused to negatively impact other anonymous users in a non-verifiable manner, we suggest that anti-abuse mechanisms be tailored to counter specific abuses in a manner that allows users to verify the actions of the enforcing entities and anonymously appeal particular decisions.

ACKNOWLEDGMENTS

Our thanks go to David D. Clark for his invaluable assistance and discussion on this topic. Wajeeha Ahmad and Ilaria Liccardi were supported by the William and Flora Hewlett Foundation.

[11]Some provided all demographic data except their countries of nationality (16) and residence (13) while 6 others had varying missing demographic data, e.g. missing gender, employment or education levels.

REFERENCES

[1] Harold Abelson, Ross Anderson, Steven M Bellovin, Josh Benaloh, Matt Blaze, Whitfield Diffie, John Gilmore, Matthew Green, Susan Landau, Peter G Neumann, Ronald L Rivest, Jeffrey I Schiller, Bruce Schneier, Michael Specter, and Daniel J Weitzner. 2015. Keys Under Doormats. *Commun. ACM* 58, 10 (2015), 24–26. DOI: http://dx.doi.org/10.1145/2814825

[2] Rob Klin Mark S. Frankel Rob Kling Ya-Ching Lee Al Teich, Mark S. Frankel. 1999. Anonymous Communication Policies for the Internet: Results and Recommendations of the AAAS Conference. *The Information Society* 15, 2 (5 1999), 71–77. DOI: http://dx.doi.org/10.1080/019722499128538

[3] Matt Apuzzo. 2019. How Strongmen Turned Interpol Into Their Personal Weapon. (3 2019). https://www.nytimes.com/2019/03/22/world/europe/interpol-most-wanted-red-notices.html?smid=nytcore-ios-share

[4] Man Ho Au and Apu Kapadia. 2012. PERM: Practical reputation-based blacklisting without TTPs. In *Proceedings of the 2012 ACM conference on Computer and communications security - CCS '12*. ACM Press, New York, New York, USA, 929. DOI: http://dx.doi.org/10.1145/2382196.2382294

[5] Man Ho Au, Apu Kapadia, and Willy Susilo. 2012. BLACR: TTP-Free Blacklistable Anonymous Credentials with Reputation. *2012 Network & Distributed System Security Symposium* (2012), 1–17. http://ro.uow.edu.au/infopapers/1903

[6] M. Ho Au, P. P. Tsang, and A. Kapadia. 2011. PEREA: Practical TTP-free revocation of repeatedly misbehaving anonymous users. *ACM Transactions on Information and System Security* 14, 4 (12 2011), 1–34. DOI: http://dx.doi.org/10.1145/2043628.2043630

[7] Michael Backes, Jeremy Clark, Aniket Kate, Milivoj Simeonovski, and Peter Druschel. 2014. BackRef: Accountability in anonymous communication networks. In *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, Vol. 8479 LNCS. Springer, Cham, 380–400. DOI: http://dx.doi.org/10.1007/978-3-319-07536-5{_}23

[8] Helmut Bäumler, Hannes Federrath, and Claudia Golembiewski. 2003. *Report on the Criminal Justice Process against "AN.ON-Anonymity.Online"*. Technical Report.

[9] Joshuan Bearman. 2015. The Rise and Fall of Silk Road. (4 2015). https://www.wired.com/2015/04/silk-road-1/?intcid=inline_amp

[10] Steven M. Bellovin, Matt Blaze, Sandy Clark, and Susan Landau. 2013. Going Bright: Wiretapping without Weakening Communications Infrastructure. *IEEE Security & Privacy* 11, 1 (1 2013), 62–72. DOI: http://dx.doi.org/10.1109/MSP.2012.138

[11] Steven M. Bellovin, Matt Blaze, Sandy Clark, and Susan Landau. 2014. Lawful Hacking: Using Existing Vulnerabilities for Wiretapping on the Internet. *Northwestern Journal of Technology and Intellectual Property* 12, 1 (2014). DOI: http://dx.doi.org/10.2139/ssrn.2312107

[12] Adam J. Berinsky, Michele F. Margolis, and Michael W. Sances. 2014. Separating the Shirkers from the Workers? Making Sure Respondents Pay Attention on Self-Administered Surveys. *American Journal of Political Science* 58, 3 (7 2014), 739–753. DOI: http://dx.doi.org/10.1111/ajps.12081

[13] Beatrice Berton. 2015. *The dark side of the web: ISIL's one-stop shop?* Technical Report. European Union Institute for Security Studies. DOI: http://dx.doi.org/10.2815/454889

[14] Alex Biryukov, Ivan Pustogarov, Fabrice Thill, and Ralf-Philipp Weinmann. 2014. Content and Popularity Analysis of Tor Hidden Services. In *2014 IEEE 34th International Conference on Distributed Computing Systems Workshops*. IEEE, 188–193. DOI: http://dx.doi.org/10.1109/ICDCSW.2014.20

[15] Virginia Braun and Victoria Clarke. *Using thematic analysis in psychology*. Technical Report.

[16] Ernie Brickell and Jiangtao Li. 2007. Enhanced privacy ID: A direct anonymous attestation scheme with enhanced revocation capabilities. In *Proceedings of the 2007 ACM workshop on Privacy in electronic society - WPES '07*. ACM Press, New York, New York, USA, 21–30. DOI: http://dx.doi.org/10.1145/1314333.1314337

[17] Quanwei Cai, Jonathan Lutes, Jingqiang Lin, and Bo Luo. 2018. A-Tor: Accountable anonymity in tor. In *Lecture Notes of the Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering, LNICST*, Vol. 238. Springer, Cham, 838–851. DOI: http://dx.doi.org/10.1007/978-3-319-78813-5{_}46

[18] Enrico Cambiaso, Ivan Vaccari, Luca Patti, and Maurizio Aiello. 2019. Darknet security: A categorization of attacks to the tor network. In *CEUR Workshop Proceedings*, Vol. 2315. http://ceur-ws.org/Vol-2315/paper10.pdf

[19] David Chaum, Farid Javani, Anna Krasnova, Aniket Kate, Joeri de Ruiter, and Alan T Sherman. 2016. cMix : Anonymization by High-Performance Scalable Mixing. *Cryptology ePrint Archive* (2016). https://bib.mixnetworks.org/pdf/chaum2016cmix.pdfhttps://eprint.iacr.org/2016/008.pdf

[20] David L. Chaum and David L. 1981. Untraceable electronic mail, return addresses, and digital pseudonyms. *Commun. ACM* 24, 2 (2 1981), 84–90. DOI: http://dx.doi.org/10.1145/358549.358563

[21] Kuanchin Chen and Alan L. Rea. 2004. Protecting personal information online: A survey of user privacy concerns and control techniques. *Journal of Computer Information Systems* 44:4 (2004), 85–92. DOI: http://dx.doi.org/10.1080/08874417.2004.11647599

[22] Michael Chertoff and Tobby Simon. 2015. *The Impact of the Dark Web on Internet Governance and Cyber Security*. Technical Report. Centre for International Governance Innovation and the Royal Institute for International Affairs. https://www.cigionline.org/sites/default/files/gcig_paper_no6.pdf

[23] Andrea Chester and Gillian Gwynne. 2006. Online Teaching: Encouraging Collaboration through Anonymity. *Journal of Computer-Mediated Communication* 4, 2 (6 2006), 0–0. DOI: http://dx.doi.org/10.1111/j.1083-6101.1998.tb00096.x

[24] Christophe. 2016. Tor – the good, the bad, and the ugly. (11 2016). https://blog.sqreen.com/tor-the-good-the-bad-and-the-ugly/

[25] Kimberly M. Christopherson. 2007. The positive and negative implications of anonymity in Internet social interactions: "On the Internet, Nobody Knows You're a Dog". *Computers in Human Behavior* 23, 6 (11 2007), 3038–3056. DOI: http://dx.doi.org/10.1016/j.chb.2006.09.001

[26] Joris Claessens, Claudia Díaz, Caroline Goemans, Bart Preneel, Joos Vandewalle, and Jos Dumortier. 2002. *Revocable anonymous access to the Internet*. Technical Report. http://www.esat.kuleuven.ac.be/cosic/http://www.law.kuleuven.ac.be/icri/

[27] Debbie Collins. 2003. Pretesting survey instruments: An overview of cognitive methods. *Quality of Life Research* 12, 3 (2003), 229–238. DOI: http://dx.doi.org/10.1023/A:1023254226592

[28] Gregory Conti and Edward Sobiesk. 2007. An honest man has nothing to fear. In *Proceedings of the 3rd symposium on Usable privacy and security - SOUPS '07*. ACM Press, New York, New York, USA, 112. DOI: http://dx.doi.org/10.1145/1280680.1280695

[29] George Danezis and Len Sassaman. 2008. How to bypass two anonymity revocation schemes. In *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, Vol. 5134 LNCS. 187–201. DOI: http://dx.doi.org/10.1007/978-3-540-70630-4{_}12

[30] Munmun De Choudhury and Sushovan De. 2014. Mental Health Discourse on reddit: Self-Disclosure, Social Support, and Anonymity. In *Proceedings of the Eighth International AAAI Conference on Weblogs and Social Media*. https://pdfs.semanticscholar.org/2db7/15a479c8961d3020fe906f7bedfa0311b937.pdf

[31] Claudia Diaz and Bart Preneel. 2007. Accountable Anonymous Communication. In *Security, Privacy, and Trust in Modern Data Management*. Springer Berlin Heidelberg, Berlin, Heidelberg, 239–253. DOI: http://dx.doi.org/10.1007/978-3-540-69861-6{_}16

[32] Jesus Diaz, David Arroyo, and Francisco B Rodriguez. 2017. Fair and Accountable Anonymity for the Tor Network. In *Proceedings of the 14th International Joint Conference on e-Business and Telecommunications (ICETE 2017)*. Volume 4: SECRYPT, 560–565. DOI: http://dx.doi.org/10.5220/0006474805600565

[33] Roger Dingledine and Nick Mathewson. 2006. *Design of a blocking-resistant anonymity system*. Technical Report. The Tor Project. 1–24 pages.

[34] Roger Dingledine, Nick Mathewson, and Paul Syverson. 2004. Tor: The Second-Generation Onion Router. In *13th USENIX Security Symposium*. US Dept of the Navy. DOI: http://dx.doi.org/10.21236/ADA465464

[35] Roger Dingledine, Nick Mathewson, and Paul Syverson. 2007. Deploying Low-Latency Anonymity: Design Challenges and Social Factors. *IEEE Security & Privacy Magazine* 5, 5 (9 2007), 83–87. DOI: http://dx.doi.org/10.1109/MSP.2007.108

[36] John E Dunn. 2013. Mevade botnet miscalculated effect on Tor network, says Damballa. (9 2013). https://www.cso.com.au/article/526724/

[37] Joan Feigenbaum, Aaron D Jaggard, Rebecca N Wright, and Hongda Xiao. 2012. *Systematizing "accountability" in computer science*. Technical Report. 1 pages. https://dedis.cs.yale.edu/dissent/papers/tr1452.pdf

[38] Andrea Forte, Nazanin Andalibi, and Rachel Greenstadt. 2017. Privacy, Anonymity, and Perceived Risk in Open Collaboration. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing - CSCW '17*. ACM Press, New York, New York, USA, 1800–1811. DOI: http://dx.doi.org/10.1145/2998181.2998273

[39] Kevin Gallagher, Sameer Patil, and Nasir Memon. 2017. *New Me: Understanding Expert and Non-Expert Perceptions and Usage of the Tor Anonymity Network*. https://ricochet.im

[40] Clement Guitton and Clement. 2013. A review of the available content on Tor hidden services: The case against further development. *Computers in Human Behavior* 29, 6 (11 2013), 2805–2815. DOI: http://dx.doi.org/10.1016/j.chb.2013.07.031

[41] Ryan Henry, Kevin Henry, and Ian Goldberg. 2010. Making a nymbler nymble using VERBS. In *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, Vol. 6205 LNCS. Springer, Berlin, Heidelberg, 111–129. DOI:http://dx.doi.org/10.1007/978-3-642-14527-8{_}7

[42] Rob Jansen, Florian Tschorsch, Aaron Johnson, and Björn Scheuermann. 2014. The Sniper Attack: Anonymously Deanonymizing and Disabling the Tor Network. In *NDSS*. DOI:http://dx.doi.org/10.14722/ndss.2014.23288

[43] Eric Jardine. 2015. *The Dark Web Dilemma: Tor, Anonymity and Online Policing*. Technical Report. Centre for International Governance Innovation and The Royal Institute for International Affairs. https://www.cigionline.org/sites/default/files/no.21.pdf

[44] Ruogu Kang, Stephanie Brown, and Sara Kiesler. 2013. Why do people seek anonymity on the internet?. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems - CHI '13*. ACM Press, New York, New York, USA, 2657. DOI: http://dx.doi.org/10.1145/2470654.2481368

[45] Sheharbano Khattak, David Fifield, Sadia Afroz, Mobin Javed, Srikanth Sundaresan, Vern Paxson, Steven J. Murdoch, and Damon McCoy. 2016. Do You See What I See? Differential Treatment of Anonymous Users. In *Proceedings 2016 Network and Distributed System Security Symposium*. Internet Society, Reston, VA. DOI: http://dx.doi.org/10.14722/ndss.2016.23342

[46] Sheril Kirshenbaum. 2011. Pseudonymity, Anonymity, And Accountability Online | WIRED. (7 2011). https://www.wired.com/2011/07/accountability-online/?fbclid=IwAR2JB3tD6y5gymnqpC7hoF72Brvf_cZzTuMm3d0WrexGjYhAtE4BPygLMr8

[47] Stefan Köpsell, Rolf Wendolsky, and Hannes Federrath. 2006. Revocable Anonymity. In *In: Müller G. (eds) Emerging Trends in Information and Communication Security. ETRICS 2006. Lecture Notes in Computer Science, vol 3995*. Springer, Berlin, Heidelberg, 206–220. DOI: http://dx.doi.org/10.1007/11766155{_}15

[48] Robert E. Kraut and Paul Resnick. 2011. *Building successful online communities : evidence-based social design*. MIT Press. 309 pages.

[49] Aditi Kumar and Eric Rosenbach. 2019. The Truth About the Dark Web. *IMF Finance and Development* 56, 3 (2019). https://www.imf.org/external/pubs/ft/fandd/2019/09/pdf/the-truth-about-the-dark-web-kumar.pdf

[50] Albert Kwon, Mashael Alsabah, David Lazar, Marc Dacier, and Srinivas Devadas. 2015. Circuit Fingerprinting Attacks : Passive Deanonymization of Tor Hidden Services. *USENIX Security* (2015), 287–302. https://www.usenix.org/system/files/conference/usenixsecurity15/sec15-paper-kwon.pdf

[51] Noam Lapidot-Lefler and Azy Barak. 2012. Effects of anonymity, invisibility, and lack of eye-contact on toxic online disinhibition. *Computers in Human Behavior* 28, 2 (3 2012), 434–443. DOI: http://dx.doi.org/10.1016/J.CHB.2011.10.014

[52] Zhuotao Liu, Yushan Liu, Philipp Winter, Prateek Mittal, and Yih Chun Hu. 2017. TorPolice: Towards enforcing service-defined access policies for anonymous communication in the Tor network. In *Proceedings - International Conference on Network Protocols, ICNP*, Vol. 2017-Octob. IEEE, 1–10. DOI: http://dx.doi.org/10.1109/ICNP.2017.8117564

[53] Lucian Constantin. 2012. Tor network used to command Skynet botnet. (12 2012). https://www.computerworld.com/article/2493980/tor-network-used-to-command-skynet-botnet.html

[54] Xiao Ma, Jeff Hancock, and Mor Naaman. 2016. Anonymity, Intimacy and Self-Disclosure in Social Media. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems - CHI '16*. ACM Press, New York, New York, USA, 3857–3869. DOI: http://dx.doi.org/10.1145/2858036.2858414

[55] Jonathan Mayer. 2017. Government Hacking. *Yale Law Journal* 127 (2017). https://digitalcommons.law.yale.edu/cgi/viewcontent.cgi?article=9272&context=ylj

[56] Nora McDonald, Benjamin Mako Hill, Rachel Greenstadt, and Andrea Forte. 2019. Privacy, Anonymity, and Perceived Risk in Open Collaboration. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems - CHI '19*. ACM Press, New York, New York, USA, 1–12. DOI:http://dx.doi.org/10.1145/3290605.3300901

[57] Daniel Moore and Thomas Rid. 2016. Cryptopolitik and the Darknet. *Survival* 58, 1 (1 2016), 7–38. DOI: http://dx.doi.org/10.1080/00396338.2016.1142085

[58] Milad Nasr, Alireza Bahramali, and Amir Houmansadr. 2018. DeepCorr: Strong flow correlation aacks on tor using deep learning. In *Proceedings of the ACM Conference on Computer and Communications Security*. 1962–1976. DOI: http://dx.doi.org/10.1145/3243734.3243824

[59] Lily H Newman. 2019. The CIA Sets Up Shop on Tor, the Anonymous Internet | WIRED. *WIRED* (5 2019), 7. https://www.wired.com/story/cia-sets-up-shop-on-tor/

[60] Andreas Pfitzmann and Michael Waidner. 1985. Networks Without User Observability — Design Options. In *Advances in Cryptology — EUROCRYPT' 85*. Springer Berlin Heidelberg, Berlin, Heidelberg, 245–253. DOI: http://dx.doi.org/10.1007/3-540-39805-8{_}29

[61] Amirali Sanatinia and Guevara Noubir. 2015. OnionBots: Subverting Privacy Infrastructure for Cyber Attacks. In *2015 45th Annual IEEE/IFIP International Conference on Dependable Systems and Networks*. IEEE, 69–80. DOI: http://dx.doi.org/10.1109/DSN.2015.40

[62] Adi Shamir. 1979. How to share a secret. *Commun. ACM* 22, 11 (11 1979), 612–613. DOI: http://dx.doi.org/10.1145/359168.359176

[63] Adario Strange. 2017. Why Apple plays China's censorship game. (8 2017). https://mashable.com/2017/08/02/apple-censorship-china/

[64] John Suler. 2004. The Online Disinhibition Effect. *Cyberpsychology & Behavior* 7, 3 (2004). https://pdfs.semanticscholar.org/c70a/ae3be9d370ca1520db5edb2b326e3c2f91b0.pdf

[65] Yixin Sun, Anne Edmundson, Laurent Vanbever, Oscar Li, Jennifer Rexford, Mung Chiang, and Prateek Mittal. 2015. RAPTOR: Routing Attacks on Privacy in Tor. In *USENIX Security*. Washington D.C., 271–286. https://www.usenix.org/node/190965

[66] Swati Khandelwal. 2014. Tor-enabled Point-of-Sale malware 'ChewBacca' stole Credit Card data from 11 Countries. (1 2014). https://thehackernews.com/2014/01/tor-enabled-point-of-sale-malware.html

[67] The SecDev Foundation. 2013. *Syrian Regime Tightens Access to Secure Online Communications*. Technical Report.

[68] Chau Tran, Kaylea Champion, Andrea Forte, Benjamin Mako Hill, and Rachel Greenstadt. 2019. Tor Users Contributing to Wikipedia: Just Like Everybody Else?. In *Proceedings on Privacy Enhancing Technologies*. http://arxiv.org/abs/1904.04324

[69] Patrick P. Tsang, Man Ho Au, Apu Kapadia, and Sean W. Smith. 2007. Blacklistable anonymous credentials: Blocking Misbehaving Users without TTPs. In *Proceedings of the 14th ACM conference on Computer and communications security - CCS '07*. ACM Press, New York, New York, USA, 72–81. DOI: http://dx.doi.org/10.1145/1315245.1315256

[70] Patrick P. Tsang, Apu Kapadia, Cory Cornelius, and Sean W. Smith. 2011. Nymble: Blocking misbehaving users in anonymizing networks. *IEEE Transactions on Dependable and Secure Computing* 8, 2 (3 2011), 256–269. DOI:http://dx.doi.org/10.1109/TDSC.2009.38

[71] Zeynep Tufekci. 2008. Can You See Me Now? Audience and Disclosure Regulation in Online Social Network Sites. *Bulletin of Science, Technology & Society* 28, 1 (2 2008), 20–36. DOI:http://dx.doi.org/10.1177/0270467607311484

[72] Eric C. Turner and Subhasish Dasgupta. 2003. Privacy on the Web: an Examination of User Concerns, Technology, and Implications for Business Organizations and Individuals. *Information Systems Management* 20, 1 (1 2003), 8–18. DOI:http://dx.doi.org/10.1201/1078/43203.20.1.20031201/40079.2

[73] Luis Von Ahn, Andrew Bortz, Nicholas J Hopper, and Kevin O'Neill. 2006. Selectively traceable anonymity. In *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, Vol. 4258 LNCS. 208–222. DOI: http://dx.doi.org/10.1007/11957454{_}12

[74] Keith D Watson. 2012. The Tor Network: A Global Inquiry into the Legal Status of Anonymity Networks. *Washington University Global Studies Law Review* 11, 3 (2012). https://openscholarship.wustl.edu/law_globalstudies/vol11/iss3/6

[75] Gabriel Weimann. 2018. *Going Darker? The Challenge of Dark Net Terrorism*. Technical Report. Woodrow Wilson Center. https://www.wilsoncenter.org/sites/default/files/going_darker_challenge_of_dark_net_terrorism.pdf

[76] Philipp Winter, Anne Edmundson, Laura M. Roberts, Agnieszka Dutkowska-Żuk, Marshini Chetty, and Nick Feamster. 2018. How Do Tor Users Interact With Onion Services? 411–428. https://www.usenix.org/node/217467

[77] Gang Xu, Leonardo Aguilera, and Yong Guan. 2012. Accountable anonymity: A proxy re-encryption based anonymous communication system. In *Proceedings of the International Conference on Parallel and Distributed Systems - ICPADS*. IEEE, 109–116. DOI: http://dx.doi.org/10.1109/ICPADS.2012.25

[78] Peter Zavlaris. 2016. Cloudflare vs Tor: Is IP Blocking Causing More Harm than Good? (4 2016). https://resources.distilnetworks.com/all-blog-posts/cloudflare-vs-tor-is-ip-blocking-causing-more-harm-than-good

[79] Ahmed T. Zulkarnine, Richard Frank, Bryan Monk, Julianna Mitchell, and Garth Davies. 2016. Surfacing collaborated networks in dark web to find illicit and criminal content. In *2016 IEEE Conference on Intelligence and Security Informatics (ISI)*. IEEE, 109–114. DOI: http://dx.doi.org/10.1109/ISI.2016.7745452