

Privacy Policies on the Fediverse: A Case Study of Mastodon Instances

Emma Tosch
Northeastern University
e.tosch@northeastern.edu

Cynthia Li
Independent Researcher
licynthia.x@gmail.com

Luis Garcia
Northeastern University
garcia.lui@northeastern.edu

Chris Martens
Northeastern University
c.martens@northeastern.edu

ABSTRACT

Free and open source social platform software has dramatically lowered the barrier to entry for anyone to set up and administer their own social network. This new population of social network administrators thus assume data management responsibilities for socio-technical systems. Administrators have the power to customize this software, including data collection and data retention, potentially leading to radically different privacy policies.

To better understand the characteristics — e.g., the variability, prohibitions, and permissions — of privacy policies on these new social networking platforms, we have conducted a case study of Mastodon. We performed a text analysis of 351 privacy policies and a survey of 104 Mastodon administrators. While most administrators used the default policy that ships with the Mastodon software, we observed that approximately ten percent of our sample tailored their privacy policies to their instances and that some administrators conflated codes of conduct with privacy policies. Our findings suggest the existing market-based individualistic frameworks for thinking about privacy policies do not adequately address this emerging community.

1 INTRODUCTION

In May 2023, the administrators of █████¹—a social networking website dedicated to political discourse with at least 5,000 active members—reported that the FBI gained access to unencrypted user data, sending shock waves through certain corners of the Internet [14]. While some members of the network denounced the administrators’ data management practices, █████ used the default privacy policy that ships with the software they run, which does not address database encryption. Although administrators may not have adhered to best practices, their actions did not violate their website’s privacy policy.

Ordinarily one would expect few consequences for those who operate █████ from users: material harm and fault from data breaches has been hard to establish in some jurisdictions [13, 87, 101] and the average Internet user has become inured to data breaches at

¹We occlude all instance names throughout so as to not draw additional attention to individual servers.

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

Proceedings on Privacy Enhancing Technologies 2024(4), 700–733

© 2024 Copyright held by the owner/author(s).

<https://doi.org/10.56553/popets-2024-0138>



companies, often attributing fault to individual choices while failing to take protective measures or choose more secure alternatives [1, 18, 45, 63, 64]. However, the data breach at █████ did not occur at a for-profit company — it occurred on an independent and informal *federated social networking* website where administrators and users have a different kind of relationship.

In a *federated social network*, a collection of independently operated, maintained, and governed entities known as *instances* run social networking software that communicate with each other over a shared protocol. Federated social networking has coalesced around using the ActivityPub protocol as the *de facto* communication standard [100]. The collection of websites and services that implement ActivityPub is commonly referred to as “the Fediverse.” Each server is maintained by independent administrator teams and the members of each instance form a sub-graph in the overall Fediverse network [5]. Instances function as autonomous communities whose rules and UX may differ dramatically from other instances running the same software [105].

Due to how federation works, privacy-threatening events such as a data breach at one instance can risk exposing personal information from another [33, 89]. Existing norms and data management practices within traditional social networking platforms were developed in the context of an oligopolistic ecosystem wherein privacy policies primarily benefit well-resourced entities such as corporations, whose primary motivation for having such documents is compliance with the law [8, 19, 72, 73, 85, 90]. Conversely, Fediverse instances may be run by hobbyists, activists, and other individual stakeholders who have personal, non-financial interests in their instances. Depending on the size and scope of an instance’s community, a new administrator may be assuming responsibilities that are ordinarily performed or overseen by a team of lawyers, information security professionals, trust and safety officers, or data engineers [16, 26]. Despite these differences, instances are governed by the same sets of laws as corporate platforms, potentially exposing administrators to unforeseen complex legal issues, e.g. Fediverse instances must have privacy policies and administrators are data controllers under GDPR [3, 29, 81].

Prior work has examined privacy policy compliance at companies [42, 96] and the (in)effectiveness of privacy policies at empowering users to make informed choices [37, 44, 83]. Scaffolding this prior work is the assumption that there exists an underlying adhesive agreement (e.g., terms of service) that may admit uses or transmission of user data in ways an end-user may not anticipate [48, 51]. Relationships between end-users and administrators

in the Fediverse are more balanced than relationships between end-users and companies. Since laws mandating privacy policies were written for a different population having different resources and motivations, it is currently unknown whether prior research findings transfer to Fediverse administrators, or how Fediverse administrators understand, select, edit, or write privacy policies for the servers they operate.

To address this research gap, we present the first study of privacy practices on the Fediverse, using Mastodon — a popular microblogging platform — as a case study. We collected a sample of Mastodon privacy policy documents for textual analysis and deployed a survey to Mastodon administrators, asking about their past experience with data management and content moderation, and about the choices they made when selecting, writing, and modifying privacy policies. Our work provides preliminary answers to the following descriptive research questions:

- RQ1 (Privacy policy variability)** How much and in what ways do different Mastodon instances' privacy policies differ from each other and do those differences matter?
- RQ2 (Privacy policy provenance)** How do administrators select or write their privacy policies and what factors cause them to modify their privacy policies?
- RQ3 (Administrator characteristics)** What are the characteristics of English-speaking Mastodon administrators in terms of background, experience, and demographics?

Our findings suggest substantial heterogeneity in how administrators view user privacy and the obligations outlined in privacy policies. Some administrators struggle to differentiate between codes of conduct, terms of service, server rules, and privacy policies. A surprising number of administrators make factually incorrect statements in their privacy policies or are unaware that their instances even have privacy policies. Those who are aware of their policies range from cautious actors who seek legal counsel before changing their privacy documents, to more *laissez-faire* administrators who view such documents as largely useless legalese. We conclude with a discussion of observed community assets and recommendations for supporting these communities' unique socio-technical needs.

2 MASTODON AND THE FEDIVERSE: BACKGROUND AND RELATED WORK

Most software in the Fediverse today exchanges data according to the ActivityPub protocol [95, 100]. This shared protocol operates similarly to how email connects users with addresses on many different domains via POP3/SMTP. ActivityPub supports a variety of user experiences and interaction models and has enabled the creation of alternatives to diverse social networks such as Mastodon, Lemmy, and Funkwhale (microblogging, link aggregation, music sharing).

End-users who wish to access the Fediverse can do so in one of two ways: join an existing server or operate their own. When end-users choose to join an existing server, they delegate operational responsibilities to instance administrators, including data management. When they choose to operate their own server, they become administrators and assume responsibility for the data management of any other users whom they allow to join their instance.

Example. Suppose Alix, a member of `mu.blog` (a microblogging instance), and Bobbi, a member of `link.agg` (a link aggregation instance), follow each other. Bobbi posts a link and to them it appears under a particular topic on `link.agg`. When Alix sees this post, to them it is interleaved in their social feed as if it were posted directly on `mu.blog`. When Alix replies to Bobbi's post, it just appears as a tagged `mu.blog` post, but to Bobbi it looks like a threaded response on `link.agg`. Although their instances run different software, they can still interact with each other. Thus, while users *can* create multiple accounts across multiple instances, the main advantage of the Fediverse is that they don't have to.

Now suppose Cori wants to join a microblogging instance in the Fediverse but doesn't like the server rules for `mu.blog`. They decide they'd rather operate their own instance. They register the domain `cori.fedi` and decide to run the Mastodon microblogging software. When Cori's friends and family join their server, they can still interact with other Fediverse instances, but they are beholden to Cori's rules, design choices, etc. Thus, when users join the Fediverse, not only must they choose not only the UX (e.g., microblogging vs. link aggregation), but also the specific community via instance/server.

Research Gap: Administrators vs. Users. Prior studies of the Fediverse — and Mastodon in particular — have largely focused on user behavior, connectivity, and server rules [41, 56, 67, 82, 104]. Our work focuses on the privacy and data management perspectives of administrators (e.g., Cori), rather than users/members (e.g., Alix, Bobbi).

2.1 Related Software and Communities

Operationally, Mastodon administration is a cross between running an email server and managing a bulletin board system (BBS): like email, each user has an individual inbox that may contain messages from other servers/instances; like a BBS, these messages are public.²

While Mastodon has been presented as a Twitter alternative, its governance structures more closely resemble Reddit, due to both consisting of small communities with their own rules and processes. Administrators often make both social and technical policy decisions, resulting in a social dynamic akin to open source software development layered on top of a traditional social network [105]. Each community on Mastodon has different norms around research, data collection, and expectations of privacy or ephemerality [9, 60]; see Appendix A for comparisons with legacy platforms.

Federated online social networks predate the family of software we now consider part of the Fediverse, many of which were introduced in or after 2017 alongside the ActivityPub protocol [95]. Diaspora is one such example that failed to thrive in its early years. Bielenberg et al. [10] found Diaspora's network wanting in terms of server reliability and end-user data security. While server reliability remains an issue for some Mastodon instances [78], overall Mastodon appears to be more reliable than Diaspora was. We speculate that this may be due in part to improved affordances for administrators: the rise of containerization and improved usability of cloud hosting has lowered the barrier to entry for the average Internet user to set up and administer fault-tolerant social software.

²Different federated social networking software provides different support for message visibility, but the default assumption should be that all messages are public knowledge.

2.2 Hosting Choices Impact Policies

The first (and possibly only) technical choice new administrators make is where and how to host their instance: locally, via cloud providers (i.e., platforms as a service (PaaS)), or via managed platforms (i.e., software as a service (SaaS)). One administrator may choose to eschew mega-corporate cloud providers and host locally, exercising complete control over the software stack. Another administrator may choose a managed hosting platform, outsourcing their software maintenance and data management. The technical choices of where to host software has legal implications that an administrator may not anticipate.

For example, there are several popular managed (i.e., software as a service (SaaS)) Mastodon hosting services, e.g., `masto.host` and `toot.io`. SaaS platforms like `masto.host` may make many choices for the administrator. They may implement barriers to administrators changing Mastodon documents or code. While this may initially seem better for users — for example, they could bar administrators from downloading unencrypted user data or perform audits on software updates vs. privacy policies — it does put both users and administrators at the mercy of the hosting platform: while at present all of the managed hosting we've identified is supported by subscription, it is possible alternative funding models will arise in the future (e.g., advertising or data-brokerage-based services). Finally, SaaS platforms have their own terms of service and privacy policies that may be relevant to instance members.

2.3 Challenges Studying Social Networks

Social scientists and network scientists have been studying online social behavior since the advent of BBSs [39, 76, 77, 103] and the Fediverse is no exception [4, 31, 57]. Prior to the rise of centralized corporate social media research, where there exists a canonical list of users, researchers had to develop alternative techniques in order to produce a true random sample from social networks. These techniques often relied on simulation to show that a sampling strategy was unbiased, given assumptions about network topology [2, 52, 53, 55]. It is not yet clear whether these assumptions hold in the Fediverse, due to documented volatility in the set of active instances [78]. Researchers seeking to collect representative samples should exercise caution when generalizing and collect new data periodically, by using, e.g., La Cava et al. [56]'s "polite data crawling" methodology.

2.4 Challenges Collecting Privacy Policies

Typically we can extract the privacy policy of an active Mastodon instance using only our knowledge of its domain (i.e., without any crawling or scraping). This is because the default Mastodon software stack ships with a default privacy policy template. This policy will be dynamically populated with the instance name and appear at `https://<DOMAIN>/privacy-policy` on servers that use the Mastodon front-end.

Of course, if administrators edit this template to link to a policy hosted elsewhere, redirect the default URL, password-protect the website, or simply decide to use a different front-end, this will frustrate attempts to locate the instance's privacy policy. Fortunately the design friction of deviating from the defaults, combined with

Mastodon's editable markdown template mean that locating privacy policies for this specific software is an easier task than in the general case [65, 68, 88, 102].

2.5 Folkloric Assumptions

Due to Mastodon's distributed nature, it is difficult to obtain a complete picture of demographics, social norms, and practices. However, Fediverse participants discuss their experiences, and individuals sometimes write longer-form documentation of their observations. Zulli et al. [105] report that content moderation discussions on Mastodon often appear with the `#federation` tag and content moderation is deeply tied to social norms and marginalized identity [32, 38]. Based on both our subjective observations and published testimonials, we approached this research with the following folkloric assumptions, which influence our study design and interpretation of collected data. We also document these assumptions in order to highlight results that explicitly challenge or refute them.

2.5.1 Most administrators have a stronger than average technical computing background. Mastodon began as a "counterculture" response to Twitter before there was significant social momentum away from the latter.³ The concept of federation itself represents both a political experiment and a technical innovation (e.g. with the ActivityPub protocol). Thus we assume that most early adopters (and by extension most of the transitive closure of their social connections) would have more than a lay interest in computing, internet architecture, and sociotechnical systems. Mastodon communities tend to attract users who are interested in free and open source software, cautious about the collection and surveillance of their data, and knowledgeable of the technical and ethical aspects of data scraping and privacy [60].

2.5.2 Over-representation of administrators identifying as LGBTQ+ and under-representation of those identifying as a race other than white. Based on the Fediverse's origins, the Mastodon Covenant, and subjective observations about the population of active users, we assume Mastodon's administrators may disproportionately represent LGBTQ+ (and especially transgender) people who are well aware of their vulnerability to harm and harassment online. We also assume white administrators are over-represented, based on colloquial users' reports of community demographics [28, 97].

2.5.3 Administrators view their roles primarily as community moderators rather than data managers. Since the Fediverse emerged in response to the experiences of high-profile marginalized community members' inability to enforce boundaries and consent around their social network interactions, we expect that administrators will for the most part be motivated by adopting the role of a community steward. By extension, they may put a great deal of care into crafting their content moderation policy, but not necessarily put the same thought towards their privacy policy and data management approach. For example, Mastodon's instructions for setting up your own server note that administrators should write out a code of conduct, but do not mention privacy policies at all [62].

³As evidence of its counterculture roots, early Mastodon was designed to be GNU Social compatible and has emphatically never been pitched as a start-up (<https://news.ycombinator.com/item?id=12646083>); GNU Social was itself founded by Free Software Foundation employees.

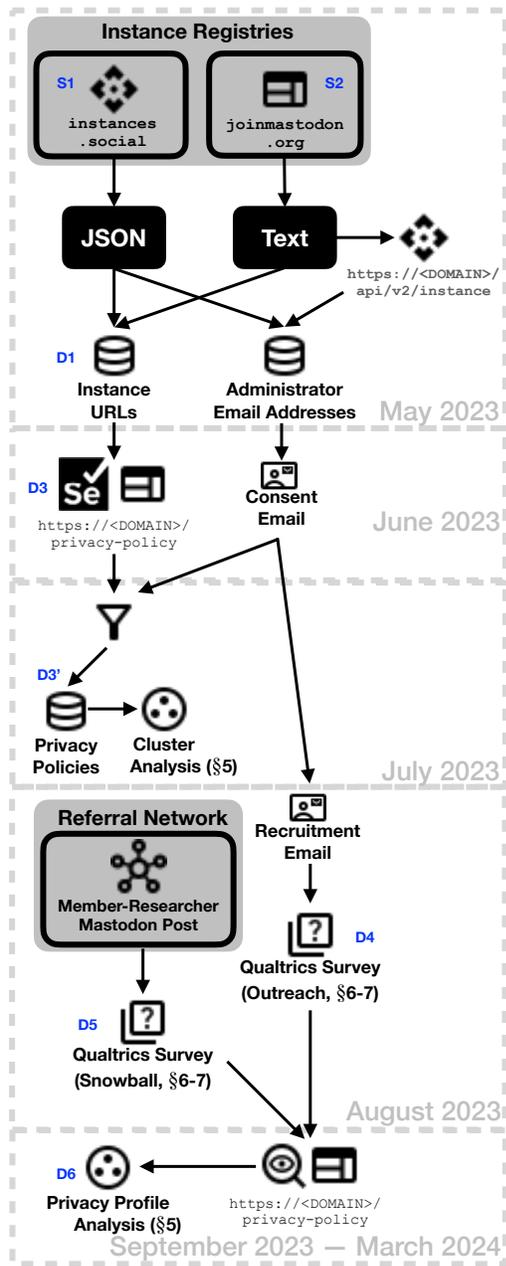


Figure 1: Timeline of the data collection process. Source data are labeled with the prefix **S**. Other data and their subsets are labeled with the prefix **D** and the suffix **'**, respectively.

3 STUDY DESIGN

To study the variability of privacy policies on Mastodon, we collected privacy policy documents for text analysis. To study administrator choices, experiences, and perspectives on privacy policies, we conducted an online survey of Mastodon administrators.⁴ Figure 1 depicts our data collection and processing workflow and timeline.

⁴Reviewed as exempt by the Northeastern University IRB.

3.1 Sampling Instances (S1, S2, D1)

All of our data collection tasks began with a sample of Mastodon instances. As discussed in Section 2.3, the most reliable way to sample network data is through crawling. However, we eschew crawling in our work due to ethical considerations, discussed in depth in Section 4. Therefore, we instead selected our initial sample of instances using long-running centralized registries: `joinmastodon.org` (S1) and `instances.social` (S2). `instances.social` provides an API for developers and researchers that returns structured data, including administrator contact information, server language, and whether the instance was live at last query. `joinmastodon.org` only provides a text list of instance domains; we must query the Mastodon API to obtain this information. We then applied the following exclusion criteria to identify 351 distinct instances (D1):

- *English not listed as a server language.* We controlled for variability in our text analysis by restricting our initial data set to English-language servers. If the API queries did not return any official languages, we excluded the server.
- *Inactive, nonresponsive or protected servers.* We excluded all servers that `instances.social` marked as inactive and all servers that returned HTTP errors at either the instance API endpoint or at the default privacy policy URL.
- *Short privacy policy.* Policies of fewer than ten lines of text after normalizing via `pandoc` [59] were excluded.

Limitations. We have no reason to believe these registries provide representative samples of Mastodon servers; we discuss norms surrounding using this data and further enumerate how they might not be representative in Appendix B. Were the registries themselves representative, our exclusion of non-responsive servers may be biased; we observed 502 and 503 errors, which is line with some of the up-time issues reported in previous work [78]. We also observed 403 errors, suggesting that we are excluding privacy policies from security-minded administrators. Finally, excluding short policies and not following links means we may exclude policies from larger or older organizations that use the default privacy policy URL as a pointer.

3.2 Policy Collection: Registries (D3')

In June 2023, we downloaded the 351 policies from the confirmed endpoint using Selenium in headless mode to access any dynamically generated text. We converted the resulting HTML to Markdown using `pandoc` [59] and manually inspected a sample of policies. After observing prohibitions against crawling and scraping on some servers (see Section 4), we emailed administrators via the contact information we'd obtained using the Mastodon API, asking permission to include their publicly available policies in our analysis. We sent one reminder email.

In July 2023 we filtered out the policies from the servers that opted out (D3'). We removed instance names, normalized white space, and converted all characters to lower case before running our text analysis RQ1, findings in Section 5.1.

Limitations. We only used the default privacy policy URL to locate policies. We did not look for policy information in other locations. Thus we believe our analyses underestimate the true variability of privacy policies and privacy-related policies.

3.3 Survey (D4, D5)

In August 2023, we launched two versions⁵ of an online survey hosted on Qualtrics to understand how administrators wrote, selected, and/or maintained their servers' privacy policies (Appendix C). We first asked respondents to identify the instance they administered before asking questions about their familiarity with their instance's privacy policy, the history of their instance's privacy policy, and their own backgrounds in an effort to answer RQ2 and RQ3.

We used Dillman et al. [24]'s tailored design approach. To maximize response quality and completion rates, we allowed respondents to skip questions, judiciously employed open responses, and implemented significant conditional branching in order to maximize the relevance of a question to the respondent. We recorded partial responses. The median response time for respondents who completed survey was 10.3 minutes for each group. Respondents were not paid.

3.3.1 Testing, piloting, deploying. Prior to deploying the survey, several members of the immediate research community tested the survey for comprehensibility and flow. Two Fediverse community members known to the authors then took the survey under observation of and in dialogue with one of the authors. After finalizing edits, the authors then rolled out the survey in stages to progressively larger audiences in August 2023 (D4).

- (1) Email individual survey links to early affirmative respondents to our request to use instance privacy policies (registries, §3.2).
- (2) Email individual survey links to later affirmative respondents to our request to use instance privacy policies (registries, §3.2).
- (3) Email individual survey links to all remaining affirmative respondents and non-respondents to our request to use instance privacy policies (registries, §3.2).
- (4) Email individual survey links to negative respondents to our request to use instance privacy policies (registries, §3.2).
- (5) Post anonymous survey link to referral network (§3.3.2).

3.3.2 Referral Network Recruitment (Snowball Sampling). "Snowball sampling" is a commonly used non-probability sampling method for networked data when there do not exist means to perform probability sampling [34, 69]. One member of our research team is also a member of the population of interest (i.e., a *member-researcher*). They both administer an instance that runs a Mastodon fork and participate as a member of a larger instance that affords them greater network connectivity. This researcher served as a "seed," posting an anonymous link to our survey and asking others to share. They fielded questions from prospective participants and made one reminder/follow up post. Members of this recruitment pool had access to a single public survey link (D5).

Limitations. Because the snowball sample is opt-in and potentially anonymous, we could not apply the exclusion criteria of Section 3.1 *a priori*. One implication of this difference is that administrators in each group may differ in systematic ways. For example, the snowball sampled group may have more technically-savvy administrators who have implemented authentication measures or

who have customized the navigation of their instances in ways that would lead us to exclude them in the outreach group. With this in mind, we report our survey results by recruitment group and test for differences. All quantitative analyses should be interpreted as exploratory and descriptive, not confirmatory and associational or causal.

3.4 Policy Inspection: Referral Networks (D6)

We manually inspected, manually assigned annotations, and manually extracted relevant quotations from the publicly available privacy policies of all survey respondents who provided valid URLs. We did not download these policies, nor use any automation to access or inspect these policies. This data in this analysis overlaps with D3', but was collected in a distinct process.

4 ETHICS

Answering our research questions required two pieces of information from instances: an administrator contact and the text of the instance's privacy policy. Neither datum is typically considered protected data when available publicly on the web. All of the privacy policies we collected were publicly accessible and all administrator contact information was obtained using public APIs; using this type of data typically does not require IRB approval. We submitted our proposed study for IRB review and were determined to be exempt.

Despite our exempt determination, we proceeded with additional care due to previous conflicts between researchers and similar communities to our community of study. In the open source community, a 2021 IRB-exempted study that involved submitting faulty patches to the Linux kernel led to the banning of University of Minnesota affiliates from future contributions to the Linux kernel [12]. Regarding Mastodon specifically, a 2018 study from the University of Milan scraped data from the Fediverse, breaking terms of service of many servers and misrepresenting the data that was scraped; this publication was met with an open letter from the Mastodon community [84].

With this context in mind, we asked minimally invasive questions in our survey to respect the privacy and social interiority of these communities. As our research objectives and interests require the trust of the Mastodon administrator community, we follow Proferes et al. [75]'s best practices for online data collection of public data that ordinarily would be considered exempted by IRBs. They advocate a "contextual" approach that seeks to understand how different online communities view the "publicness" of their public data. Based on our understanding of the community context, an ethical investigation of Fediverse practices necessarily involves the consent from administrators' individual instances, even for publicly-available data, adhering to a standard beyond an IRB-approved protocol.

Scraping. During our preliminary analysis of privacy policies, we noted that at least two instances prohibited scraping in their privacy policy:

We love scientific research, but data must be gathered with consent. You may not scrape data from █████ for research without express consent of each user whose data you gather.

⁵The two versions of the survey were nearly identical to participants, differing only in wording to reflect the recruitment strategy for each group.

Researchers who wish to study █████ or our users by collecting data using the API or through any other means that does not involve an "opt-in" from individual users are required to submit their protocol. This does apply to database scraping software, or any means of recording user activity where our users might be surprised that they were included afterward, because they were not given a chance to consent.

We did not collect any non-administrator user data for this project. The only administrator user data we collected was the listed contact email address via the instance API endpoint.

Although our IRB would not consider privacy policies to be user data, we deemed it in alignment with these ethical considerations to involve administrators in this process. Of the 84 administrators who responded to our request for permission to use their privacy policies, 8 opted out (i.e., denied permission for their policies to be used in our analyses). Their privacy policies are not included in our textual analysis below, nor do we include their servers in our analyses of instance characteristics. However, our manual inspection reveals that all 8 policies use the default privacy policies with either no modifications (5 instances) or minor markdown modifications (3 instances, no semantic differences); the default policy is the default privacy policy of Mastodon and is covered by AGPLv3.

Given that one of our findings was administrator confusion over governance documents, it is possible that those who requested their privacy policies not be used were actually thinking of their server rules, codes of conduct, or some other public-facing document(s). This observation is important in the context of social media companies effectively closing off academic researchers' access to their data. Researchers will move along with users to alternative social media while expecting data collection norms to transfer as well. This is problematic precisely because the Fediverse, in its origins, is a rejection of the corporate social media ecosystem.

Crawling. Mastodon provides an API endpoint that lists all of the servers accessible via one hop from the instance's user network — i.e., the set of servers which host all of the accounts that the users on the current server follow. This information is *derived* from user data. Some administrators identified via centralized registries responded to our emails to ask how we came to identify them; we provided them with this information when asked. This experience led us to eschew crawling the network ourselves. We did not believe we would obtain sufficiently useful additional information beyond that offered by the centralized registries. Furthermore, due to the aforementioned community trust issues, we felt the risk of damaging future relationships to administrators was not worth the potential reward of additional instance data.

Participant Privacy and Confidentiality Considerations vs. Data Integrity. As discussed in Section 2.5, the Fediverse has a reputation for providing communal space for vulnerable communities, e.g., transgender populations. Because administrators are often also members of their instances' communities (i.e., not just operating a service), as points of contact for the instance they are particularly vulnerable. We thus assume that administrators place high value on privacy and confidentiality. Thus, we sought to collect as little personal or identifying data as possible.

We do, however, need to collect *some* data in order to validate our results, since we cannot assume that all survey respondents will act in good faith and [22, 30]. In particular, we mitigated the threat of spammers in our design through (1) limiting recruitment channels to direct outreach and posting in the Fediverse (e.g., we did not post on major corporate social media platforms), (2) not paying respondents to participate in this research, (3) requiring respondents to provide the name of the instance they administer.

5 RESULTS AND FINDINGS

To answer **RQ1**, we analyzed the text of privacy policies (**D3'**) on 351 identified using the centralized registries (§3.1). To answer **RQ2** and **RQ3**, we analyzed the results from our two survey groups:

Outreach group (D4). From instances identified via centralized registries, the research team recruited 348 distinct administrators to respond to the survey by email. From this group, 55 administrators opened the survey, 45 began the survey, and 40 reached and responded to the server name question. We considered all 40 of these responses high quality and did not exclude any from our analyses.

Snowball group (D5). One-hundred individuals attempted to respond to the survey that was distributed via referral network. Seventy provided text for the server name; of these, 64 were confirmed to be running Mastodon or a closely-related Mastodon fork. While we are generally interested in platforms that participate in the Fediverse, in our analysis below we exclude instances confirmed to operate other Fediverse platforms (e.g., Friendica) in an attempt to treat the platform affordances as a controlled variable.

5.1 Policy Variability (D3', D6)

To better understand the landscape of privacy policies in Mastodon, we conducted a preliminary analysis of the variability among servers' policy documents (**RQ1**). Mastodon provides a default template for a privacy policy, but since the software is open source, administrators are free to use or create one of their own. Our analysis revealed that this affordance yields privacy policies that can differ from each other not only in obvious ways, but also in subtler ways that users can miss on a casual reading.

5.1.1 Cluster analysis (D3'). We performed three rounds of cluster analysis over the initial privacy policy dataset.

Round1 We used the MeanShift algorithm provided by scikit-learn over bags of words from the policies [70].

Round2 From each cluster produced in **Round1**, one researcher selected exemplar documents. From these exemplar documents, that same researcher produced keywords by first using nltk [11] to extract nouns from the pre-processed and normalized versions of these documents, and then manually selecting out terms that were most relevant to privacy and security (e.g., keeping "password" and "admin," but dropping "video"). That same researcher then used the MeanShift algorithm again to produce a new set of clusters over keywords extracted from all of the documents.

Round3 From each cluster produced in **Round2**, the same researcher observed that some policies within each cluster

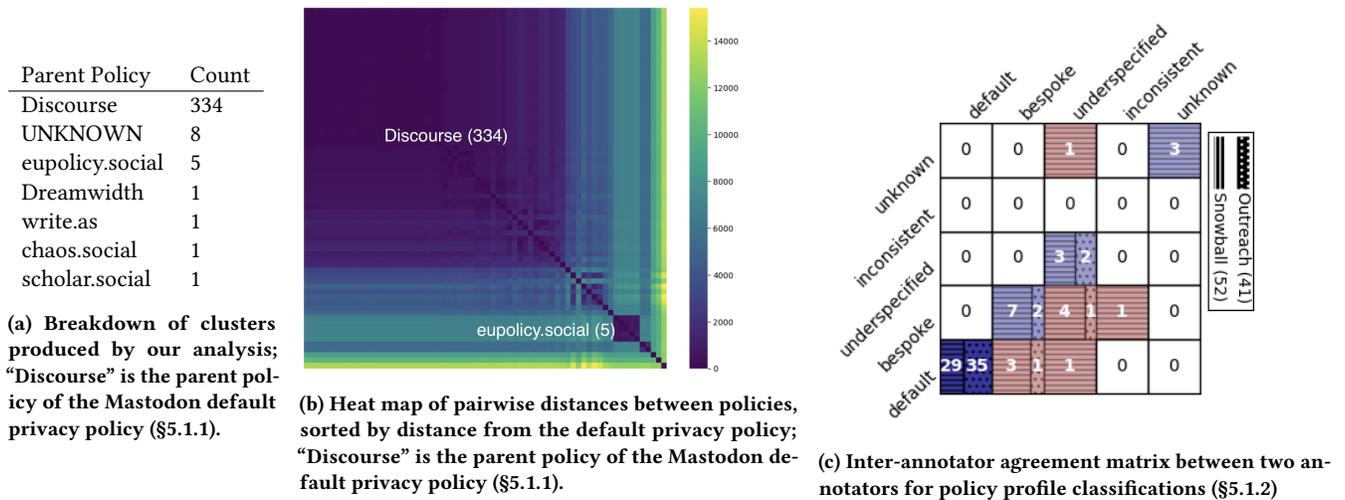


Figure 2: Policy variability analysis (RQ1). Tab. 2a depicts the counts per policy cluster. Policies that did not cite any template or inspiration for their text are UNKNOWN. The policies from the eight instances that opted out all use the default with no modifications, so we left them in the dataset. Fig. 2b depicts a heat map of pairwise Levenshtein distances between policies, with the policies (represented by the rows and columns) sorted by their Levenshtein distance from the default Mastodon policy. Origin is the top left. The two parent policies that map to these regions are labeled. We identified an additional four clusters that collectively contain 12 documents from our sample. Fig. 2c depicts an inter-annotator agreement matrix between two annotators (Cohen’s $\kappa \approx 0.67$) for policy profile classifications of survey respondents’ privacy policy documents. Eighty-four respondents’ instances had publicly identifiable privacy policies. Nine instances included privacy information on their about page or on another web page linked from the privacy policy. Each matrix entry is broken down by the recruitment method. The data in this figure overlaps with but is not the same as the data in Tab. 2a and Fig. 2b; these data include instances from the snowball sample.

explicitly referred to another policy. That researcher manually inspected example documents (not necessarily the exemplars from Round2) from each cluster and noted which policy was cited as being an inspiration or template for the documents.

Findings. Table 2a depicts the results of this analysis. We call the cited policy a *parent policy*, using them to label each cluster. To ensure that two documents of the same cluster were not too different in content from each other, we performed a “sanity check” by calculating Levenshtein distances on every pair of documents using the pre-processing steps of Section 3.2. Figure 2b depicts a heat map of these distances. Dark regions centered along the diagonal of the graph suggest that there are “families” of policies such that parent policies differ from the default policy of Mastodon. There were no cases of one cluster citing two different parent policies.

The researcher who conducted this work disseminated and discussed their findings with the other authors. Because a single researcher performed that part of the analysis, we were not able to produce measures for inter-rater reliability.

Significance. These different parent policies (or the decision to not use one) provide different starting points from which administrators write their own privacy policies, if they choose to modify them at all. These starting points may not necessarily be germane to their instances’ needs.

5.1.2 Theorized Policy Profiles (D3). Upon manual inspection of a sample of all privacy policies, we partitioned privacy policies into those that hew closely to the default Mastodon privacy policy and those that do not. Among those that do not, there are legally-enforceable policies that address the storage and transmission of personal data and then there are documents that would not meet any legal standard for notice and choice. Among those that do not meet the legal standard, there are documents that attempt to address privacy and there are those that make outright false or contradictory claims. We discuss each of these four profiles below and connect them to the cluster analysis when appropriate.

Default. Mastodon ships with a default privacy policy that was initially based on and cites a reference privacy policy.⁶ This default policy lists the specific personal identifying information (PII) the software collects, how long that information is kept, and how it is used. We would expect administrators who do not modify the default Mastodon setup to only add non-contradictory statements to the privacy policy, since they do not change the software’s basic operations.

There were 334 default-derived policies in our dataset, of which 50 (15%) differed from the default text (Levenshtein distance within {2, 388}). Of those 50 differing policies, 43 policies differed only in ways unrelated to privacy; examples include spelling localization

⁶Mastodon’s privacy policy shares a common ancestor with Discourse’s privacy policy (<https://www.discourse.org/privacy>).

(e.g., *favorite* vs. *favourite*), removing the acknowledgment of or link to the reference/parent policy (e.g., Discourse, see Table 2a) or modifying the date that the policy was last edited. The remaining seven included removals, additions, and replacements to the text that we argue represent semantic differences from the source policy. These changes include adding or removing the following text:

- ✦ a disclosure of which hosting platform the instance runs on,
- ✦ a disclosure of how collected data will be used for analysis,
- ✦ a disclosure of the server administrators’ jurisdictions,
- ✦ a claim that advertisements are banned,
- ✦ a claim that the instance does not use cookies,
- default information on cookie usage,
- a default section on how the server will *not* sell, trade, or transfer data to outside parties, and
- textual references to regulations (e.g., COPPA) that may be jurisdictionally irrelevant.

Bespoke. Seventeen policies differed substantially from **Default**. Common among these policies are more detailed explanations of what information is at stake and how that information is processed. Documents ranged in their formality, from earnest educational attempts at explaining to the end-user some of the underlying processes of the underlying server, to more typical legalese including technical definitions of terms used in the policy (e.g., “User”, “Account”).

These policies display diversity among themselves as well as the differences they share when compared with the default policy. We found policies that offer different attitudes toward privacy and privacy-related subjects. For example, one policy “reserves the right to display advertisements on [a user’s] content unless [the user has] purchased an Ad-free Upgrade or a Services account”. Another policy is written to be a “non-privacy policy”, asking users to “assume that everything [they] contribute to [the server] is public”. A third example offers users the ability to review any subpoenas that the subject instance receives that compels it to disclose personal information. These examples involve distinct categories of information and information processing actions, representing distinct concerns of administrators and their communities.

Inconsistent. After our initial analysis we manually re-examined some of the policies we’d excluded due to length. We found that some privacy policies included statements that are not consistent with how the Mastodon software actually operates. For example, the entirety of one privacy policy reads:

Your data is yours, █████ and █████ will never process or use your data.

Mastodon must store messages on the instance’s server in order to function properly, so this statement is factually incorrect. The following excerpt from a privacy policy claims that the instance does not use cookies:

Do we use cookies?
No.

We found that this claim was false by creating an account on the instance that made the claim and finding that two cookies were indeed saved onto our browser. One cookie, called `_session_id`,

stores our current login session on the server, and is necessary for our login status to persist between page reloads. The other cookie, called `_mastodon_session_id`, is set both for authenticated and unauthenticated users, and is used by Mastodon to track our browsing behaviors as a security precaution.

Under-specified. All privacy policies are, to some degree, under-specified. However, some are noticeably lacking in content. Others should not be considered privacy policies at all; for example, this is the entirety of one privacy policy:

Accounts dormant for more than 12 months are subject to suspension.

Another under-specified policy avoids the problems of inconsistent policies by gesturing toward whatever is necessary to operate an instance:

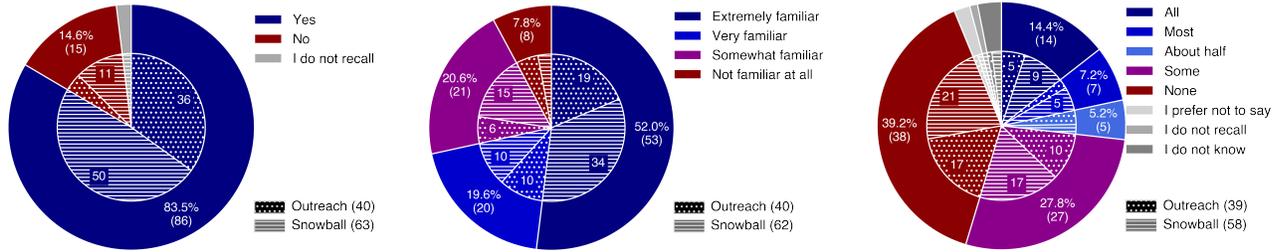
█████ will process the data you provide in order to provide you with a Mastodon instance.
We will make best efforts to protect your data, and we won’t transfer it, sell it or use it for anything except providing you with this instance.
If you want to exercise your rights under the GDPR ping us an email at: █████.

5.1.3 *Observed Policy Profiles (D6).* After the end of the survey (Nov. 2023), one researcher identified URLs for privacy policies of 84 instances represented by survey respondents. Nine of these instances contained privacy policy information in two locations. Two researchers annotated the set of 93 documents according the four profiles defined here (Cohen’s $\kappa \approx 0.67$), plus “unknown.” All “unknown” classifications were for policies written in languages the annotators could not understand. We found the greatest disagreements regarding how far a policy could deviate from the default before it became something else.

5.2 Policy Provenance and Incidents (D4, D5)

In our survey we asked respondents whether they were *personally involved* in the crafting of their instance’s privacy policy. This allowed us to get a sense of what proportion of respondents were speaking from personal experience vs. institutional knowledge. We then asked them to rate their familiarity with their privacy policies and to report when they last read through them (Q8–Q10) before proceeding to our specific questions regarding instances’ privacy policy origins (Q13, Q18, and Q11).

Most respondents across both groups reported being involved in establishing their instances’ privacy policies (85%), with no evidence of a statistically significant difference between the two groups ($\chi^2 \approx 0.68$, p-value ≈ 0.41 ; excludes “I do not recall” responses, Figure 3a). A smaller percentage — but still a majority — of administrators reported being “Extremely familiar” with their instances’ privacy policies (52%; Figure 3b); this was the mode answer for both samples. When coded on a 0-3 scale, the sample means were both “Very familiar,” with no statistical difference between the two groups ($F \approx 0.43$, p-value ≈ 0.51). Only the medians of the samples differed (“Very familiar” vs. “Extremely familiar” for the outreach and snowball sampled groups respectively).



(a) Were you involved in the initial crafting of [instance]'s privacy policy? (Q8) (b) How familiar are you with [instance]'s privacy policy? (Q9) (c) How much of [instance]'s privacy policy was written from scratch...? (Q11)

Figure 3: Self-reported administrator involvement, familiarity, and authorship suggest misconceptions when compared with reported policy origins and last reported policy review.

These results might suggest that most administrators are actively involved in privacy policy creation and maintenance. When asked how the initial privacy policy was crafted, respondents gave seemingly contradictory responses. For example, 14 administrators reported having authored the entirety of their privacy policy (Figure 3c). However, when we navigated to their instance's policy endpoint, we found that seven instances were using the default policy. One instance had replaced the default text with a code of conduct. Two other instances replaced the default text with much shorter and less specific text discussing data. The remaining policies were inaccessible (404 errors, or behind a login page).

We also observed this phenomenon among respondents who reported authoring at least some (but not all) of their privacy policies. Thirteen respondents reported being involved in originating their instances' privacy policies using an existing policy that they modified beyond the name of their server. However, when we navigate their instances' privacy policies, six of the 13 display the default policy. One instance replaced the text of the default privacy policy with a link to the terms of service. Another replaced the text with a link to a subsection of their about page, which points to masto.host's privacy policy. It is possible that in this case, the privacy policy text may have been auto-generated by masto.host.

We received forty-two responses to our open-ended question (Q24) asking administrators if they would like to share any additional information about their policies' origins. Among these responses we saw evidence of both laissez-faire:

I literally wrote it stoned out of my mind while I was setting up the server. lol

– and hands-on administrator behavior:

We try to review and update it, but only do so based on sources that we would trust, for instance at the encouragement of a law focused Mastodon instance that posted what they believed to be good policies.

One respondent substantiated our suspicion that some administrators the use names of different documents interchangeably or confuse them:

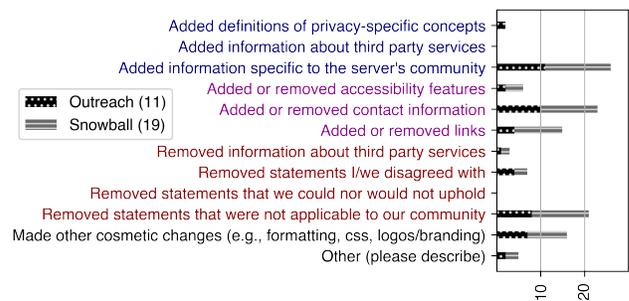


Figure 4: Aggregated counts of types of edits to existing privacy policies or tool/template output (Q15, Q20). Blue text indicates additions, red deletions, and magenta either or both. Most options were available for both questions captured by this histogram; for more information, see Appendix D.

I may have confused the privacy policy with the server rules in these questions?

Among respondents who reported *not* being involved in the privacy policy's creation, most reported using a managed hosted service and all but two reported that they were using the default policy. Table 3 in Appendix D aggregates responses across questions Q8, Q11, Q13, and Q18.

Policy alterations and privacy incidents. We suspect that administrators who had experience with privacy incidents would incorporate their insights into their instance's policies. To better understand this trajectory, we asked about alterations to policies and past privacy incidents (Q14–Q15, Q19–Q20, Q25–Q28).

Thirty respondents reported making changes (other than the name of the server) to their reference policies or to the output of the tool(s) they used to generate the policies (Figure 4). Given the evidence that some administrators were confused about which governance document we were referring to and that we may not have identified the actual privacy policy endpoint correctly, these

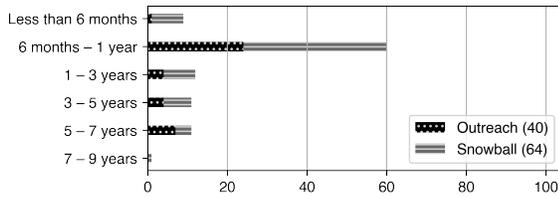


Figure 5: The length of time that respondents have been functioning as administrators for their servers (Q2).

data should be interpreted as modifications to the documents that the administrators *thought* were the privacy policies.

Four respondents provided us with additional information and context for their changes. We corroborated these changes against the instance’s privacy policies. In all cases, administrators made updates to the default policy to clarify terminology or provide additional context for how the Mastodon software works. For example, one policy describes the function and scope of the `_mastodon_session_id` cookie.

We then asked respondents if they had experienced scenarios that would require them to engage with their privacy policy, ranging from rare events like requests to disclose information about users, servers, etc. (Q25), to more mundane questions from their instance members (Q28). Only four respondents report having ever received a request to disclose information about their members or server; we did not ask follow up questions about these incidents out of concern that it might deter participants from continuing. For a comparison of frequencies of each of these incidents, see Table 4 in Appendix D, which aggregates over questions Q25–Q28.

5.3 Administrator Characteristics (D4, D5)

Mastodon administrators are not a monolith; their characteristics and values inform the sociotechnical choices that impact users. We seek to better understand the characteristics of administrators who select and craft governance documents like privacy policies to inform our understanding of how privacy is conceptualized and operationalized on Mastodon (RQ3).

5.3.1 Administrator Tenure. Most of the 104 respondents who answered our question about the length of time they had been administrators had been in their roles for less than a year. The average, median, and mode tenures for both groups all fall within the 6 months – 1 year bucket, with the exception of the outreach group’s mode, which rounds to the 1 – 3 years bucket. There is no evidence of a statistical difference in the variability of length of tenure between the two groups ($F \approx 1.96$, $p\text{-value} \approx 0.16$).

Figure 5 depicts the response counts for each of the brackets we provided. Note that these brackets are not of equal size: regardless of exogenous factors, we would expect unpaid server administration to have high abandonment rates, so we chose a finer granularity for the first year of administrative experience (6 mos.) and more coarse-grained brackets after the first year (2 yrs.). Mastodon was founded in 2016; only one administrator selected this option and it was for an online community that long predates Mastodon.

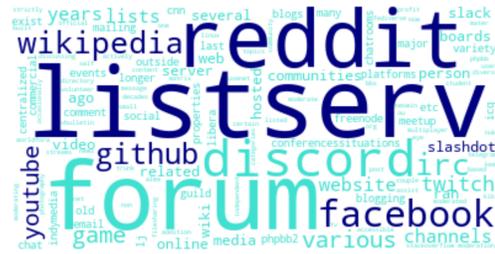


Figure 6: Platforms on which administrators have had prior content moderation experience. The seven platforms in navy are the seven options we provided; free-text responses are in turquoise (Q30).

Do you have any background or prior experience in (X)?

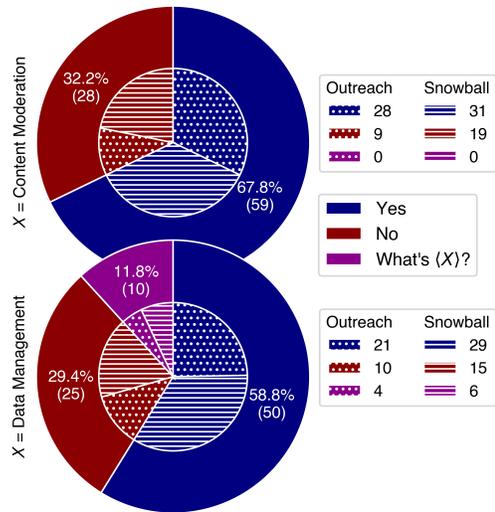


Figure 7: Respondents had similar rates of experience with both content moderation and data management (Q29, Q37).

5.3.2 Content Moderation and Data Management Experience. We found no evidence of a statistical difference in the proportion of administrators who report having a background or prior experience in content moderation vs. data management ($\chi^2 \approx 1.14$, $p\text{-value} \approx 0.29$, Figure 7). The 59 respondents who reported having prior experience in content moderation were then asked to specify the platforms on which they acquired this experience (Q30). Fifty-five out of the 58 administrators who responded to this question reported more than one platform (95%). Figure 6 depicts the breadth of these responses as a word cloud.

5.3.3 Technical Backgrounds. Administrators who rely on existing policies, tools, or templates to choose or construct their privacy policies are effectively outsourcing legal expertise. We were interested in the backgrounds of administrators who instead chose to write some portion of their policies themselves. Therefore, we asked these respondents about their teams’ technical and educational experiences in the five domains listed in Table 1 (Q12).

Domains	Formal educational			Informal educational			Professional background			No one had this background		
	Outreach	Snowball	%Δ	Outreach	Snowball	%Δ	Outreach	Snowball	%Δ	Outreach	Snowball	%Δ
Law	2 (11%)	5 (18%)	-7	3 (16%)	11 (39%)	-23	2 (11%)	4 (14%)	-4	13 (68%)	12 (43%)	26
Policy	2 (11%)	2 (7%)	3	4 (21%)	15 (54%)	-33	3 (16%)	5 (18%)	-2	11 (58%)	10 (36%)	22
Security	2 (11%)	6 (21%)	-11	8 (42%)	12 (43%)	-1	10 (53%)	14 (50%)	3	4 (21%)	3 (11%)	10
Privacy	2 (11%)	3 (11%)	-0	8 (42%)	18 (64%)	-22	5 (26%)	9 (32%)	-6	5 (26%)	3 (11%)	16
Operations	1 (5%)	4 (14%)	-9	3 (16%)	11 (39%)	-23	10 (53%)	11 (39%)	13	6 (32%)	7 (25%)	7

Table 1: Administrator educational or technical backgrounds. Forty-eight (19 direct outreach; 29 snowball) of the 53 respondents (20 direct outreach; 33 snowball) who were asked about their educational or technical backgrounds responded (Q12). Two responses were excluded from analysis for selecting both “No one had this background” and one of the other training options. The blue highlighted cells mark backgrounds that are represented by more than 50% of the respondents (darker) or more than 33% of the respondents (lighter). Absolute differences between the two groups are highlighted (red: ≥ 30%, orange: ≥ 20%, yellow: ≥ 10%, white: <10%). Respondents could also provide free-text explanations of their backgrounds. Six did, naming “Licenses,” “Media,” “Fine Arts Degree,” “Web Design,” “Information Technology,” and “Computer Science,” as relevant backgrounds for their teams.

For each domain, we asked about the nature of their training in terms of three categories: formal educational, informal educational, and professional. We defined formal educational backgrounds as degrees or coursework in the domain, while informal educational backgrounds included workshops and self-study. Respondents could choose more than one domain and more than one type of background. We gave respondents the option to indicate that no one on their team had a particular background; recall that because respondents could skip any question except for the consent form and server name, we should not necessarily interpret the complement of “No one had this background” as an affirmation that someone had this background. We also did not prohibit respondents from selecting both “No one had this background” and one of the training options. One respondent from the snowball sampled group did not respond to some of these questions; another from this group responded inconsistently and was removed from the analysis.

Finding: Limited experience in policy or law. Most of the teams making decisions about privacy policies reported having no formal background in policy or law, confirming our assumption from section 2.5. However, a significant number reported *informal* education in policy. This number also represents the most significant difference between our two samples: while only 21% of the outreach group reported independently learning about policy, a whopping 54% of the snowball sampled group did.

Finding: Autodidacts unevenly represented. Across all areas, the snowball sampled group showed a greater propensity for informal education than the outreach sample. The member-researcher who seeded this sample is an academic who may be attracting autodidacts in these specific areas due to their Fediverse activity.

Finding: Operations + Security = DevOps strongly represented? Both groups reported substantial proportions of their administrative teams having professional experience in both security and operations. This association is especially strong in the outreach group, where the majority of the respondents report these backgrounds. Our interpretation is that these individuals are DevOps professionals. While we did not anticipate this background, it is unsurprising given the technical expertise required to set up and successfully run a Mastodon instance.

5.3.4 Demographics. One of the biggest questions facing the Fediverse right now is being framed as a trade-off between user privacy and content moderation and the impact it has on marginalized users [25, 28, 91]. All respondents were asked whether they identify as having a minoritized or marginalized identity (M/M, Q43–Q44). Respondents were free to interpret their minority or marginalized identity as they saw fit; our intention was to give space for respondents to express their demographic identifiers relative to their own cultural context. Figure 8(a) depicts the results. Respondents who answered “Yes” or “I’m not sure” to Q43 were then asked to select from a subset of provided axes of M/M, including an “Other” option (Figure 8(b); counts of “Other” not depicted, but text appears grouped with additional clarifying free text responses in 8(c)). Sexual orientation or preference is the most common axis in both groups, followed by gender identity or expression. We note that in the snowball sampled group, disability or neurodivergence was one rank higher than it was in the direct outreach group. Due to the small sample sizes, we do not compute whether this is a statistically significant difference.

Figure 8(d) depicts the frequencies of the *number* of distinct axes on which respondents report having M/M identities. Over half of those who have M/M identities report multiple axes (7/11 in the outreach group; 14/19 in the snowball sampled group). Different marginalized identities aren’t independent categories. For example, gender and sexuality are deeply enmeshed categories and one recent study found that in the US, disability was 156% more common in among those who identified as LGBTQ+ than among those who did not [92].

6 DISCUSSION

Our initial text analysis suggested that most administrators simply use the default Mastodon privacy policy that ships with the Mastodon software. This is unsurprising in the context of prior work on software configuration choices in privacy-sensitive contexts and contemporaneous efforts with our work [43, 66]. However, our survey results suggest a more nuanced perspective. A sizable minority of our sample of Mastodon administrators (≈ 10%) either made minor edits or wrote new policies from scratch. Several administrators

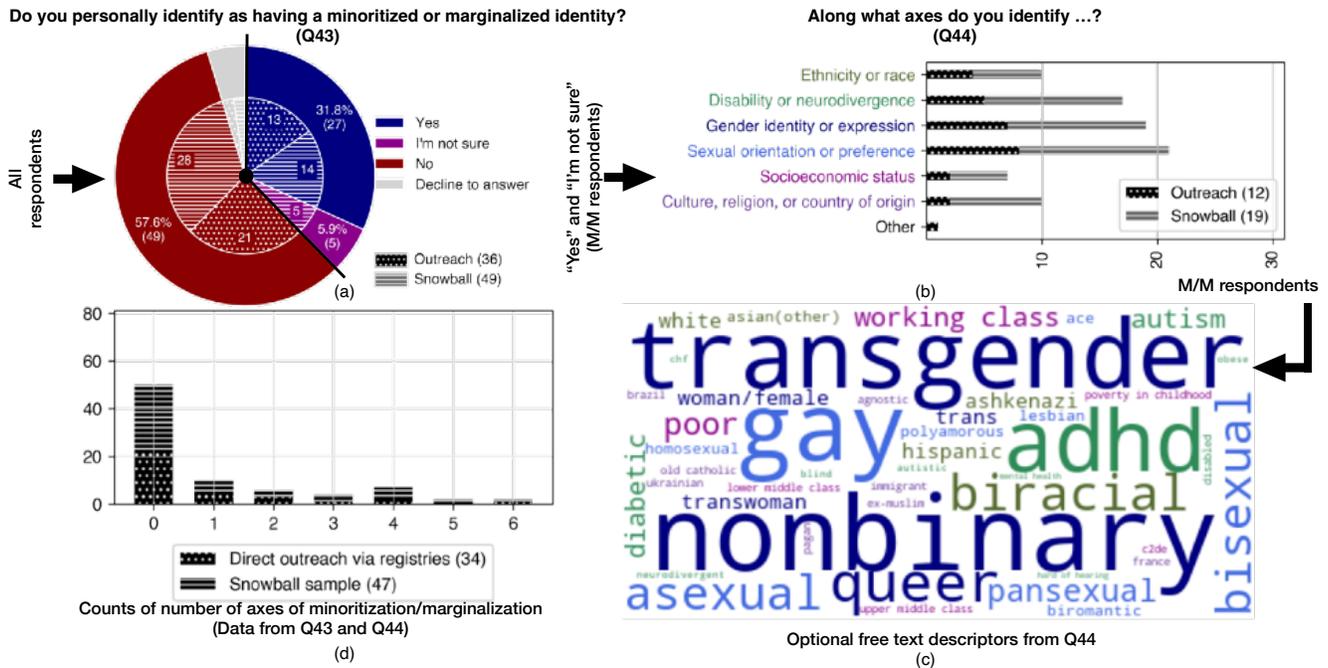


Figure 8: All respondents were asked whether they identified as having a marginalized or minoritized (M/M) identity (a). Depending on their responses, they were then asked to identify on which axes they had such an identity (b) and were given the option of selecting more than one axes, as well as providing explanatory text (c). While most respondents did not identify as having a M/M identity, those who did often specified more than one axis (d).

expressed past or future interest in tailoring their privacy policies but were concerned about the legal implications, e.g.,

We currently use the default [...] We have considered modifying it, but have always been scared off by the legal risks that might pose – all the things we might need to account for, might be responsible for. We have been so worried we might miss some things, we haven’t added any at all.

Appendix Section D.6 expands on similar responses to Q24 and describes some of the interactions the authors have had with administrators throughout this study. Our initial text analysis, survey responses, and interactions all suggest that the folkloristic assumption that Mastodon administrators have strongly than average technical computing backgrounds fails to capture the true heterogeneity of administrator backgrounds, interests, and subsequent needs.

While there were certainly some administrators who believed privacy policies for their small, recreational servers to be irrelevant, just as many administrators wrote about seeking governance guidance and expertise from other other instances. We now discuss opportunities for sociotechnical infrastructure to grow and strengthen in terms of assets-based community design (ABCD) [54], an alternative design approach for when purely needs-based approach may fail to deliver sustainable impact on account of neglecting strengths and affordances that already exist [71].

Asset: Administrator-members. Mastodon administrators are typically members of the communities for which they provide services. This membership gives them personal insight into the effects of technical choices (e.g., where to host the software) on social outcomes (e.g., the consequences of data seizure by law enforcement). Furthermore, different instances represent different communities, each having different privacy needs. Even if administrators lack expertise in the law, they have unique competencies in what the users of their service want and need. Given that our folkloristic assumption of over-representation in the Fediverse of LGBTQ+ identities relative to the population at large has held up, we see administrator-members as uniquely positioned to advocate for the privacy and needs of these uniquely vulnerable communities.

Challenge: Double frustration. Because administrators’ incentives are aligned with members, they may experience frustration about privacy policies from both the client and service provider perspective. Indeed, we saw both the hopelessness of trying to adhere to what Horstmann et al. [42] described as the unidirectional guidance of privacy compliance *and* the well-documented lack of interest in privacy policies from users [7, 46, 86].

Asset: Instance cooperation, not competition. Despite operating in a regulatory ecosystem not built for them, several Mastodon administrators reported being members of administrator groups and seeking out mentorship and resources to support their community-building efforts. Administrators talk to each other and have a culture of sharing resources as a community: for example, some

respondents reported circulating and discussing governance documents with other instances that had shared privacy interests. Furthermore, the authors of one major Mastodon fork (Hometown⁷) package technical and social documentation together [49].

While our survey respondents expressed apprehension about the implications of editing privacy policies, many who were interested in doing so reported expertise in security or operations. Furthermore, administrators who run their own servers or customize Mastodon must have sufficient technical knowledge to keep their instances running.

Recommendation: community privacy guides/templates. The broader Fediverse community could benefit from building on these efforts to create, e.g. privacy templates that incorporate technical privacy concerns relevant to members, *even if they do not immediately address any current regulations*. These documents could be integrated across instances and jurisdictions into an interactive guide to choosing from a set of privacy policies, similar to Github’s `choosealicense.com` guide for open source license selection.

We believe we have seen some limited evidence to support the folkloristic assumption that most administrators see themselves as content moderators (if they support content moderation at all), rather than data managers. However, this folkloristic assumption was itself based on the idea that administrators are managing their own software, either locally or through a cloud provider. The growth of managed hosting complicates our understanding of what it means to “run your own server,” since hosting platforms often provide their own licenses. Thus, administrators who report little engagement with user data or user privacy considerations are not necessarily eschewing their duties if they have outsourced data management to a hosting platform. Thus, the recommendation of community resources applies more to administrators who manage their own servers.

Asset/Need: Managing context collapse. Ten respondents reported administering more than one server, six of whom provided the names of their other servers (Q4–Q7). We did not ask respondents why they administered more than one server, but we manually investigated the six cases where respondents provided their other servers’ names and inferred the following possible reasons:

- *Grouping distinct interests.* Two respondents each operate two distinct servers that differ in both topic typology: locality/interest and identity/interest.
- *Separating work and personal access to the Fediverse.* One respondent operates a Mastodon server for their place of employment; they are the only member of the other server they operate.
- *Access control for sensitive content.* One respondent operates several fetish-themed servers that vary in their content (e.g., whether NSFW).

Given the representation of M/M identities in our data set, this behavior suggests that federated model of having just one account does not address issues of context collapse [61, 98]. On the “asset” side, Mastodon provides a flexible way for users to find communities that support their distinct privacy needs, but on the “need”

side, current tools for policy creation and understanding do not foreground information relevant to those needs.

Recommendation: Narrative information formats. Tang et al. [93] refined prior work on privacy policy comprehensibility [74] to show that non-expert end-users often misunderstand not only the purpose of privacy policies, but the terms therein. For example, Reidenberg et al. [79] showed that end-users interpreted the absence of explicit enumerations of permissions as ambiguous, while privacy experts interpreted the omission as permissive (an interpretation that is correct in the sense of legal support). Work on formalizing privacy policies [6, 23, 35] can help disambiguate these documents in a form that’s amenable to computational adaptation, while usable privacy efforts are needed to translate them into user-oriented desiderata [17, 18, 50, 58]; however, most “usable privacy” interventions prescribe a one-size-fits-all approach to interpreting privacy documents that does not account for the Fediverse’s diversity of purposes nor parity of incentives and needs across administrators and users. This gap can be bridged by adapting privacy document formalization to meet user-facing needs through *narrative generation*, which allows for interactive exploration of specific scenarios that users may modify and query [21], based on findings that narrative-structured information supports more accurate mental models and recall [94]. When paired with a collective approach to synthesizing adaptable policy templates that serve the diverse needs of Fediverse communities (per previous recommendation), we see synthesis of narrative-structured situations generated from policy information as a promising avenue towards tool support that will help both users and administrators.

7 CONCLUSION

We conducted a text analysis of 351 Mastodon privacy policies and subsequent online survey of 104 administrators to understand how they select or craft privacy policies for their Mastodon servers. To our knowledge, this is the first comprehensive analysis of privacy policies in the Fediverse and the first study of administrator privacy choices.

Our findings support some folkloric assumptions about demographics and knowledge backgrounds of administrators, such as strong representation of technical interests and more experience with community moderation than with data management (though not as much as we thought!). Despite respondents’ technical backgrounds, they do not necessarily engage with the Mastodon software at the level we had expected going into this study, especially if they use managed hosting systems. As a result, some administrators edit or replace boilerplate language, leading to inconsistencies with the underlying code. Since policy selection and management tasks are already challenging for expert compliance professionals, it is no surprise that the average Mastodon administrator would struggle, too. These findings point toward a need for “privacy-enhancing technology” support not just for user understanding of privacy policies, but also policy creation on the part of network administrators.

ACKNOWLEDGMENTS

This material is based upon work supported by the National Science Foundation under Grants No. 1846122. Emma Tosch was partially

⁷<https://github.com/hometown-fork/hometown>

supported by National Science Foundation Grant No. 2330961 while performing this work.

The authors thank Erin McBride for her background in law and contracts and John Foley for editing feedback. We are especially grateful to the Fediverse administrators who participated in this study and communicated their feedback and insight throughout this research project.

REFERENCES

- [1] Lillian Ablon, Paul Heaton, Diana Catherine Lavery, and Sasha Romanosky. 2016. *Consumer Attitudes toward Data Breach Notifications and Loss of Personal Information*. Rand Corporation, New York.
- [2] Lada Adamic and Eytan Adar. 2005. How to Search a Social Network. *Social Networks* 27, 3 (2005), 187–203.
- [3] Raghav Ahooja. 2023. *Section 230 and the Fediverse: The ‘Instances’ of Mastodon’s Immunity and Liability*. SSRN.
- [4] Danielle Allen, Woojin Lim, Eli Frankel, Joshua Simons, Divya Siddarth, and Glen Weyl. 2023. *Ethics of Decentralized Social Technologies: Lessons from Web3, the Fediverse, and Beyond*. Harvard College.
- [5] Jacopo Anderlini and Carlo Milani. 2022. Emerging Forms of Sociotechnical Organisation: The Case of the Fediverse. *Digital Platforms and Algorithmic Subjectivities* 167 (2022), 167–181.
- [6] Grigoris Antoniou, David Billington, Guido Governatori, and Michael J Maher. 1999. On the modeling and analysis of regulations. (1999).
- [7] Manon Arcand, Jacques Nantel, Mathieu Arles-Dufour, and Anne Vincent. 2007. The impact of reading a web site’s privacy statement on perceived control over privacy and perceived trust. *Online Information Review* 31, 5 (2007), 661–681.
- [8] Glenn Bass, Lorraine Kenny, Guinee Sarah, and Madeline Wood. 2019. Symposium: The Tech Giants, Monopoly Power, and Public Discourse, Glenn Bass (Ed.). Knight First Amendment Institute at Columbia University.
- [9] Michael Bernstein, Andrés Monroy-Hernández, Drew Harry, Paul André, Katrina Panovich, and Greg Vargas. 2011. 4chan and/b: An Analysis of Anonymity and Ephemerality in a Large Online Community. In *Proceedings of the International AAAI Conference on Web and Social Media*, Vol. 5, 50–57.
- [10] Ames Bielenberg, Lara Helm, Anthony Gentilucci, Dan Stefanescu, and Honggang Zhang. 2012. The Growth of Diaspora — A Decentralized Online Social Network in the Wild. In *2012 Proceedings IEEE INFOCOM Workshops*. IEEE, 13–18.
- [11] Steven Bird, Ewan Klein, and Edward Loper. 2009. *Natural language processing with Python: analyzing text with the natural language toolkit*. O’Reilly Media, Inc.
- [12] Monica Chin. 2021. How a University Got Itself Banned from the Linux Kernel. (30 April 2021). <https://www.theverge.com/2021/4/30/22410164/linux-kernel-university-of-minnesota-banned-open-source>
- [13] Danielle Keats Citron and Daniel J Solove. 2022. Privacy Harms. *Boston University Law Review* 102 (2022), 793–863.
- [14] Cindy Cohn and Rory Mir. 2023. *FBI Seizure of Mastodon Server Data is a Wakeup Call to Fediverse Users and Hosts to Protect their Users*. <https://www.eff.org/deeplinks/2023/07/fbi-seizure-mastodon-server-wakeup-call-fediverse-users-and-hosts-protect-their>
- [15] Wayne L Cornelius. 1985. UTILIZATION OF AN ELECTRONIC BULLETIN BOARD. (1985).
- [16] Geoffrey Cramer. 2023. *An Empirical Study of Trust & Safety Engineering in Open-Source Social Media Platforms*. Ph.D. Dissertation. Purdue University.
- [17] Lorrie Faith Cranor. 2003. P3P: Making privacy policies more useful. *IEEE Security & Privacy* 1, 6 (2003), 50–55.
- [18] Lorrie Faith Cranor. 2005. Giving Notice: Why Privacy Policies and Security Breach Notifications aren’t Enough. *IEEE Communications Magazine* 43, 8 (2005), 18–19.
- [19] Lorrie Faith Cranor. 2012. Necessary but not sufficient: Standardized mechanisms for privacy notice and choice. *J. on Telecomm. & High Tech. L.* 10 (2012), 273.
- [20] Kimberlé Crenshaw. 1991. Mapping the Margins: Intersectionality, Identity Politics, and Violence against Women of Color. *Stanford Law Review* 43, 6 (1991), 1241–1299. <http://www.jstor.org/stable/1229039>
- [21] Chinmaya Dabral, Emma Tosch, and Chris Martens. 2022. Exploring Consequences of Privacy Policies with Narrative Generation via Answer Set Programming. *Conference on Programming Languages and the Law*.
- [22] Anastasia Danilova, Alena Naiakshina, Stefan Horstmann, and Matthew Smith. 2021. Do you Really Code? Designing and Evaluating Screening Questions for Online Surveys with Programmers. In *2021 IEEE/ACM 43rd International Conference on Software Engineering (ICSE)*. IEEE, 537–548.
- [23] Henry DeYoung, Deepak Garg, Limin Jia, Dilsun Kaynar, and Anupam Datta. 2010. Experiences in the logical specification of the HIPAA and GLBA privacy laws. In *Proceedings of the 9th annual ACM workshop on Privacy in the electronic society*. 73–82.
- [24] Don A. Dillman, Jolene D. Smyth, and Leah Melani Christian. 2014. *Internet, phone, mail, and mixed-mode surveys: The tailored design method*. John Wiley & Sons.
- [25] Elizabeth Dwoskin. 2023. *Fleeing Elon Musk’s X, the quest to re-create ‘Black Twitter’*. <https://www.washingtonpost.com/technology/2023/08/06/musk-black-twitter-spill/>
- [26] Alex Feerst. 2019. *Your Speech, Their Rules: Meet the People Who Guard the Internet*. <https://medium.com/p/ab58fe6b9231>
- [27] Dror G Feitelson. 2023. “We do not appreciate being experimented on”: Developer and Researcher Views on the Ethics of Experiments on Open-Source Projects. *Journal of Systems and Software* (2023), 111774.
- [28] Jonathan Flowers. 2022. *The Whiteness of Mastodon*. <https://techpolicy.press/the-whiteness-of-mastodon/>
- [29] Seth Frey, PM Krafft, and Brian C Keegan. 2019. “This Place Does What It Was Built For” Designing Digital Institutions for Participatory Change. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019), 1–31.
- [30] Ujwal Gadiraju, Ricardo Kawase, Stefan Dietze, and Gianluca Demartini. 2015. Understanding Malicious Behavior in Crowdsourcing Platforms: The case of online surveys. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. 1631–1640.
- [31] Robert W Gehl. 2015. The case for alternative social media. *Social Media+ Society* 1, 2 (2015), 2056305115604338.
- [32] Sarah Gilbert. 2023. Towards Intersectional Moderation: An Alternative Model of Moderation Built on Care and Power. *Proceedings of the ACM on Human-Computer Interaction* 7, CSCW2 (2023), 1–32.
- [33] Dan Goodin. 2023. Mastodon Fixes Critical “TootRoot” Vulnerability Allowing Node Hijacking. (6 7 2023). <https://arstechnica.com/security/2023/07/mastodon-fixes-critical-tootroot-vulnerability-allowing-node-hijacking/>
- [34] Leo A Goodman. 1961. Snowball sampling. *The annals of mathematical statistics* (1961), 148–170.
- [35] Guido Governatori, Vineet Padmanabhan, Antonino Rotolo, and Abdul Sattar. 2009. A defeasible logic for modelling policy-based intentions and motivational attitudes. *Logic Journal of IGPL* 17, 3 (2009), 227–265.
- [36] Thomas Groß. 2021. Validity and reliability of the scale internet users’ information privacy concerns (iupc). *Proceedings on Privacy Enhancing Technologies* (2021).
- [37] Hana Habib, Sarah Pearman, Jiamin Wang, Yixin Zou, Alessandro Acquisti, Lorrie Faith Cranor, Norman Sadeh, and Florian Schaub. 2020. “It’s a scavenger hunt”: Usability of Websites’ Opt-Out and Data Deletion Choices. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–12.
- [38] Oliver L Haimson, Daniel Delmonaco, Peipei Nie, and Andrea Wegner. 2021. Disproportionate removals and differing content moderation experiences for conservative, transgender, and black social media users: Marginalization and moderation gray areas. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW2 (2021), 1–35.
- [39] M Hassett, C Lowder, and D Rutan. 1992. Use of computer network bulletin board systems by disabled persons.. In *Proceedings of the Annual Symposium on Computer Application in Medical Care*. American Medical Informatics Association, 151.
- [40] Nick Hatley and Courtney Kennedy. 2022. *How we keep our online surveys from running too long*. Pew Research Center. <https://www.pewresearch.org/decoded/2022/12/08/how-we-keep-our-online-surveys-from-running-too-long/>
- [41] Jiahui He, Haris Bin Zia, Ignacio Castro, Aravindh Raman, Nishanth Sastry, and Gareth Tyson. 2023. Flocking to mastodon: Tracking the great twitter migration. In *Proceedings of the 2023 ACM on Internet Measurement Conference*. 111–123.
- [42] Stefan Albert Horstmann, Samuel Domiks, Marco Gutfleisch, Mindy Tran, Yasemin Acar, Veelasha Moonsamy, and Alena Naiakshina. 2024. “Those things are written by lawyers, and programmers are reading that.” Mapping the Communication Gap Between Software Developers and Privacy Experts. *Proceedings on Privacy Enhancing Technologies* 1 (2024), 151–170.
- [43] Sohyeon Hwang, Priyanka Nanayakkara, and Yan Shvartzshnaider. 2023. Whose Policy? Privacy Challenges of Decentralized Platforms. In *CHI’23 Workshops: Designing Technology and Policy Simultaneously: Towards a Research Agenda and New Practice*.
- [44] Jane Im, Ruiyi Wang, Weikun Lyu, Nick Cook, Hana Habib, Lorrie Faith Cranor, Nikola Banovic, and Florian Schaub. 2023. Less is Not More: Improving Findability and Actionability of Privacy Controls for Online Behavioral Advertising. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–33.
- [45] Ramkumar Janakiraman, Joon Ho Lim, and Rishika Rishika. 2018. The Effect of a Data Breach Announcement on Customer Behavior: Evidence from a Multichannel Retailer. *Journal of Marketing* 82, 2 (2018), 85–105.
- [46] Carlos Jensen, Colin Potts, and Christian Jensen. 2005. Privacy practices of Internet users: Self-reports versus observed behavior. *International Journal of Human-Computer Studies* 63, 1-2 (2005), 203–227.

- [47] Shagun Jhaver, Seth Frey, and Amy X Zhang. 2023. Decentralizing Platform Power: A Design Space of Multi-level Governance in Online Social Platforms. *Social Media+ Society* 9, 4 (2023), 20563051231207857.
- [48] Michael Karanicolas. 2021. Too Long; Didn't Read: Finding Meaning in Platforms' Terms of Service Agreements. *University of Toledo Law Review* 52 (2021), 1–25.
- [49] Darius Kazemi. 2019. *Run Your Own Social*. Hometown. <https://runyourown.social>
- [50] Patrick Gage Kelley, Joanna Breese, Lorrie Faith Cranor, and Robert W Reeder. 2009. A "nutrition label" for privacy. In *Proceedings of the 5th Symposium on Usable Privacy and Security*. 1–12.
- [51] Nancy S Kim. 2022. Adhesive Terms and Reasonable Notice. *Seton Hall Law Review* 53 (2022), 85–147.
- [52] Jon Kleinberg. 2001. Small-world phenomena and the dynamics of information. *Advances in neural information processing systems* 14 (2001).
- [53] Jon M Kleinberg. 2000. Navigation in a small world. *Nature* 406, 6798 (2000), 845–845.
- [54] John Kretzmann and John P McKnight. 1996. Assets-based Community Development. *Nat'l Civic Rev.* 85 (1996), 23.
- [55] Ravi Kumar, Jasmine Novak, and Andrew Tomkins. 2006. Structure and evolution of online social networks. In *Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining*. 611–617.
- [56] Lucio La Cava, Sergio Greco, and Andrea Tagarelli. 2021. Understanding the growth of the Fediverse through the lens of Mastodon. *Applied network science* 6 (2021), 1–35.
- [57] Lucio La Cava, Sergio Greco, and Andrea Tagarelli. 2022. Network analysis of the information consumption-production dichotomy in mastodon user behaviors. In *Proceedings of the International AAAI Conference on Web and Social Media*, Vol. 16. 1378–1382.
- [58] Jialiu Lin, Shahriyar Amini, Jason I Hong, Norman Sadeh, Janne Lindqvist, and Joy Zhang. 2012. Expectation and purpose: understanding users' mental models of mobile app privacy through crowdsourcing. In *Proceedings of the 2012 ACM conference on ubiquitous computing*. 501–510.
- [59] John MacFarlane. 2006. Pandoc. <https://pandoc.org>
- [60] Aymeric Mansoux and Roel Roscam Abbing. 2020. Seven theses on the fediverse and the becoming of FLOSS. (2020).
- [61] Alice E Marwick and Danah Boyd. 2011. I tweet honestly, I tweet passionately: Twitter users, context collapse, and the imagined audience. *New media & society* 13, 1 (2011), 114–133.
- [62] Mastodon. 2019. *Setting up your new instance*. Mastodon gGmbH. <https://docs.joinmastodon.org/admin/setup/#info>
- [63] Peter Mayer, Yixin Zou, Byron M. Lowens, Hunter A. Dyer, Khue Le, Florian Schaub, and Adam J. Aviv. 2023. Awareness, Intention, (In) Action: Individuals' Reactions to Data Breaches. *ACM Transactions on Computer-Human Interaction* (2023).
- [64] Peter Mayer, Yixin Zou, Florian Schaub, and Adam J Aviv. 2021. "Now I'm a bit (angry:)" Individuals' Awareness, Perception, and Responses to Data Breaches that Affected Them. In *30th USENIX Security Symposium (USENIX Security 21)*. 393–410.
- [65] Abraham Mhaidli, Selin Fidan, An Doan, Gina Herakovic, Mukund Srinath, Lee Matheson, Shomir Wilson, and Florian Schaub. 2023. Researchers' Experiences in Analyzing Privacy Policies: Challenges and Opportunities. *Proceedings on Privacy Enhancing Technologies* 4 (2023), 287–305.
- [66] Abraham H Mhaidli, Yixin Zou, and Florian Schaub. 2019. "We Can't Live Without {Them!}" App Developers' Adoption of Ad Networks and Their Considerations of Consumer Risks. In *Fifteenth Symposium on Usable Privacy and Security (SOUPS 2019)*. 225–244.
- [67] Matthew N Nicholson, Brian C Keegan, and Casey Fiesler. 2023. Mastodon Rules: Characterizing Formal Rules on Popular Mastodon Instances. In *Companion Publication of the 2023 Conference on Computer Supported Cooperative Work and Social Computing*. 86–90.
- [68] Razieh Nokhbeh Zaeem and K Suzanne Barber. 2021. A Large Publicly Available Corpus of Website Privacy Policies Based on DMOZ. In *Proceedings of the Eleventh ACM Conference on Data and Application Security and Privacy*. 143–148.
- [69] Charlie Parker, Sam Scott, and Alistair Geddes. 2019. Snowball sampling. *SAGE research methods foundations* (2019).
- [70] Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, et al. 2011. Scikit-learn: Machine learning in Python. *the Journal of machine Learning research* 12 (2011), 2825–2830.
- [71] Lucy Pei and Bonnie Nardi. 2019. We Did it Right, but It was Still Wrong: Toward Assets-based Design. In *Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–11.
- [72] Irene Pollach. 2007. What's wrong with online privacy policies? *Commun. ACM* 50, 9 (2007), 103–108.
- [73] Andrea Prat and Andrea Valletti. 2022. Attention Oligopoly. *American Economic Journal: Microeconomics* 14, 3 (2022), 530–557.
- [74] Robert W Proctor, M Athar Ali, and Kim-Phuong L Vu. 2008. Examining usability of web privacy policies. *Intl. Journal of Human-Computer Interaction* 24, 3 (2008), 307–328.
- [75] Nicholas Proferes, Naiyan Jones, Sarah Gilbert, Casey Fiesler, and Michael Zimmer. 2021. Studying reddit: A systematic overview of disciplines, approaches, methods, and ethics. *Social Media+ Society* 7, 2 (2021), 20563051211019004.
- [76] Sheizaf Rafaeli. 1984. The electronic bulletin board: A computer-driven mass medium. *Social Science Micro Review* 2, 3 (1984), 123–136.
- [77] Sheizaf Rafaeli and Robert J LaRose. 1993. Electronic bulletin boards and "public goods" explanations of collaborative mass media. *Communication Research* 20, 2 (1993), 277–297.
- [78] Aravindh Raman, Sagar Joglekar, Emiliano De Cristofaro, Nishanth Sastry, and Gareth Tyson. 2019. Challenges in the Decentralised Web: The mastodon case. In *Proceedings of the Internet Measurement Conference (ACM IMC)*. 217–229.
- [79] Joel R Reidenberg, Travis Breaux, Lorrie Faith Cranor, Brian French, Amanda Grannis, James T Graves, Fei Liu, Aleccia McDonald, Thomas B Norton, Rohan Ramanath, et al. 2015. Disagreeable Privacy Policies: Mismatches between Meaning and Users' Understanding. *Berkeley Tech. Lj* 30 (2015), 39.
- [80] Amaury Rousseau. 2023. instances.social. [instances.social](https://github.com/TheKinrar/instances.com). <https://github.com/TheKinrar/instances.com>
- [81] Alan Z Rozenshtein. 2023. Moderating the fediverse: Content moderation on distributed social media. *J. Free Speech L.* 3 (2023), 217.
- [82] Eduard Sabo, Mirela Riveni, and Dimka Karastoyanova. 2023. Decentralized Networks Growth Analysis: Instance Dynamics on Mastodon. In *International Conference on Complex Networks and Their Applications*. Springer, 366–377.
- [83] Florian Schaub, Rebecca Balebako, Adam L. Durity, and Lorrie Faith Cranor. 2015. A design space for effective privacy notices. In *Eleventh symposium on usable privacy and security (SOUPS 2015)*. 1–17.
- [84] Mastodon Administrators, Scholars, and Users. 2020. *An Open Letter from the Mastodon Community*. <https://www.sunclipse.org/wp-content/downloads/2020/01/open-letter.pdf>
- [85] H. Jeff Smith. 1993. Privacy policies and practices: Inside the organizational maze. *Commun. ACM* 36, 12 (1993), 104–122.
- [86] H Jeff Smith, Tamara Dinev, and Heng Xu. 2011. Information privacy research: an interdisciplinary review. *MIS quarterly* (2011), 989–1015.
- [87] Daniel J Solove and Danielle Keats Citron. 2017. Risk and Anxiety: A theory of Data-breach Harms. *Texas Law Review* 96 (2017), 737–786.
- [88] Mukund Srinath, Soundarya Sundareswara, Pranav Venkit, C Lee Giles, and Shomir Wilson. 2023. Privacy Lost and Found: An Investigation at Scale of Web Privacy Policy Availability. In *Proceedings of the ACM Symposium on Document Engineering 2023*. 1–10.
- [89] Alex Stamos and Sara Shah. 2023. *Common Abuses on Mastodon: A Primer*. Stanford University. <https://fsi.stanford.edu/news/common-abuses-mastodon-primer>
- [90] Marshall Steinbaum. 2022. Establishing Market and Monopoly Power in Tech Platform Antitrust Cases. *The Antitrust Bulletin* 67, 1 (2022), 130–145.
- [91] Morgan Sung. 2023. *For Bluesky to thrive, it needs sex workers and Black Twitter*. <https://techcrunch.com/2023/05/02/bluesky-black-twitter-sex-workers-culture/>
- [92] Chris R. Surfus. 2023. A Statistical Understanding of Disability in the LGBT Community. *Statistics and Public Policy* 10, 1 (2023), 2188056. <https://doi.org/10.1080/2330443X.2023.2188056>
- [93] Jenny Tang, Hannah Shoemaker, Ada Lerner, and Eleanor Birrell. 2021. Defining Privacy: How Users Interpret Technical Terms in Privacy Policies. *Proceedings on Privacy Enhancing Technologies* (2021).
- [94] Perry W Thorndyke. 1977. Cognitive structures in comprehension and memory of narrative discourse. *Cognitive psychology* 9, 1 (1977), 77–110.
- [95] Sean Tilley. 2017. *A Quick Guide to The Free Network*. Medium. <https://medium.com/we-distribute/a-quick-guide-to-the-free-network-c069309f334>
- [96] Christine Utz, Sabrina Amft, Martin Degeling, Thorsten Holz, Sascha Fahl, and Florian Schaub. 2023. Privacy Rarely Considered: Exploring Considerations in the Adoption of Third-Party Services by Websites. *Proceedings on Privacy Enhancing Technologies* 1 (2023), 5–28.
- [97] Ana Valens. 2021. *Mastodon is crumbling — and many blame its creator*. DailyDot. <https://www.dailydot.com/debug/mastodon-fediverse-eugen-rochko/>
- [98] Jessica Vitak. 2012. The impact of context collapse and privacy on social network site disclosures. *Journal of broadcasting & electronic media* 56, 4 (2012), 451–470.
- [99] Nina Wallerstein and Bonnie Duran. 2003. *The theoretical, historical, and practice roots of CBPR*. Jossey-Bass.
- [100] Christine Lemmer Webber, Jessica Tallon, Erin Shepherd, Amy Guy, and Evan Prodromou. 2018. *ActivityPub*. W3C.
- [101] Alan F Westin. 1968. *Privacy and Freedom*. Ig Publishing.
- [102] Shomir Wilson, Florian Schaub, Aswarth Abhilash Dara, Frederick Liu, Sushain Cherivirala, Pedro Giovanni Leon, Mads Scharup Andersen, Sebastian Zimmeck, Kanthashree Mysore Sathyendra, N Cameron Russell, et al. 2016. The creation and analysis of a website privacy policy corpus. In *Proceedings of the 5th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. 1330–1340.

- [103] Kou Zhongbao and Zhang Changshui. 2003. Reply networks on a bulletin board system. *Physical Review E* 67, 3 (2003), 036117.
- [104] Matteo Zignani, Sabrina Gaito, and Gian Paolo Rossi. 2018. Follow the “mastodon”: Structure and evolution of a decentralized online social network. In *Proceedings of the International AAAI Conference on Web and Social Media*, Vol. 12. 541–550.
- [105] Diana Zulli, Miao Liu, and Robert Gehl. 2020. Rethinking the “social” in “social media”: Insights into topology, abstraction, and scale on the Mastodon social network. *New Media & Society* 22, 7 (2020), 1188–1205.

A NETWORK FEATURE COMPARISON

Mastodon is often compared to other software and platforms in its functionality and user experience. For readers seeking to understand how Mastodon differs from well-known social software and platforms, we provide Table 2. We focus on general-purpose (i.e., not tailored to specific topics or activities like Strava for exercise or ResearchGate for scientific research) social networking platforms. We compare each entity in terms of the social networks (i.e., connectivity between users), subgraph software affordances (i.e., tools for managing subgroups within the network), and post visibility. These three concepts influence privacy policies via what Jhaver et al. [47] refer to as the “middle level” of a multi-level governance structure commonly seen in online social networking platforms.

B REGISTRIES

Many researchers use registries to obtain a sample of Fediverse instances. There are further ethical and statistical issues to consider when using registries, which we describe below.

The Mastodon organization provides a curated centralized registry of servers that abide by the “Mastodon covenant”⁸ on their website. Listings are strictly opt-in and administrators must submit an application for approval. In addition to meeting content requirements (e.g., no hate speech), servers listed on `joinmastodon.org` must be open and have low barriers for new member sign-ups.⁹

⁸<https://joinmastodon.org/covenant>; this document was added to `joinmastodon.org` in October 2022 and describes the necessary criteria to be listed in the `joinmastodon.org` registry, e.g. no hate speech.

⁹At the time of this writing, only about 5% of Mastodon servers crawled by `joinmastodon.org` appear to register with `joinmastodon.org`. We do not know how many of these servers are *qualified* to register.

While `joinmastodon.org` advises that they do not boost new servers and that instances looking to grow should depend on other methods to do so, it is unlikely that servers that are explicitly *not* looking to grow would register. Therefore, we expect `joinmastodon.org` to undercount smaller, more private servers that are seeking stability over growth. Thus we believe `joinmastodon.org` does not produce a representative sample of servers: we believe data set may be biased due to the additional effort of registration and the server requirements for listed instances.

We augmented our recruitment pool with another centralized registry: `instances.social`. An individual developer-administrator (@`thekinrar`) maintains the `instances.social` registry [80]. The source code available for `instances.social` only includes the web application code; neither the data itself nor the data collection code are available. From the Github issue history and personal communication with @`thekinrar`, we know that the data are a combination of voluntary registrations and instance crawling.¹⁰ Through personal communication with @`thekinrar`, we learned that their crawler only uses instance-level public information; i.e., no user information such as profiles, follows, or posts.

Despite only using public information, we know from the `instances.social` issue history that some administrators whose instances are listed would prefer they not be.¹¹ This lack of community consensus over what makes something on the internet “public” as been discussed in other contexts than the Fediverse (see Section 4; Figure 9 depicts administrator reactions to being listed in a registry, including a request for a privacy policy). Knowing there could be community trust issues, we looked to prior work in analogous contexts to guide us on ethical best practices when working with administrators and engaging with instances as research subjects. We followed the principles outlined in Feitelson [27]’s post-incident survey of open source developers; these principles align with previously-published recommendations for research on Reddit, which evoke community-based participatory research [75, 99].

We strongly discourage other researchers from assuming that approaches that made sense in other contexts (e.g., Reddit or Twitter)

¹⁰<https://github.com/TheKinrar/instances/issues/120>

¹¹<https://github.com/TheKinrar/instances/issues/159>

	Platform	Connectivity	Coordination	Visibility (Post)	Subgraph	Subgraph Tooling	Visibility (Group Post)
Web 1.0	email	decentralized	federated	private	n/a	n/a	n/a
	BBS	decentralized	siloed	public	message board	*	public
Web 2.0	Mastodon	decentralized	federated	public	instance	yes	public
	Facebook	centralized	siloed	private	group	no	private
	Reddit	centralized	siloed	public	subreddit	yes	public
	Twitter	centralized	siloed	public	n/a	n/a	n/a
	Instagram	centralized	siloed	private	n/a	n/a	n/a
	Wikipedia	centralized	siloed	public	n/a	n/a	n/a
	TikTok	centralized	siloed	public	n/a	n/a	n/a

Table 2: Comparison of general social networking software in terms of selected features. All visibility features are labeled according to their historic levels, since these features influence administrator and user experiences when they move to other platforms. Reddit and Mastodon both offer sophisticated tooling (i.e., Mastodon’s API and direct code manipulation and Reddit’s regular expressions language) to subgraph moderators. Since bulletin board systems (BBSs) have their origins amongst hobbyists who often ran “homebrewed” software [15, 76], similar tooling was possible but is undocumented.

would apply in the Fediverse. For example, Nicholson et al. [67] analyze instance rulesets by sampling the registry instances .social and scraping the server public pages. They make the case that because this data is publicly facing, designed for newcomers, and because their results are reported in aggregate, scraping is permissible. They argue that they follow Proferes et al. [75]’s approach to contextual ethical data collection – a higher standard than IRBs require – and take community context into account. Critically, Nicholson et al. [67] argue that community context suggests that asking administrators for permission to scrape public pages is not required.

Before conducting the study presented in this paper, we would have agreed with Nicholson et al. [67]. However, our interactions with administrators suggest that some would not consent to this. As researchers, we are concerned that transferring norms from centralized social media could erode trust with Fediverse administrators and hurt the health of the Fediverse in the process. We consider this an evolving topic and best iterated upon with the input and consent of Fediverse members.

C SURVEY INSTRUMENT

For each question in the survey, we indicate whether the question was conditionally visible. We code each question type using Pew Research’s online question typology [40]. The following question types appear in our survey:

- Open-ended** Respondents are presented with a textbox; this is commonly known as a “freetext” question.
- Stand-alone** Respondents are presented with a select-one (i.e., radio button) question, possibly with an “Other” option that is freetext.
- Check-all** Respondents may select any number of options (i.e., checkbox) that may include an “Other” option having a textbox field.
- Battery items** These are grouped questions, each of which may be Stand-alone or Check-all; they all have the same stem and must be presented to the respondent as a group.

Participants were shown a progress bar and presented with one question per page (i.e., they only saw one question at a time), unless presented with a battery question. This allowed us to implement fine-grained conditional branching/display logic and minimize the extraneous information presented to participants. While respondents could skip questions, they could not view every question available due to the conditional branching we used. Questions were grouped by topic. Respondents were shown signposts at the start of each topic section, e.g.:

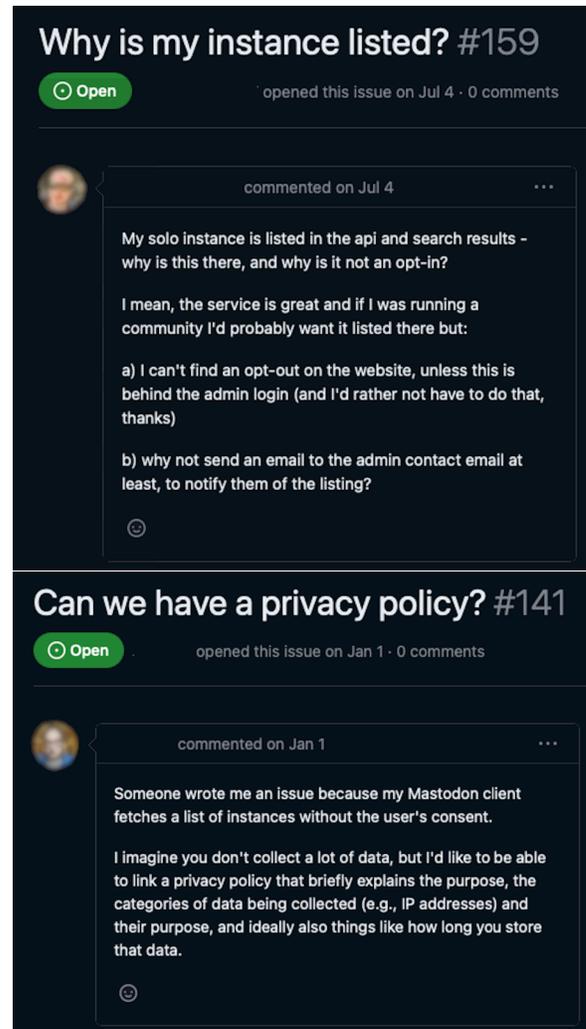
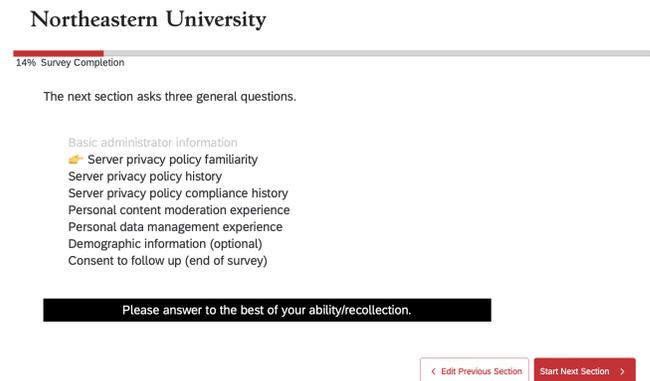


Figure 9: Some Mastodon administrators react to inclusion in instance registries with unease, prompting one administrator to request a privacy policy for instances .social. At the time of this writing, the issue requesting a privacy policy has been open for over one year.



-
- Q1** What Mastodon server are you an administrator for?
Conditional Display: None
Question type: Open-ended
- Q2** How long have you been an administrator for [your instance]?
Conditional Display: None
Question type: Select one
O1 Less than 6 months
O2 6 months – 1 year
O3 1 – 3 years
O4 3 – 5 years
O5 5 – 7 years
O6 7 – 9 years
O7 More than 9 years
O8 I do not recall
O9 I prefer not to say
- Q3** In what year was [your instance] established?
Conditional Display: In **Q2**, **O6** or **O7** selected
Question type: No
O1 2023
O2 2022
O3 2021
O4 2020
O5 2019
O6 2018
O7 2017
O8 2016
O9 2015
O10 2014
O11 Before 2014
O12 I do not know
O13 I do not recall
O14 I prefer not to say
- Q4** Are you an administrator on any other Mastodon servers?
Conditional Display: None
Question type: Stand-alone
O1 Yes
O2 No
- Q5** How many other Mastodon servers are you an administrator on?
Conditional Display: In **Q4**, **O1** selected
Question type: Stand-alone
O1 0
O2 1
O3 2
O4 3
O5 Greater than 3
- Q6** You answered “Yes” to the question “Are you an administrator on any other Mastodon servers?” but responded “0” to the question “How many other Mastodon servers are you an administrator on?” If this was a mistake, please navigate back and correct your responses. If it was not a mistake, please describe what you mean by this.
Conditional Display: In **Q5**, **O1** selected
Question type: Open-ended
- Q7** Please list the names of the top 3 other Mastodon servers for which you are an administrator, if you feel comfortable doing so.
Conditional Display: In **Q4**, **O1** selected
Question type: Open-ended
- Q8** Were you involved in the initial crafting of [your instance]’s privacy policy?
Conditional Display: None
Question type: Stand-alone
O1 Yes
O2 No
O3 I do not recall
- Q9** How familiar are you with [your instance]’s privacy policy?
Conditional Display: None
Question type: Stand-alone
O1 Not familiar at all
O2 Somewhat familiar
O3 Very familiar
O4 Extremely familiar
- Q10** When was the last time you read through [your instance]’s privacy policy?
Conditional Display: None
Question type: Stand-alone
O1 Less than 6 months ago
O2 6 months – 1 year ago
O3 3 – 5 years ago
O4 5 – 7 years ago
O5 7 – 9 years ago
O6 9 – 11 years ago
O7 More than 11 years ago
O8 Never
- Q11** How much of [your instance]’s privacy policy was written from scratch (by you or someone on your team)?
Conditional Display: None
Question type: Stand-alone
O1 None
O2 Some
O3 About half
O4 Most
O5 All
O6 I do not know
O7 I do not recall
O8 I prefer not to say
- Q12** To the best of your knowledge, did any of [your instance]’s privacy policy authors have background knowledge or experience in any of the following areas?
Conditional Display: In **Q11**, none of **O1**, **O6**, **O7**, nor **O8** selected
Question type: Check-all
O1 Law, Formal Educational (e.g., degree or coursework)
O2 Law, Informal Educational (e.g., workshops, self-study)
O3 Law, Professional Background
O4 Law, No one had this background
O5 Policy, Formal Educational (e.g., degree or coursework)
O6 Policy, Informal Educational (e.g., workshops, self-study)

- O7 Policy, Professional Background
 - O8 Policy, No one had this background
 - O9 Security, Formal Educational (e.g., degree or coursework)
 - O10 Security, Informal Educational (e.g., workshops, self-study)
 - O11 Security, Professional Background
 - O12 Security, No one had this background
 - O13 Privacy, Formal Educational (e.g., degree or coursework)
 - O14 Privacy, Informal Educational (e.g., workshops, self-study)
 - O15 Privacy, Professional Background
 - O16 Privacy, No one had this background
 - O17 Operations, Formal Educational (e.g., degree or coursework)
 - O18 Operations, Informal Educational (e.g., workshops, self-study)
 - O19 Operations, Professional Background
 - O20 Operations, No one had this background
 - O21 Other Background 1, Formal Educational (e.g., degree or coursework)
 - O22 Other Background 1, Informal Educational (e.g., workshops, self-study)
 - O23 Other Background 1, Professional Background
 - O24 Other Background 1, No one had this background
 - O25 Other Background 2, Formal Educational (e.g., degree or coursework)
 - O26 Other Background 2, Informal Educational (e.g., workshops, self-study)
 - O27 Other Background 2, Professional Background
 - O28 Other Background 2, No one had this background
 - O29 Other Background 3, Formal Educational (e.g., degree or coursework)
 - O30 Other Background 3, Informal Educational (e.g., workshops, self-study)
 - O31 Other Background 3, Professional Background
 - O32 Other Background 3, No one had this background
- Q13** Did you or someone on your team create [your instance]’s privacy policy using a tool or template service?
Conditional Display: In Q11, O5 is not selected
Question type: Stand-alone
 O1 Yes
 O2 I think so
 O3 No
 O4 I do not know
 O5 I do not recall
 O6 I prefer not to say
 O7 What’s a tool or template service?
- Q14** Did you or someone on your team make changes to the tool or template output?
Conditional Display: In Q13, O1 or O2 selected
Question type: Stand-alone
 O1 Yes
 O2 I think so
 O3 No
 O4 I do not know
 O5 I do not recall
 O6 I prefer not to say
- Q15** What kinds of changes did [you *[sic]*] or your team you make to the output of the tool? Please check all that apply.
Conditional Display: In Q14, O1 or O2 selected
Question type: Check-all
 O1 Added or removed links
 O2 Added or removed accessibility features
 O3 Added or removed contact information
 O4 Added information specific to the server’s community
 O5 Added information about third party *[sic]* services
 O6 Added definitions of privacy-specific concepts
 O7 Removed statements I/we disagreed with
 O8 Removed statements that we could not or would not uphold
 O9 Removed statements that were not applicable to our community
 O10 Removed information about third party services *[sic]*
 O11 Made other cosmetic changes (e.g., formatting, css, logos...[,] branding)
 O12 Other (please describe)
- Q16** Were there any tools or templates you or your team considered but did not use?
Conditional Display: In Q14, O1 or O2 selected
Question type: Stand-alone
 O1 Yes
 O2 I think so
 O3 No
 O4 I do not know
 O5 I do not recall
 O6 I prefer not to say
- Q17** Please share why you or your team did not use the other tools or templates you’d considered. Alternative(s) ...
Conditional Display: In Q16, O1 or O2 selected
Question type: Check-all
 O1 was/were paywalled (not gratis)
 O2 was/were proprietary (not libre)
 O3 were unenforceable
 O4 put too much responsibility on server administrators
 O5 put too much responsibility on users
 O6 insufficiently addressed privacy concerns
 O7 premitted advertising
 O8 prohibited advertising
 O9 too prescriptive about third party *[sic]* services
 O10 too proscriptive about third party *[sic]* services
 O11 were not compatible with other policy documents (e.g., terms of service, code of conduct, business plan, community guidelines, etc.)
 O12 were too hard to understand
 O13 were not compliant with relevant laws
 O14 Other Reason (please describe)
- Q18** Did you or someone on your team create [your instance]’s privacy policy text by modifying or using the text of one or more existing privacy policies?
Conditional Display: In Q11, O5 not selected
Question type: Stand-alone
 O1 Yes
 O2 I think so

- O3** No
O4 I do not know
O5 I do not recall
O6 I prefer not to say
- Q19** Other than the name of the server, did you or someone on your team make changes to the existing privacy policy when crafting [your instance]'s privacy policy?
Conditional Display: In **Q18**, **O1** or **O2** selected
Question type: Stand-alone
O1 Yes
O2 I think so
O3 No
O4 I do not know
O5 I do not recall
O6 I prefer not to say
- Q20** Other than the name of server, what kinds of changes did you make to the existing policy? Please check all that apply.
Conditional Display: In **Q19**, **O1** or **O2** selected
Question type: Check-all
O1 Added or removed links
O2 Added information specific to the server's community
O3 Added or removed accessibility features
O4 Added or removed contact information
O5 Removed statements I/we disagreed with
O6 Removed statements that we could nor or would not uphold
O7 Removed statements that were not applicable to our community
O8 Made other cosmetic changes (e.g., formatting, css, logos...[,] branding)
O9 Other (please describe)
- Q21** Were there any existing policies you or your team considered but did not choose to use?
Conditional Display: In **Q18**, **O1** or **O2** selected
Question type: Stand-alone
O1 Yes
O2 I think so
O3 No
O4 I do not know
O5 I do not recall
O6 I prefer not to say
- Q22** Please share why you or your team did not use certain other policies. Alternative(s)...
Conditional Display: In **Q21**, **O1** or **O2** selected
Question type: Check-all
O1 were unenforceable
O2 put too much responsibility on server administrators
O3 put too much responsibility on users
O4 insufficiently addressed privacy concerns
O5 permitted advertising
O6 prohibited advertising
O7 were not compatible with other policy documents (e.g., terms of service, code of conduct, business plan, community guidelines, etc.)
O8 were too hard to understand
O9 were not compliant with relevant laws
- O10** Other Reason (please describe)
- Q23** What is/are the names or urls of the policy/policies on which [your instance]'s privacy policy is based.
Conditional Display: In **Q18**, **O1** or **O2** selected
Question type: Open-ended
- Q24** Is there anything else you'd like to share about the origins of [your instance]'s privacy policy?
Conditional Display: None
Question type: Open-ended
- Q25** Have you or your team ever gotten a request to share information about the content, users, or usage of [your instance]?
Conditional Display: None
Question type: Stand-alone
O1 Yes
O2 I think so
O3 No
O4 I do not know
O5 I do not recall
O6 I prefer not to say
- Q26** Have you or your team ever needed to invoke [your instance]'s privacy policy in your capacity as an administrator of it?
Conditional Display: None
Question type: Stand-alone
O1 Yes
O2 I think so
O3 No
O4 I do not know
O5 I do not recall
O6 I prefer not to say
- Q27** Have any members of [your instance] raised privacy issues with you or your team?
Conditional Display: None
Question type: Stand-alone
O1 Yes
O2 I think so
O3 No
O4 I do not know
O5 I do not recall
O6 I prefer not to say
- Q28** Have any members of [your instance] asked questions regarding your privacy policy?
Conditional Display: None
Question type: Stand-alone
O1 Yes
O2 I think so
O3 No
O4 I do not know
O5 I do not recall
O6 I prefer not to say
- Q29** Do you have any background or prior experience as a content moderator?
Conditional Display: None
Question type: Stand-alone
O1 Yes

- O2** No
O3 What's content moderation?
- Q30** On what platforms other than Mastodon have you had experience in content moderation? Please check all that apply.
Conditional Display: In Q29, O1 is selected
Question type: Check-all
O1 Reddit
O2 Facebook
O3 Youtube
O4 Wikipedia
O5 Slashdot
O6 Github
O7 BBS/listserv (please specify context, if willing)
O8 Other (please specify)
- Q31** Please select the content moderation issues that you have experience with. Do not search for definitions of unfamiliar terms; try to answer this question quickly and confidently without additional resources.
Conditional Display: None
Question type: Check-all
O1 Hate speech
O2 Cyberbullying
O3 Harassment/Stalking
O4 Doxxing
O5 Disinformation
O6 Misinformation
O7 Spam
O8 Scams
O9 Sock puppets
O10 Cryptocurrency
O11 NFTs
- Q32** Have you ever encountered (personally or as a witness) what you would personally consider to be a privacy violation as part of your content moderation experience?
Conditional Display: In Q29, O1 selected
Question type: Stand-alone
O1 Yes
O2 Possibly yes
O3 Not sure
O4 Definitely no
O5 I do not recall
O6 I prefer not to say
- Q33** To the best of your knowledge, was this privacy violation prohibited by the platform's privacy policy?
Conditional Display: In Q32, O1 or O2 selected
Question type: Stand-alone
O1 Yes
O2 Possibly yes
O3 Not sure
O4 Definitely no
O5 I was not sufficiently familiar with the platform's privacy policy
O6 I prefer not to say
- Q34** To the best of your knowledge, was this privacy violation prohibited by the platform's code of conduct?
- Conditional Display:** In Q32, O1 or O2 selected
Question type: Stand-alone
O1 Yes
O2 Possibly yes
O3 Not sure
O4 Definitely no
O5 I was not sufficiently familiar with the platform's code of conduct
- Q35** Do you believe this privacy violation should have been covered in some way by the platform's governance documents? Select all that apply.
Conditional Display: In Q32, O1 or O2 selected
Question type: Check-all
O1 Yes, in the privacy policy
O2 Yes, in the code of conduct
O3 Yes, in the terms of service
O4 Yes, other
O5 No
O6 Not sure
- Q36** Is there any additional specific information you would like to share about your content moderation experience?
Conditional Display: In Q29, O3 not selected
Question type: Open-ended
- Q37** Do you have any background or prior experience in data management?
Conditional Display: None
Question type: Stand-alone
O1 Yes
O2 No
O3 What's data management
O4 I prefer not to say
- Q38** Please select the data management issues, concerns, or techniques that you are familiar with. Do not search for definitions of unfamiliar terms; try to answer this question quickly and confidently without additional resources.
Conditional Display: None
Question type: Check-all
O1 PII
O2 Secure servers
O3 Airgapping
O4 Differential privacy
O5 TLS
O6 De-identification
O7 GDPR compliance
O8 CORI compliance
O9 COPPA compliance
O10 CCPA compliance
O11 CPRA compliance
O12 Encryption
O13 Databases
- Q39** In what contexts other than Mastodon have you had experience in data management? Please check all that apply.
Conditional Display: In Q37, O1
Question type: Check-all
O1 Work (industrial)

- O2 Work (medical)
 O3 Work (government)
 O4 Work (academic)
 O5 Personal
 O6 Other (please specify)
- Q40** Do you have a background in or prior experience with collecting PII?
Conditional Display: In Q37, O1
Question type: Stand-alone
 O1 Yes, a lot
 O2 Yes, a little
 O3 No
 O4 What's PII?
 O5 I prefer not to say
- Q41** Do you have a background in or prior experience with data encryption?
Conditional Display: In Q37, O1
Question type: Stand-alone
 O1 Yes, a lot
 O2 Yes, a little
 O3 No
 O4 What's data encryption?
 O5 I prefer not to say
- Q42** Is there any additional specific information you would like to share about your data management experience?
Conditional Display: In Q37, O1
Question type: Open-ended
- Q43** Do you personally identify as having a minoritized or marginalized identity?
Conditional Display: None
Question type: Stand-alone
 O1 Yes
 O2 No
 O3 I'm not sure
 O4 Decline to answer
- Q44** Along what axes do you identify as having a minoritized or marginalized identity? If you feel comfortable, please specify categories in the text boxes.
Conditional Display: In Q43, O1 or O3 selected
Question type: Check-all/Open-ended
 O1 Gender identity or expression
 O2 Ethnicity or race
 O3 Sexual orientation or preference
 O4 Socioeconomic status
 O5 Culture, religion, or country of origin
 O6 Disability or neurodivergence
 O7 Other
- Q45** There are specific demographic questions we ask when studying populations in the United States. Please indicate whether any of the following are true.
Conditional Display: None
Question type: Battery items
 [Your instance] is hosted in the United States.
 O1 Yes
- O2 No
 O3 Not sure
 O4 Decline to answer
 You are based in the United States.
- O5 Yes
 O6 No
 O7 Not sure
 O8 Decline to answer
 O9 Yes
 O10 No
 O11 Not sure
 O12 Decline to answer
- Q46** Please indicate which of the following U.S. Census racial categories you identify as. You may leave this question blank if you do not wish to answer.
Conditional Display: In Q45, at least one of O1, O5, or O9 is selected
Question type: Check-all
 O1 White
 O2 Black or African American
 O3 American Indian and Alaska Native
 O4 Native Hawaiian and Other Pacific Islander
 O5 Other
- Q47** Please indicate whether or not you identify as having Latin American ethnicity, culture, or origins.
Conditional Display: In Q45, at least one of O1, O5, or O9 is selected
Question type: Stand-alone
 O1 Yes
 O2 No
 O3 Prefer not to say
- Q48** Please indicate your gender identity, but only if you feel comfortable doing so.
Conditional Display: In Q45, at least one of O1, O5, or O9 is selected
Question type: Check-all
 O1 Woman
 O2 Man
 O3 Nonbinary
 O4 Prefer to self-describe
 O5 Prefer not to say
- Q49** If you feel comfortable sharing, let us know how often your gender identity aligns with your gender assignment at birth.
Conditional Display: In Q45, at least one of O1, O5, or O9 is selected
Question type: Stand-alone
 O1 Always
 O2 Most of the time
 O3 Sometimes
 O4 Never
 O5 Prefer not to say
- Q50** Is there anything else you'd like to share about your identity?
Conditional Display: None
Question type: Open-ended

Q51 Would you be interested in participating in any follow-up studies at a later point in time, such as an extended interview or beta-testing policy authoring tools?

Conditional Display: None

Question type: Stand-alone

O1 Yes

O2 Maybe

O3 No

Q52 Is there a different email address you would like us to use for future contact? Leaving this field blank indicates the email address we have for you is acceptable for further outreach.

Conditional Display: In Q51, O1 or O2 selected

Question type: Open-ended

D COMPLETE DATA ANALYSIS

We include here the complete descriptive and exploratory analyses of all survey questions by topic, in sequential question order. Unless otherwise noted, we reproduce relevant charts and tables and include back-references to aid in flow and legibility.

D.1 Basic Information (Q1–Q7)

Q1: *What Mastodon server are you an administrator for?* This question had additional instructions that varied depending on the recruitment method we used. All participants were asked this question. It was the only required question after the consent form, i.e., participants could not continue the survey without responding.

In review, concerns were raised that requiring respondents to name their instance would have a deleterious effect on response rates, due to findings in prior work [66, 96]. We have no evidence to suggest that requiring respondents to provide their instance name deterred participation. In fact, our evidence suggests that this question is correlated with high quality responses. This judgment is based on both the literature about reliability of online surveys and the authors’ past experiences conducting online surveys.

Consider the following: we directly emailed 349 administrators. Thirty-four emails bounced (18 soft, 16 hard). We know that there is a high churn in the availability of Mastodon instances. We do not know how many of the email addresses we sent our outreach to are even active. As one administrator said in response to a recent spam attack:

Some instance admins got reminded that they had an instance. And we also learned there are A LOT of abandoned instances out there with their door wide open for registration without approval.^a

^a<https://techcrunch.com/2024/02/20/spam-attack-on-twitter-x-rival-mastodon-highlights-fediverse-vulnerabilities/>

Thus while at 17% (55/315), our response rate is lower than our target response rate of 20%, we did not consider this to be remarkably low for a recruitment method via weak social ties.

All recruitment messages included the statement:

This survey is not tailored to your instance. It does not collect IP addresses. It is a separate request from another email you may have received from us regarding consent to use the text of your privacy policy.

Critically, while the 315 administrators who received this message (i.e., email did not bounce, had not previously opted out of Qualtrics emails) via email were told which instance we believed they were administrators for, they were not told via email that they would need to provide any identifying information.

Fifty-five of the administrators we reached out to followed the survey link. The survey landing page is the consent form. Respondents must affirm that they are Mastodon administrators over the age of 18 who understand the content of the consent form in order to progress to the body of the survey. Only 45 of these respondents progressed passed the consent form.

This consent form does not explicitly state that respondents must provide the instance name. The first question in the survey after the consent form asks for the instance name. Therefore, since this is respondents’ first exposure to the idea that they will have to provide the instance name, survey breakoff at this first question would provide us with evidence that doing so adversely affects participation.

Forty-five respondents saw this question. That is, of the 55 respondents who clicked on their emailed links, ten did not complete the consent form. Barring collusion, we believe it is reasonable to assume that only 45 administrators were aware that the survey asked them to provide their instance name.

Only six respondents out of the 45 did not provide a verifiable instance name. Although we turned off recording IP addresses and thus do not have access to this information directly, Qualtrics does do some proprietary analysis of IP addresses for spam detection. Qualtrics marked three of the six respondents who did not progress as spam. No other respondents were marked as spam. Therefore, $\approx 7\%$ (3/42) of respondents broke off at the instance name question.

We note that two of the respondents gave server names that were not in our original list. One of these was a legitimate Mastodon server name. The other was a single character response; it is likely that this behavior was an attempt to occlude the respondents’ identity. Since these two individual responses did not appear to include any other inconsistencies or behavior that would suggest an attempt to subvert the integrity of the data collection process, we included them in the analysis.

We performed a similar analysis toward the end of our data collection period for the snowball sampled survey. The proportion of useful surveys for the snowball sample at that time was lower than the proportion of useful surveys for our direct outreach sample: 70% vs. 87%. We use as the denominator the total number of responses that made it past the consent page and were not marked by Qualtrics as spam. We use as the numerator the total number of valid/usable responses. The standards we use for determining the numerator differ slightly between groups: we included responses that obfuscated server in the first group, but not for the second.

The snowball sampled survey had the potential to reach non-English-speaking individuals and we have reason to suspect that it did: the two respondents who did not make it past the consent form indicated that they did not understand the content of the consent

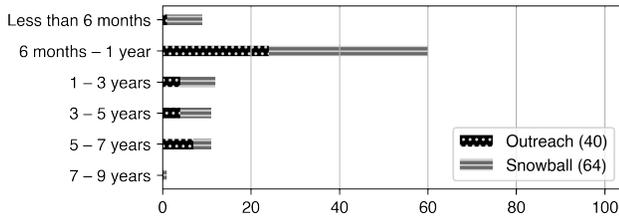
form. Furthermore, the instance name of one respondent appears in both centralized repositories that we used (i.e., `joinmastodon.org` and `instances.social`), but was not included in the outreach group due to non-English server languages.

Four respondents provided servernames that we could not verify as valid or invalid. This was due to either intentionally obscuring the servername or because the server now appears to be offline. There were two respondents not deemed to be “hackers” who were in the former group: one respondent supplied input that was stripped somewhere in the Qualtrics processing (i.e., it appears as the empty string to us), one wrote “n/a.”

Among our snowball sample, three respondents in total triggered our validation questions by selecting tenure as administrators longer than Mastodon has existed. One respondent explicitly stated in a freetext response: I wanted To see What you’re asking our admins! The other respondent stopped replying earlier. Neither provided valid Mastodon server names (`fgngfn` and `bitchface.cunt`).

Q2: How long have you been an administrator for [your instance]?

Discussed in Section 5.3.1:



Q3: In what year was [your instance] established? We used this question for validating unexpected responses to Q2. Mastodon was first released in 2016. Our survey ran in 2023. Therefore, we would not have expected respondents to choose either the “7–9 years” or the “More than 9 years.” Zero respondents in the direct outreach group selected either of these options. Several respondents selected this option in the snowball sampled group, but only one belonged to a high integrity response. We discuss some of the low integrity responses in our discussion of Q1.

Q4–Q7: [Questions about administering other servers.] We discuss our findings regarding multiple server administration and its implications in Section 6 under **Asset/Need: Managing context collapse** and summarize this information. Ten respondents reported administering more than one server (Q4); we have highlighted responses that provided server names:

Source	Other Servers Administered (Q5)	Other Servers Named (Q7)
1 outreach	2	3*
2 outreach	1	0
3 outreach	1	0
4 snowball	1	1
5 snowball	2	2
6 snowball	1	0
7 snowball	1	1
8 snowball	2	1*
9 snowball	3	0
10 snowball	1	1

Respondents could not view Q5 if they did not answer Q4 (Are you an administrator on any other Mastodon servers?) affirmatively. Respondent 1 appeared to interpret Q7 as asking for the names of *all* servers administered, including the server for which we identified them in their list. We believe respondent 8 interpreted Q5 as asking for the total number of servers administered. It is possible that respondent 9 administers either 3 or 4 servers in total. Given the phrasing of Q4, we believe the remaining responses represent an accurate interpretation of the questions.

Question Q6 is a data integrity check that was not triggered by any of the respondents.

D.2 Privacy Policy Familiarity (Q8–Q10)

Q8: Where you involved in the initial crafting of [instance]’s privacy policy? Discussed in Section 5.2; chart appears in Figure 3a:

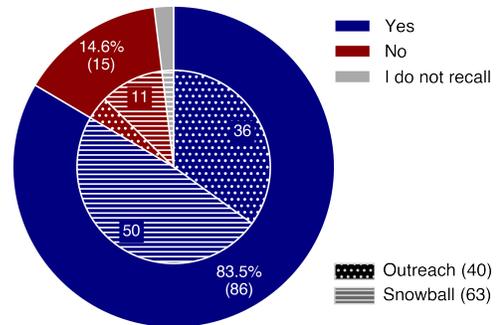
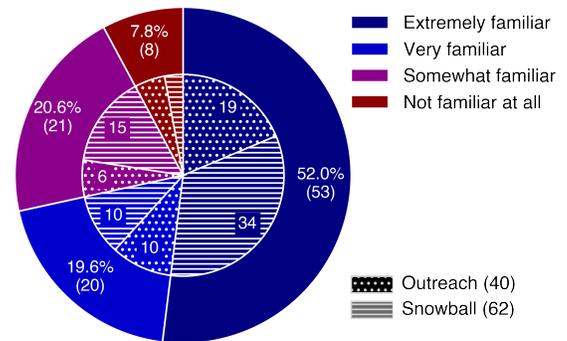


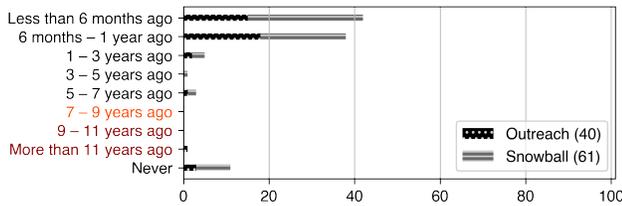
Table 3 aggregates over Q8, Q11, Q13, and Q18.

Q9: How familiar are you with [your instance]’s privacy policy? Discussed in Section 5.2; chart appears in Figure 3b:



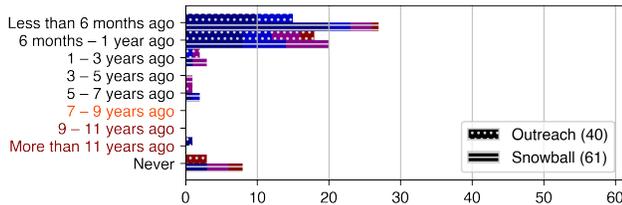
Q10: When was the last time you read through [your instance]’s privacy policy? We did not include detailed discussion of this question in Section 5.2 since the results are consistent with what we would expect, given the findings of Q2.

We summarize responses in the chart below; the date ranges highlighted in orange and red represent possibly invalid responses, since they predate the existence of the Mastodon software:



To further validate consistency, we flagged responses where respondents reported having read the privacy policy prior their reported start of tenure as administrators. Only one response fell into this category; it is the same instance represented in the “More than 11 years ago” category above. This instance supports an online community that we confirmed predates the existence of Mastodon. Since the community predates Mastodon, it would stand to reason that many of the governance documents predate Mastodon.

We can compare the last time an administrator read their instance’s privacy policy against their reported familiarity with the policy using the response color-coding of Q9:



Of note here are respondents reporting that they are “Extremely familiar” or “Somewhat familiar” with the privacy policy, despite reporting never having read it.

D.3 Policy Authorship (Q11)

Q11: How much of [your instance]’s privacy policy was written from scratch (by you or someone on your team)? Discussed in Section 5.2; chart appears in Figure 3c:

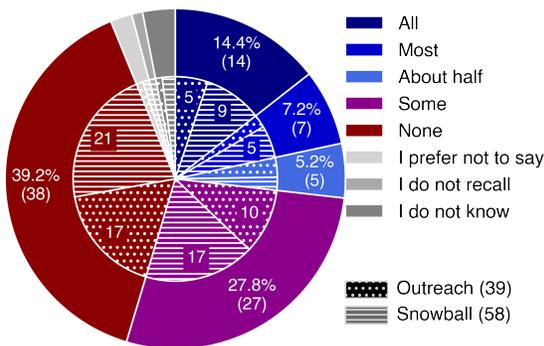


Table 3 aggregates over Q8, Q11, Q13, and Q18. Respondents who answered that they wrote their entire policies from scratch were not asked about tools, templates, or reference policies; all other respondents were asked whether they used tools, templates, or reference policies.

D.4 Technical/Educational Background (Q12)

Q12: To the best of your knowledge, did any of [your instance]’s privacy policy authors have background knowledge or experience in any of the following areas? We discuss this question in detail in Section 5.3.3, aggregating the data in Table 1 (reproduction omitted for readability).

D.5 Tools or Templates (Q13–Q17)

Q13: Did you or someone on your team create [your instance]’s privacy policy using a tool or template service? Discussed in Section 5.2; the following summary chart does not appear in the body of the paper:

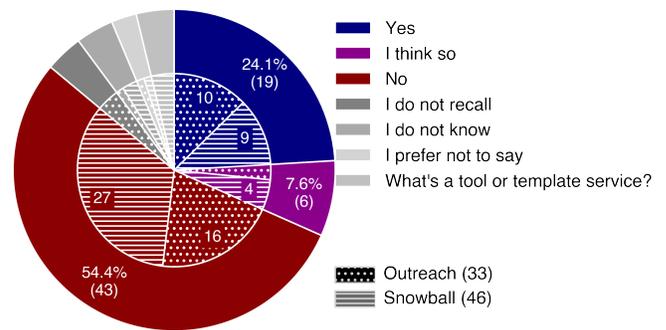
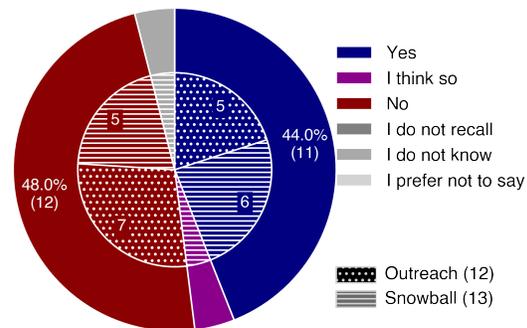
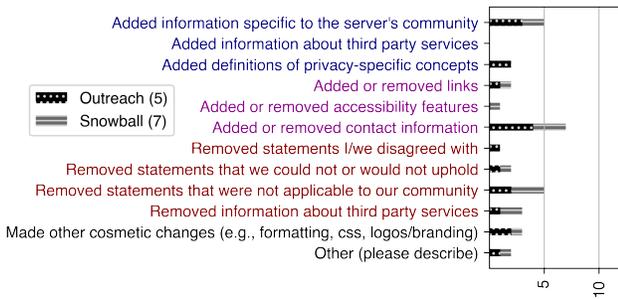


Table 3 aggregates over Q8, Q11, Q13, and Q18.

Q14 Did you or someone on your team make changes to the tool or template output? Discussed in Section 5.2; the following summary chart does not appear in the body of the paper:



Q15: What kinds of changes did you or your team make to the output of the tool? Please check all that apply. Discussed in Section 5.2; responses to this question are aggregated with the responses to Q20 in Figure 4. The following summary chart of only the changes made to the output of a tool or template does not appear in the body of the paper:



The authors devised the options listed here from their observations when manually inspecting privacy policies. Respondents had the opportunity to clarify or suggest other edits that they made. Only two respondents provided additional information:

Option	Count
... were too hard to understand	1
... put too much responsibility on server administrators	2
... were unenforceable	1
... were not compliant with relevant laws	1
... Other Reason (please describe)	1

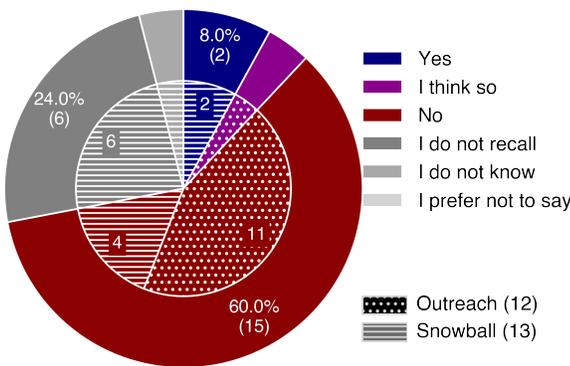
Multiple templates were considered, but we chose the one that worked best for our needs/users.

The first policy was used in 2018. We revised it in 2022 as we revamped the service. Generally, we don't believe there were any major deviations from the source but did some wordsmithing and clarifying.

The authors devised the presented options from their manual inspections of privacy policies and from manual inspection of online tools for generating privacy policies.

Added some specific details concerning how the instance works (how long logs are stored for, for example). the rest is mostly still Mastodon's default privacy policy.

Q16: Were there any tools or templates you or your team considered but did not use? Only two respondents definitively reported having considered tools or templates that they did not ultimately use:



Q17: Please share why you or your team did not use the other tools or templates you'd considered. Alternative(s)... [respondents multi-select from a list of options, including one freetext box]. Only three respondents saw this question. We allude to this question and its responses elsewhere in the paper, but do not directly report or analyze it, due to the low response rate. Respondents selected the following options (freetext response below):

D.6 Reference Policies (Q18–Q23)

Q18: Did you or someone on your team create [your instance]'s privacy policy text by modifying or using the text of one or more existing privacy policies? Discussed in Section 5.2; the following summary chart does not appear in the body of the paper:

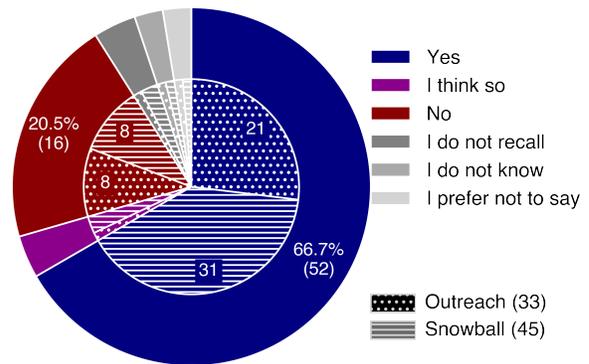


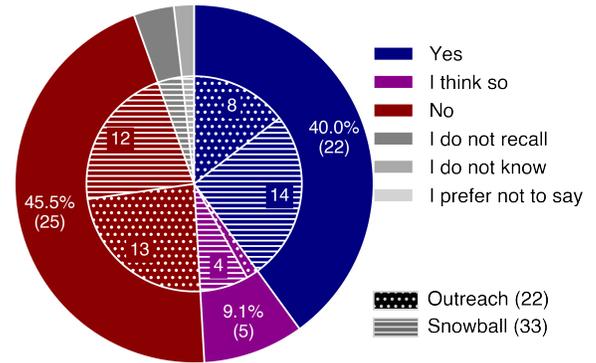
Table 3 aggregates over Q8, Q11, Q13, and Q18.

Q19: Other than the name of the server, did you or someone on your team make changes to the existing privacy policy when crafting [your instance]'s privacy policy? Discussed in Section 5.2; the following summary chart does not appear in the body of the paper:

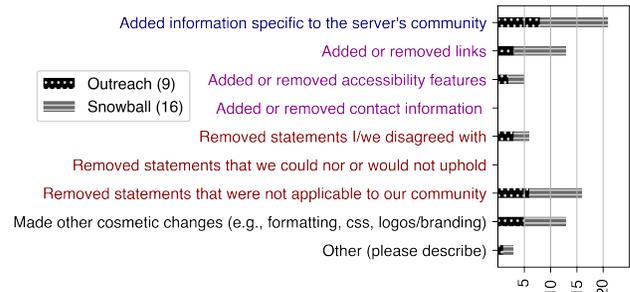
	Involvement (Q8)	Tool or Template (Q13)	Other Policy (Q18)	From Scratch (Q11)	Count
1	Yes	N/A	N/A	All	14
2	Yes	SKIP	SKIP	Some	3
3	Yes	Yes	Yes	≈ Half	3
4	Yes	Yes	Yes	Some	7
5	Yes	Yes	Yes	None	4
6	Yes	Yes	No	None	3
7	Yes	I think so	Yes	Some	1
8	Yes	I think so	I think so	???	1
9	Yes	I think so	No	None	3
10	Yes	No	SKIP	None	1
11	Yes	No	Yes	Most	4
12	Yes	No	Yes	≈ Half	1
13	Yes	No	Yes	Some	13
14	Yes	No	Yes	None	9
15	Yes	No	I think so	≈ Half	1
16	Yes	No	No	Most	1
17	Yes	No	No	Some	1
18	Yes	No	No	None	2
19	Yes	No	???	Most	1
20	Yes	???	Yes	None	2
21	Yes	???	I think so	None	1
22	Yes	??	??	??	1
23	Yes	-	-	-	2
24	Yes	?	Yes	None	2
25	No	SKIP	SKIP	Some	1
26	No	Yes	No	None	1
27	No	I think so	Yes	??	1
28	No	No	Yes	None	3
29	No	No	No	None	3
30	No	No	No	??	1
31	No	No	???	Most	1
32	No	No	???	None	1
33	No	??	Yes	None	1
34	No	??	??	None	1
35	???	Yes	Yes	Some	1
36	???	?	No	None	1

≈ Half ≙ “About half” - ≙ “I prefer not to say”
 ? ≙ “What’s a tool or template service?”
 ?? ≙ “I do not know” ??? ≙ “I do not recall”

Table 3: Cross-tabulation of administrator involvement in and origins of instance privacy policies. Responses to Q8 are color-coded using the legend of Figure 3a. Responses to Q11 are color-coded using the legend of Figure 3c, with the responses to Q13 and Q18 color-coded analogously. Six respondents broke off from the survey after answering whether they were involved in their policy’s genesis; these data are excluded from above.



Q20: Other than the name of server, what kinds of changes did you make to the existing policy? Please check all that apply. Discussed in Section 5.2; the following chart does not appear in the body of the paper:

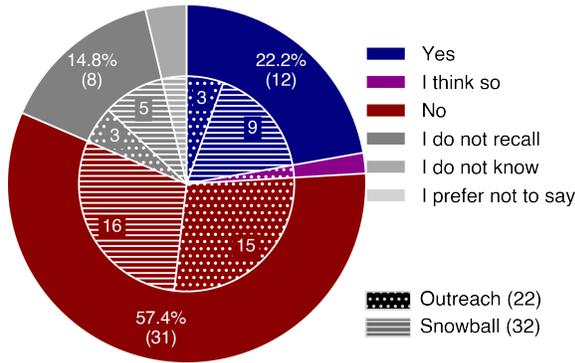


Two respondents provided additional information:

Adapted existing privacy policy to be specific to the Mastodon software’s behavior

Added simple language descriptions of basis for data processing and of what exactly we do with data, for what reason, and with who we share the data, to attempt best effort GDPR compliance

Q21: Were there any existing policies you or your team considered but did not choose to use? Discussed in Section 5.2; the following chart does not appear in the body of the paper:



Q22: Please share why you or your team did not use certain other policies. Alternative(s)...[respondents multi-select from a list of options, including one freetext box]. Discussed in Section 5.2; the following chart does not appear in the body of the paper:

Option	Count
...were unenforcable	2
...put too much responsibility on server administrators	2
...insufficiently addressed privacy concerns	5
...permitted advertising	5
...prohibited advertising	2
...were not compatible with other policy documents	4
...were too hard to understand	5
...were not compliant with relevant laws	9

Three respondents provided additional information:

were not relevant to geographical jurisdiction

Modifying the Mastodon template means typing up a policy. When I do that, which I'll get around to, I want to do it right.

The default Mastodon privacy policy is too vague.

The authors devised the list options of presented to survey respondents from their observations when manually inspecting a sample of privacy policies.

Q23: What is/are the names or urls of the policy/policies on which [your instance]'s privacy policy is based. Recall that 52 respondents answered confidently that their privacy policy was based off a reference policy. Forty-four respondents provided additional information about their reference policy. Twenty-five respondents listed the Mastodon or Hometown¹² default policy, or Discourse (i.e., the privacy policy on which Mastodon's privacy policy is based) as their reference policy. Among this group, two respondents relayed that their policies were based on both the default policy and another policy. In one case, the other policy was Matrix's and for the other, it was the "the organization's primary privacy policy."

One respondent cheekily linked to a specific line of the Discourse's privacy policy, which reads:

¹²Hometown is a popular Mastodon fork

[Do we disclose any information to outside parties?](#disclose)

Another six respondents provided links to other Mastodon instances' About pages, which contain information about codes of conduct, block/allow lists, and other server information unrelated to privacy policies. Upon navigating the publicly available pages on each instance, we found that four of these instances used the default privacy policy and one used a modified default privacy. We could not identify an available privacy policy endpoint for the sixth instance.

Nine respondents listed non-default-derived privacy policies. In most cases, these privacy policies were bundled with terms of service. Two respondents reported that they based their privacy policies off two distinct sources. One respondent reported using the rocketlawyer.com website to generate their documents.

The final four respondents reported:

a lot. we read a lot of other servers policies.

Various templates found via a Google search.

It diesnt exist anymore / the instance is no longer online

don't recall

We did not perform any additional analyses over these responses. None of the nine non-default-derived policies appeared in our initial text analysis of privacy policies. Responses to this question contributed to our hypothesis that some administrators struggle to differentiate different governance documents, discussed in Section 6.

Q24: Is there anything else you'd like to share about the origins of [your instance]'s privacy policy? All respondents saw this question. Of those who chose to reply, most noted that their instances used the default policy and mentioned if they made changes. Several respondents included additional information of interest. One respondent mentioned governance in general, citing a resource previously known to the authors:

Darius Kazemi's "How to run a small social network" <https://runyourown.social/> was very inspirational in setting up a community for ourselves.

Some feedback potentially belied administrator's lack of technical understanding of how Mastodon works. For example, unless the respondent who wrote

it doesn't exist. Our server is informal and in it's initial stages

has written their own software or has overwritten the privacy policy that ships with popular social software, there likely is a privacy policy somewhere. Similarly, statements such as

We never collect or share user data

after often not accurate, since most social software requires at least some retention of IP addresses (i.e., user data) in order provide a usable experience.

Other respondents displayed engagement with both the technical (in terms of software) and legal aspects of privacy policies:

To my knowledge, our record of the processing activities was one of the first and most complete at the time of writing [REDACTED]

The privacy policy, such as it is, essentially states that [REDACTED] is not going to fight the law on behalf of its users. Little else is guaranteed.

We currently use the default Mastodon privacy policy. We have considered modifying it, but have always been scared off by the legal risks that might pose – all the things we might need to account for, might be responsible for. We have been so worried we might miss some things, we haven’t added any at all.

To the best of my knowledge, it was mostly written from scratch by the previous Admin. I made a few minor adjustments when I took over the reigns. There was, at the time, some discussion with the admin of [REDACTED], and I’ve maintained a working relationship with the current [REDACTED] [...] admin.

I largely gave the text of the default Mastodon Privacy Policy to a lawyer for verification that the policy is sufficient and valid under European Law, a few suggestions were made but not implemented, as their impact would have been negligible.

The authors of this paper have heard from several administrators over email throughout this study. Our communications informed our use of the ABCD approach discussed in Section 6. One administrator informed us by email that it was “not possible” to modify the default Mastodon policy, while another wrote in response to Q24:

The Mastodon software ships with a standard privacy policy. Most instance owners and administrators do not modify it.

While it does appear to be true that most administrators do not modify the default policy, there is clearly a sizable subset of administrators who either do modify these policies or want to. One respondent relayed that there exists a super group of regional administrators who actively share knowledge:

there is a group of [REDACTED] MastoAdmins who share this sort of info; big thanks to members of that group for their help.

We have also heard from one administrator who has recently added support for the ATProtocol to their instance, and from another administrator who had realized that they had a rendering error in their privacy policy after having read a draft of this paper. We believe these interactions support the notion that there is substantial heterogeneity in both the background in and desire to engage in the technical and legal aspects of Mastodon administration.

D.7 Privacy Policy Incidents (Q25–Q28)

We ask three questions about privacy incidents:

Q25 Have you or your team ever gotten a request to share information about the content, users, or usage of [your instance]?

Q26 Have you or your team ever needed to invoke [your instance]’s privacy policy in your capacity as an administrator of it?

Q27 Have any members of [your instance] raised privacy issues with you or your team?

Q28 Have any members of [your instance] asked questions regarding your privacy policy?

We mention our findings in Section 5.2 and aggregate in Table 4; this following chart *does not* appear in the body of the text.

	Disclosure Request (Q25)	Privacy Policy Violation (Q26)	Privacy Issues (Q27)	Privacy Policy Questions (Q28)	#
1	???	No	Yes	Yes	1
2	Yes	Yes	Yes	Yes	2
3	Yes	Yes	Yes	No	1
4	Yes	Yes	No	No	1
5	Yes	No	No	Yes	2
6	I think so	Yes	No	No	1
7	No	Yes	No	Yes	3
8	No	Yes	No	No	3
9	No	No	SKIP	SKIP	1
10	No	No	Yes	Yes	1
11	No	No	No	???	4
12	No	No	No	Yes	7
13	No	No	No	No	60

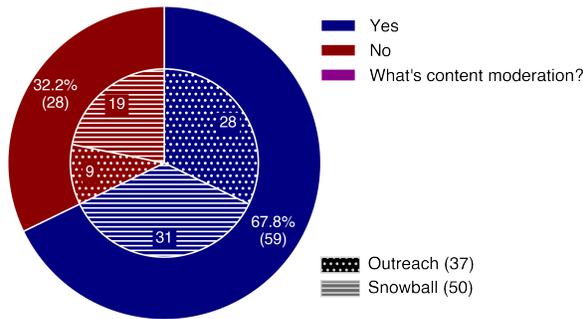
??? ≙ “I do not recall”

Table 4: Eighty-seven respondents answered questions about privacy policy compliance, ranging from requests to share information about users to instance member questions about privacy (Q25–Q28).

We had hypothesized that the degree of specialization or “bespoke-ness” of a policy would be correlated with the severity of legal jeopardy an administrator had previously encountered, i.e., that administrators who had experienced formal requests for user data would have more specialized policies than administrators who had only dealt with “informal” privacy queries, who would in turn have more specialized policies than administrators who never engaged with their privacy policies. We were not able to collect enough data to meaningfully examine this hypothesis; most respondents reported never having encountered situations requiring that they engage with their privacy policies.

D.8 Content Moderation and Data Management Experience (Q29–Q42)

Q29: Do you have any background or prior experience as a content moderator? Discussed in Section 5.3.2; a compressed version of this chart appears in Figure 7.

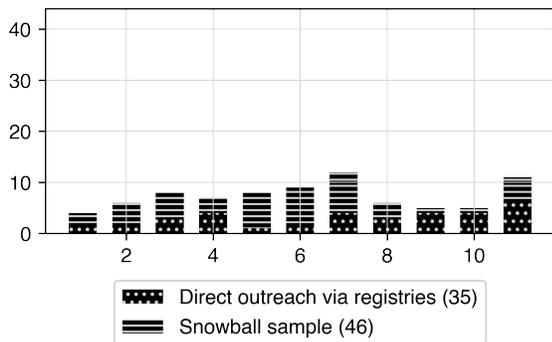


Q30: On what platforms other than Mastodon have you had experience in content management? Please check all that apply. Discussed in Section 5.3.2; word cloud appears in Figure 6 and is not reproduced here.

Q31: Please select the content moderation issues that you have experience with. Do not search for definitions of unfamiliar terms; try to answer this question quickly and confidently without additional resources. We originally designed this question for validation purposes, with the intention of listing a mixture of issues commonly understood to fall under the purview of content moderation (e.g., hate speech) and other issues that are not necessarily content moderation issues (e.g., NFTs). Thus, the original question design was meant to elicit a two-present: first, whether the respondent recognizes the term, and then whether the respondent believes the term describes a content moderation issue.

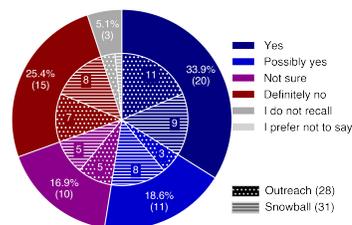
During testing and piloting, we found that respondents assumed our intention was that all of the provided issues were content moderation issues. We still found this question useful for priming respondents, so we did not remove it. However, we do not believe it produced data that was useful for any inferences, so we did not report on it in the body of the paper.

Figure 10 represents both Q31 and Q38. We did not use this data when discussing self-reported familiarity with content moderation or data management (i.e., we only reported on Q29 and Q37 in the body of the paper). We had initially hypothesized that the number of categories respondents reported being familiar with would be correlated with reported content moderation experience.

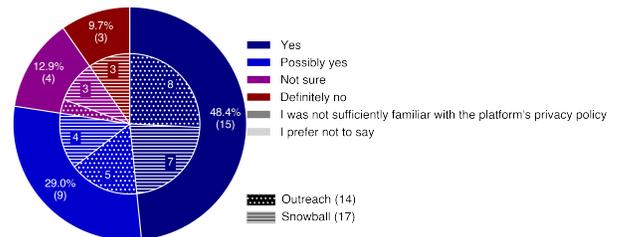


Unfortunately, we do not believe we can justify the battery of terms used as a sufficiently representative sample of content moderation terms to be able to use it to corroborate self-reported experience or expertise in content moderation. We had also considered comparing the distribution of familiar terms for content moderation with the distribution of familiar terms for data management. However, we do not believe these batteries are comparable and explain our reasons for this in our description of Q38.

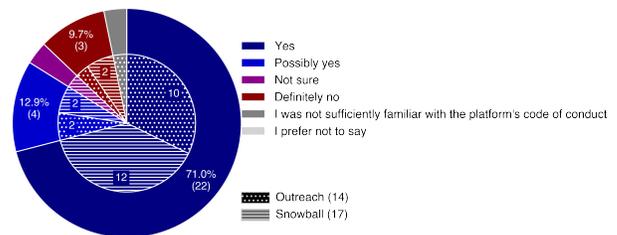
Q32: Have you ever encountered (personally or as a witness) what you would personally consider to be a privacy violation as part of your content moderation experience? We did not discuss this question in the paper, nor report on the following data; following up on this question in interviews is part of an ongoing study:



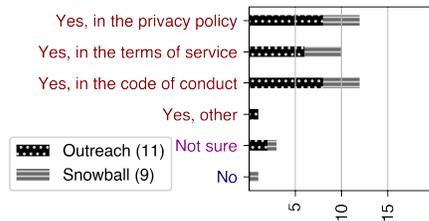
Q33: To the best of your knowledge, was this privacy violation prohibited by the platform's privacy policy? We did not discuss this question in the paper, nor report on the following data; following up on this question in interviews is part of an ongoing study:



Q34: To the best of your knowledge, was this privacy violation prohibited by the platform's code of conduct? We did not discuss this question in the paper, nor report on the following data; following up on this question in interviews is part of an ongoing study:



Q35: Do you believe this privacy violation should have been covered in some way by the platform's governance documents? Select all that apply. We did not discuss this question in the paper, nor report on the following data; following up on this question in interviews is part of an ongoing study:



Q36: Is there any additional specific information you would like to share about your content moderation experience? Twenty-two respondents provided additional commentary. Most emphasized that while they had encountered or were aware of content moderation incidents elsewhere, they had no problems on Mastodon. Five respondents reported operating very small servers, which they felt made moderation easy. Four respondents mentioned using Mastodon tooling (e.g., banning accounts, blocking servers) to handle entities that were unambiguously violating their codes of conduct. Four other respondents discussed content moderation in terms of governance documents, community-building, and communication:

Some of it dates back to the time before ToS / Privacy statements or Codes of Conduct. We mostly made judgement calls, but having pre-agreed rules is absolutely better than that.

mostly I read the book "how to respond to reports of code of conduct violations" by Mary Gardiner and Valerie Aurora

My experience with a hybrid format (online portions of physical events and vice versa) has allowed me to clarify things with people that I might not otherwise meet or converse with in an online-only moderation environment. I believe getting direct feedback is valuable, and I would suggest it for anyone wanting to learn more about content moderation.

I have encouraged the owners of chats and services to make the rules and expectations set with users very clear up front, as I am not a fan of taking disciplinary action based on feelings towards others that do not fall under the rules. This has caused friction with other moderators in the past.

Interestingly, only one respondent wrote about advertising:

Unauthorized advertising of business services and products are the #1 moderation issue. It's not spam because they don't contact anyone. They just set up an electronic billboard on the server.

Finally, two respondents objected to content moderation as being relevant to privacy policies:

Privacy policies as far as I can tell are useless; codes of conduct are mostly signals to build trust; actual enforcement and day to day moderation decisions are far far more impactful

Most content moderation is because of the site rules and terms of service and not the privacy policy

This section is weird to me because I am not sure about the terminology. A "privacy violation" can be some user violating another users privacy, but this is not a privacy violation / incident in the sense of "the privacy policy was violated". The privacy policy applies to us, the website and the people running it, and obliges us / is a promise by us to act in a certain way. It doesnt concern how users should act. I have not witnessed a violation of the privacy policy, either by myself, or by an administrator of another fediverse instance.

As we mentioned in our initial discussion of Q31, we had not intended for respondents to infer that the battery of items were definitively and unambiguously content moderation issues. One respondent raised this very issue:

It is not fair to bucket all of the cryptocurrency space in the same category as bullying, hate speech, etc. Scams, spam, sure. But not all of us in the crypto space are scammers.

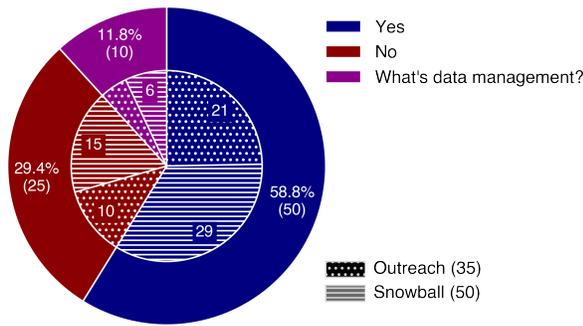
In future iterations of this survey and similar ones, we recommend asking about content moderation history *after* asking about data management experience. We believe that the current ordering had the potential to undermine respondent faith in our competence, due perceived conflation of content moderation and privacy policies. The selection of topics in the battery question (Q31) had the potential to contribute to this perception.

Finally, we note that the following sentiment represents a fundamental tension we have observed in the Fediverse, given the demographics we'd observed:

Content moderation on Mastodon is seen as a major encroachment on freedom of expression by a lot of users.

We do not delve into this issue in the body of the paper because our focus was on privacy policies, not content moderation.

Q37: Do you have any background or prior experience in data management? Discussed in Section 5.3.2; a compressed version of this chart appears in Figure 7.



Q38: Please select the data management issues, concerns, or techniques that you are familiar with. Do not search for definitions of unfamiliar terms; try to answer this question quickly and confidently without additional resources. We had originally designed this question (Q38), along with Q31 for validation purposes. Unlike the battery for Q31, this battery included one item that is unambiguously not a data management issue: CORI, which can refer to both a law and a database for Criminal Offender Record Information in the Commonwealth of Massachusetts in the United States.

The data management battery contains legal acronyms, technologies, and security concepts. It was our intention to differentiate privacy concepts from security concepts, but like our content moderation battery, we found during testing that respondents took a broader view of what we were asking and appeared to assume that we considered most or all of the terms germane to the topic.

Again, we can plot a histogram of the counts of distinct (legitimate) issues with which each respondent reports having familiarity:

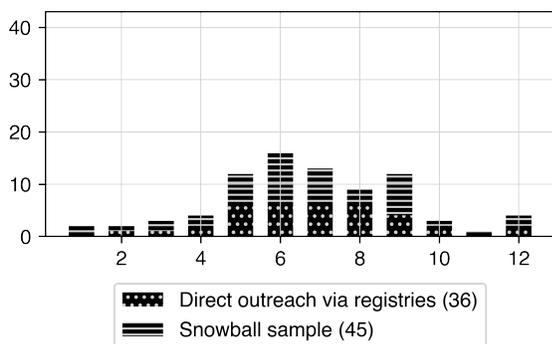


Figure 10 plots the total counts of respondents who reported familiarity with content moderation and data management concepts by battery item. One major challenge when comparing between the two groups of terms is that the content moderation terms are more widely used in common parlance. The data management terms are more technical and legal and may require more specialized knowledge to understand. While one could argue that this represents a fundamental facet of content management vs. data management, we did not feel sufficiently confident in our results to report them in the body of the paper. Groß [36] recently raised concerns about construct validity in widely known and used survey instruments.

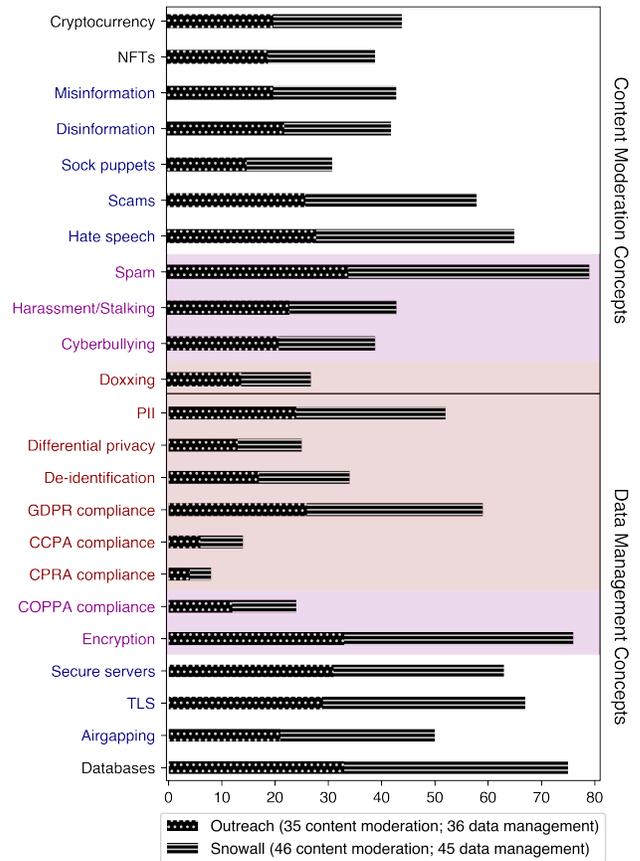
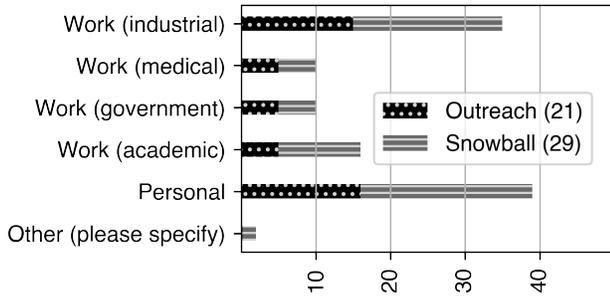


Figure 10: Counts of content moderation and data management issues respondents reported recognizing. Items and chart regions are color-coded by whether we consider them security concepts, privacy concepts, both, or neither. The center items (CPRA compliance, CCPA compliance, GDPR compliance, De-identification, Differential privacy, PII, Doxxing) are color-coded as privacy concepts. The adjacent pink bands (Encryption, COPPA compliance, Cyberbullying, Harassment/Stalking, Spam) are color-coded as both security and privacy concepts. The blue text items (Misinformation, Disinformation, Sock puppets, Cryptocurrency, NFTs, Scams, Hate speech, Airgapping, TLS, Secure servers) are security concepts. The remaining concepts (Databases, NFTs, Cryptocurrency) are neither.

Therefore, we feel this section of the survey deserves additional analysis and scrutiny before it can be used to infer meaningful findings.

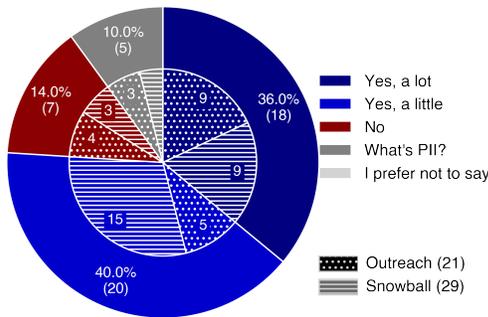
In future studies, we recommend developing battery of items supported by existing literature and user studies in order to validate self-reported confidence of respondents' familiarity with the content moderation and data management. All we can say at present is that our findings do not contradict respondents' self-reporting.

Q39: *In what contexts other than Mastodon have you had experience in data management? Please check all that apply.* We did not discuss this question in the body of the paper.

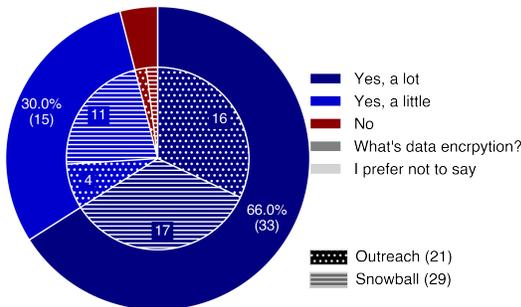


Two respondents provided additional contexts in which they had experience in data management: “Work (non-profit)” and “Work clergy.”

Q40: *Do you have a background in or prior experience with collecting PII?* This question was not discussed in the body of the paper:



Q41: *Do you have a background in or prior experience with data encryption?* This question was not discussed in the body of the paper:



Q42: *Is there any additional specific information you would like to share about your data management experience?* Several respondents provided additional information on their professional backgrounds

(we omit all quotations referring to current employment to protect respondents’ privacy):

I worked at a medical device company for a while and thought a lot about cyber security there, including building secure HIPAA compliant systems

Professionally, I worked at a medtech company for several years, where PII (specifically, PHI) management was existentially important for the business...

Been CTO of several large companies

I have worked with both PII and data encryption at a tech company that is a regular target of state-sponsored attacks

I am an infosec engineer by trade, so this is very familiar territory.

Some respondents also provided additional information about the other privacy-related technologies they use:

we also run matrix for e2e chats in our community

As a hobbyist, I used to be very into the PGP web of trust, and to this day, I use PGP encrypted email.

D.9 Demographics and Identity (Q43–Q50)

Q43–Q44 (Minoritized/Maringalized identities). We discuss these questions at a high level in Section 5.3.4 and Figure 8. Due to space considerations, we could not discuss multiply-marginalized identities in the body of the paper.

Some intersections may have straightforward explanations: gender and sexuality are deeply enmeshed categories, such that transgender and gender-variant people also often do not identify as heterosexual. Disability is also highly implicated with LGBTQ+ identity due to discrimination and poor access to healthcare [92]. Crenshaw [20] attests to the way multiple axes of marginalization affect each other, compounding on each other in ways that often exacerbate the harm that structural inequity causes. Accordingly, identity-based marginalization (like LGBTQ+ identity, race, or ethnicity) means increased vulnerability to mental/physical disability and poverty, which may be another reason for the high percentage of multiply marginalized respondents.

Q45: *There are specific demographic questions we ask when studying populations in the United States. Please indicate whether any of the following are true...* This question was used solely for determining US-anchored demographic characteristics.

Respondents who met the conditions of our US-anchored battery item question (Q45) were asked to identify themselves in terms of US Census Bureau-defined racial and ethnic categories. We also included questions about gender identity in this section (questions Q45–Q49).

Q46–Q49: (US-anchored questions about racial, ethnic, and gender identity.) Forty respondents gave racial demographic information. Thirty-seven identified as “White” only. The remaining three each identified as “Asian” only, “Other” (no additional text provided) only, and both “White” and “Native Hawaiian and Other Pacific

Islander.” Forty-five respondents gave Latin American ethnicity information; two responded “Yes” (one from each group), three responded “Prefer not to say” (two from the outreach group, one from the snowball group), and the rest responded “No.”

Forty-six respondents gave information about their gender identity. Five selected “Prefer not to say” for either their gender identity or its alignment with their gender assignment at birth. Of the 39 respondents who provided information, 33 identified as cisgender men. The remaining six respondents included one female transgender administrator, one female cisgender administrator, two female-leaning genderfluid administrators, one genderfluid non-binary administrator, and one genderfluid agender administrator. Despite the several administrators reporting varying identification with their genders assigned at birth (e.g., “Sometimes,” “Most of the time”), no one reported multiple genders.

Q50. Finally, we asked all respondents whether they would like to share anything else about their identities. Most responses were captured by other data. We will not reproduce the content of these freetext responses, since some of them may be identifying and include personal information. We are grateful to the administrators for providing this information, since it does help us contextualize who administrators are and how their identities inform their views on community and privacy.

However, we do highlight one response that could be interpreted in a variety of ways:

Previous question - I think I'd ban any Americans who use my mastodon server.

We do not know if this response reflects suspicion toward demographic data collection or if it reflects concerns about institutional American privacy violations. Regardless, it is worth noting that this sentiment exists and presents potential challenges when building trust with the Fediverse community.