



PDF Download
3567568.pdf
01 April 2026
Total Citations: 13
Total Downloads: 943

Latest updates: <https://dl.acm.org/doi/10.1145/3567568>

Published: 29 December 2022

RESEARCH-ARTICLE

Who Moderates on Twitch and What Do They Do?: Quantifying Practices in Community Moderation on Twitch

[Citation in BibTeX format](#)

JOSEPH SEERING, Stanford University, Stanford, CA, United States

SANJAY R. KAIRAM

Open Access Support provided by:

Stanford University

Who Moderates on Twitch and What Do They Do? Quantifying Practices in Community Moderation on Twitch

JOSEPH SEERING, Stanford University, USA

SANJAY R. KAIRAM, Reddit, Inc.*, USA

Volunteer moderators are an increasingly essential component of effective community management across a range of services, such as Facebook, Reddit, Discord, YouTube, and Twitch. Prior work has investigated how users of these services become moderators, their attitudes towards community moderation, and the work that they perform, largely through interviews with community moderators and managers. In this paper, we analyze survey data from a large, representative sample of 1,053 adults in the United States who are active Twitch moderators. Our findings – examining moderator recruitment, motivations, tasks, and roles – validate observations from prior qualitative work on Twitch moderation, showing not only how they generalize across a wider population of livestreaming contexts, but also how they vary. For example, while moderators in larger channels are more likely to have been chosen because they were regular, active participants, mods in smaller channels are more likely to have had a pre-existing connection with the streamer. We similarly find that channel size predicts differences in how new moderators are onboarded and their motivations for becoming moderators. Finally, we find that moderators’ self-perceived roles map to differences in the patterns of conversation, socialization, enforcement, and other tasks that they perform. We discuss these results, how they relate to prior work on community moderation across services, and applications to research and design in volunteer moderation.

CCS Concepts: • **Human-centered computing** → **Collaborative and social computing**; **Empirical studies in HCI**.

Additional Key Words and Phrases: content moderation, volunteer moderators, Twitch

ACM Reference Format:

Joseph Seering and Sanjay R. Kairam. 2023. Who Moderates on Twitch and What Do They Do? Quantifying Practices in Community Moderation on Twitch. *Proc. ACM Hum.-Comput. Interact.* 7, GROUP, Article 18 (January 2023), 18 pages. <https://doi.org/10.1145/3567568>

1 INTRODUCTION

Many major social media services rely at least in part on users to moderate content. Facebook and Reddit, for example, have both relied on users to moderate groups for more than a decade, and even networked services like Twitter and YouTube have been actively increasing the role that volunteers play in moderation in recent years.^{1,2}

A significant body of research in the social computing space has explored practices in volunteer community moderation across multiple services, identifying patterns with respect to topics like

¹<https://support.google.com/youtube/answer/9826490?hl=en>

²<https://help.twitter.com/en/using-twitter/communities#what-is-mod>

* This author completed this work in its entirety while employed at Twitch/Amazon.

Authors’ addresses: **Joseph Seering**, seeringj@stanford.edu, Stanford University, Stanford, CA, USA, 94305; **Sanjay R. Kairam**, sanjay.kairam@gmail.com, Reddit, Inc.*, San Francisco, CA, USA, 94104.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.

2573-0142/2023/1-ART18 \$15.00

<https://doi.org/10.1145/3567568>

how moderators are recruited, how they learn how to moderate, their philosophies for moderating, and how they interact with internal trust and safety employees (e.g. [3, 8, 24, 26, 28, 30, 36]). These insights have been developed primarily through interview studies using convenience sampling; it is not yet known how these findings generalize across broader moderator populations.

In this paper, we address this gap on one platform – Twitch – by analyzing and presenting results from a large, representative survey developed around frameworks from prior work for describing processes [30] and social roles [36] in volunteer moderation respectively. The questionnaire captures a variety of aspects of volunteer moderation, broadly summarized according to these three primary research questions:

- (1) How are individuals commonly recruited and onboarded as moderators within Twitch communities?
- (2) What motivates these individuals to become moderators?
- (3) What is the relationship between Twitch moderators' self-described roles and the tasks they perform within their communities?

We deployed this questionnaire to a sample of adults located in the United States who are active Twitch moderators, gathering 1,053 responses capturing moderator demographics, recruitment contexts, motivations, common tasks, and self-identified roles within channels. Our findings broadly support many of the assertions made in small-scale, qualitative studies of volunteer moderation on Twitch, demonstrating that these generalize across the adult US-based moderator population.

By incorporating aggregate audience data about the channels which participants moderated, our findings also illustrate patterns in how moderator motivations, practices, and tasks can differ across contexts, including communities of varying size and moderators in differentiated roles. We highlight some of these findings here:

- Moderators in large channels are most likely to have been chosen because they were regular, active participants in the channel, while those in smaller channels more likely had a prior connection with the streamer.
- The moderators most likely to rely on in-product information or first-party resources when learning how to moderate are those in smaller channels.
- Moderators in smaller streams are more likely to refer to 'friends' and feeling 'close' when asked why they started moderating, compared with those in larger streams, who more typically referred to 'community', 'viewers', or 'content'.
- The tasks that moderators perform vary with channel size, with tasks related to enforcement and coordination becoming more common as channels grow.
- Moderators' self-perceived roles within the channel map to systematic differences in the tasks performed, raising opportunities for the design of role-specific tooling.

By adapting frameworks developed in prior work and evaluating them quantitatively, this work serves as an example of how different methods can build on each other to provide complementary insights into topics in social computing. By drawing, in part, on prior work identifying moderation practices that extend across services, we can more clearly build connections to volunteer moderation contexts outside of Twitch. We conclude by discussing how these findings can be used to inform the design of moderation tools and features, both on Twitch and more broadly across services supporting online communities.

2 PRIOR WORK

Volunteer moderators occupy an important place in the moderation ecosystem for a number of large social media services. According to the Twitch H1 2021 Transparency Report, the fraction of live minutes watched in channels moderated by volunteer moderators and/or user-created moderation

bots in the first half of 2021 was 87.8%, meaning that the vast majority of Twitch users participate primarily in spaces that are moderated in part by other users.³ During this same time period, more than 37 million messages were manually removed by users or user-created moderation bots. These users' impact, however, extends far beyond the volume of messages removed. Their work provides a valuable complement to the company's internal moderation processes by proactively cultivating positive communities.

Since the beginning of the social internet, myriad online social spaces have been moderated by volunteer community members. In the late 1970s and early 1980s, these moderators were frequently academics or industry computer enthusiasts [13, 18], but through the mid 1980s and the 1990s, the internet began to diversify and so did its moderators [35]. In the 1990s, a variety of ethnographic work studied the processes of moderation and governance on Usenet and Multi-User Dungeons, among other services, identifying the complex relationships between communities and their moderators and how moderation processes can shape social interactions [27, 32, 33]. Many of these relationships mirror those found today in communities on spaces like Twitch.

While "community-reliant" approaches, per Caplan [5], have received less public attention in recent years as compared with larger, more centralized models (see e.g., [15, p. 124–128]), they remain prevalent and have been the subject of much recent research in social computing that has emphasized the social depth of volunteer moderators' roles across Twitch, Reddit, Facebook groups, and Discord, among others. [6, 19, 20, 22, 23, 26, 28, 30, 37]. Several recent papers have explored volunteer moderation practices specifically on Twitch, analyzing topics ranging from social influence [29], large-scale event moderation [24], community-focused moderation practices [3, 36], and approaches to handling harassment [25, 34].

In this note, we aim to complement findings from prior qualitative work on volunteer moderation through a survey of a larger sample. In designing our survey and discussing our findings, we drew primarily from a set of frameworks developed within the past three years for describing volunteer moderation processes and practices. Our survey design draws significantly from Seering et al.'s work describing volunteer moderators from Twitch, Facebook Groups, and Reddit [30], and we draw our four categories of moderator roles from Wohn's framework for Twitch moderator social roles [36]. We used these papers both because they document a breadth of approaches and processes and because they have complementary perspectives — one studies moderation on Twitch specifically and the other studies moderation broadly, tying Twitch moderation to moderation on other services. We aim to use this work to complement these papers' findings, providing evidence of the prevalence of the processes and practices they document at scale. In particular, we focus on quantifying the prevalence of major aspects of moderation processes identified in prior work such as motivations for moderating, moderator onboarding and learning, tasks and actions performed by moderators, and moderators' social roles within their communities.

3 METHODS

We distributed a survey in December 2020, which was completed by 1,053 logged-in Twitch users located in the United States who had moderated at least one channel within the preceding 28 days. We grouped candidates into four segments based on whether they (1) had moderation activity in a single channel or multiple channels, and whether they had (2) had or had not used Twitch's 'Mod View' feature, both over the preceding 28 days. A pilot wave of 1,000 survey contacts was used to estimate response rates within each segment, allowing us to balance our candidate pool.

The survey was distributed via email to 28,850 members of this population who had opted-in to communications from Twitch and kept open until over 1,000 completed responses had been received.

³https://safety.twitch.tv/s/article/Transparency-Reports?language=en_US#2H12021TransparencyReport

Self-Reported Gender	% Respondents
Man	71%
Woman	25%
Non-Binary / Third Gender	3%
Prefer to Self-Describe	1%
Prefer not to Disclose	2%

Table 1. Respondent self-reported gender. Respondents could select more than one category. These statistics roughly match demographics for adult Twitch users in the United States captured in prior work [31]

The first question in the survey asked participants to report their age; respondents indicating that they were below 18 were directed to a page indicating that they were ineligible. Of the 2,051 individuals who initiated the survey (7.1% attempt rate), we received 1,053 completed responses (3.6% completion rate). The fraction of survey respondents in each of the four segments closely matched those in our larger candidate pool, based on behavioral logs, indicating that the sample is representative of the overall adult, US-based, Twitch moderator population. Respondents received a \$10 Amazon gift card as compensation for completing the survey.

Respondents were prompted to provide the name of one channel they had moderated within the past month. 918 respondents (87%) provided a string matching a channel that had been actively streaming as of November 2020. Usage logs describing aggregate behavior in these channels were used to segment our findings by channel size below. Prior work has demonstrated that changes in the social experience within channels on Twitch vary logarithmically with audience size (e.g. [8, 12, 21, 31]); we thus use the following categories to group channels by size: (**T0**: 0-2 CCU⁴; **T1**: 3-10 CCU; **T2**: 11-100 CCU; **T3+**: 100+ CCU).

3.1 Survey Participants

Respondents' ages ranged from 18 to 63 ($M = 27.3$, $SD = 7.5$), with 85% indicating that they were under the age of 35. Respondents self-reported their gender by choosing one or more option from the list presented in Table 1, with 1.7% selecting more than one option. 71% of our respondents self-identified as men, 25% as women, and 4% chose a non-binary option or preferred to self-describe.

Though official demographics for the Twitch user population are not available, these statistics roughly match demographics for adult Twitch users in the United States as reported in prior surveys distributed directly to Twitch users (e.g. [31]). Although both the present study and [31] excluded individuals under the age of 18 from participation, the minimum age required for a Twitch account is 13, meaning that the population of Twitch users aged 13 to 18 is not captured here. The survey also captured self-reported information about race/ethnicity and disability status; while these are not used as part of our analysis in our paper, we have included a summary in the Appendix for readers to reference.

As self-reported, 462 (44%) respondents had moderated exactly one channel within the prior 28 days, 307 (29%) exactly two channels, and the remaining 284 (27%) three or more channels. In reference to their self-reported recently-moderated channel, respondents indicated how long ago it had been since they first became a moderator on that channel using multiple choice. The median tenure on the channel was between 3 and 6 months, with 124 (12%) having started within the past month and 288 (27%) having moderated on the channel for at least a year. 119 (11%) indicated that they were the only active moderator on the channel, 219 (21%) indicated that the channel had two active mods, and the remaining 715 (68%) indicated working in teams of 3 or more.

⁴Concurrent Users (CCU): defined as the average number of logged-in or anonymous viewers present in a typical minute that a channel is streaming. This number is visible next to the stream title when the streamer is broadcasting.

	Overall	Channel Tier			
		T0	T1	T2	T3+
I was a regular participant within this channel on Twitch.**	69%	52%	66%	76%	88%
I had a personal connection with the streamer from outside of Twitch.**	50%	65%	59%	39%	26%
The streamer reached out to me to ask if I wanted to be a moderator.	46%	38%	47%	53%	38%
I already had a reputation as a good moderator in other channels.*	14%	7%	13%	18%	19%
I was already a moderator for a closely-related channel on Twitch.	12%	11%	12%	13%	12%
I was suddenly appointed as a moderator without being asked.*	12%	12%	9%	11%	21%
I was already a moderator for this community on a service other than Twitch (e.g. Discord).	11%	8%	12%	12%	9%
I reached out to the streamer and asked if I could be a moderator.	9%	6%	9%	10%	11%
I was a streamer who agreed to “mutually moderate” channels with another streamer.**	8%	12%	9%	6%	3%
The streamer was looking for additional moderators to cover more time zones.**	8%	3%	3%	14%	15%
A mod on this channel reached out to me to ask if I wanted to be a moderator.	5%	3%	4%	6%	7%
The channel had an open call for new moderators, and I volunteered.	5%	3%	4%	9%	4%
I was “fast tracked” as a mod due to a sudden increase in viewership.	5%	3%	3%	8%	7%

Table 2. Paths for becoming a moderator separated out by channel size category. Significance testing was performed via logistic regression against CCU, indicating that the odds of a moderator selecting this option in reference to a channel varied significantly with the size of the channel’s audience (* [dark gray] indicates $p < 0.01$; ** [light gray] indicates $p < 0.001$).

4 HOW DO INDIVIDUALS BECOME MODERATORS?

In this section, we report findings about the common contexts in which users become moderators and the sources of information on which they draw when learning how to moderate.

4.1 Moderator recruitment

After providing the name of a recently-moderated channel, participants were prompted with: “Which of the following statements apply to your relationship with the channel or streamer prior to your becoming a moderator on the channel? Please select all that apply” Respondents used check boxes to indicate agreement with a set of 14 options (including an ‘Other (please specify)’ option).

Items were drawn from common recruitment contexts identified in prior qualitative work on community moderation. These included having been a long-term participant in the community (e.g. [8, 24, 26, 30]), having an existing personal connection or friendship with a streamer or moderator (e.g. [26, 30]), having waited for a request or responded to an open call for moderators (e.g. [26, 36]), or extending moderation responsibilities to a new service or context (e.g. [22]).

Table 2 shows results for the 918 respondents who provided a valid channel name. We find that all of these previously identified pathways to becoming a moderator exist on Twitch and that their relative prevalence was in line with what had been asserted in prior work.

	Overall	Channel Tier			
		T0	T1	T2	T3+
Advice or guidance from streamers	76%	68%	79%	79%	77%
Advice or guidance from other moderators **	69%	63%	66%	74%	70%
In-product information or tutorials (e.g. Mod View onboarding) *	49%	48%	54%	48%	41%
Twitch-provided resource pages or help articles	46%	49%	49%	41%	42%
Resources and help articles found on other sites (e.g. Reddit, Medium)	42%	41%	44%	43%	38%
Experience from moderating on other sites or services (outside of Twitch)	37%	39%	37%	36%	32%

Table 3. Percentage of respondents indicating that a source was either “important” or “very important” in learning how to moderate effectively. Significance testing performed using ordinal regression on the 5-point Likert response against $\log_{10}CCU$ (* indicates $p < 0.01$; ** indicates $p < 0.001$)

For each recruitment context, we analyze the relationship with channel size using a logistic regression model, which predicts the likelihood of a particular statement being true given a single predictor variable: the size of the channel ($\log_{10}CCU$); we indicate in Table 2 recruitment contexts for which the coefficient estimate for this predictor is significantly different from zero. We address below some themes which emerged from our findings.

4.1.1 Relationships with channel size. Moderators in larger channels (T3+: 88%; T2: 76%) are more likely to be chosen from the set of regular participants within the channel ($\beta = 0.74, p < 0.001$). This is still common, though less so, in channels of smaller size (T1: 66%, T0: 52%). Moderators in smaller channels (T0: 65%; T1: 59%) are more commonly recruited through prior relationships with the streamer ($\beta = -0.78, p < 0.001$), compared with those in larger channels (T2: 39%, T3+: 26%).

While less common overall, mods in smaller channels (T0: 12%, T1: 9%) are much more likely to be streamers themselves, engaged in “mutual moderation” agreements with other streamers ($\beta = -0.73, p < 0.001$), compared with mods in larger channels (T2: 6%, T3+: 3%). In an opposing trend, it’s much more likely that mods in larger channels (T3+: 15%; T2: 14%) are recruited to help cover multiple time zones ($\beta = 0.69, p < 0.001$) than those in smaller channels (T1: 3%, T0: 3%).

4.1.2 Directionality of requests. Moderators in channels of all sizes were much more likely to report that streamers had reached out to them to ask if they wanted to moderate the channel (44% overall) than that they had reached out to the streamer first (9% overall) or had volunteered in an open call for mods (5% overall). In interviews with Twitch moderators, Wohn [36] identified a potential ‘stigma’ associated with asking to be a moderator, which aligns with our quantitative findings. These findings also illustrate how recruitment can vary across services; Matias [26] found that formal application processes and mod ‘resumes’ are commonplace on Reddit, showing how service-specific cultural differences may shape who becomes a moderator and how.

4.2 Moderator Onboarding

Prior qualitative work has found that moderators’ primary avenues for learning about how to moderate are observing others’ moderation practices, learning by doing, and discussing how to moderate with other moderators [8, 30]. In interviews, moderators were less likely to say that they learned about how to moderate from on-service informational materials [30]. In the survey, moderators responded to the question “When learning how to be an effective moderator on Twitch, how important were each of the following types of information sources to you personally?” using 5-point Likert scales for several categories of information sources(ranging from ‘Not at all’ to

Overall	Top Terms by Tier of Channel Moderated			
	T0	T1	T2	T3+
friend	friend	friend	streamer	commun*
streamer	channel	streamer	commun*	moder*
channel	streamer	channel	channel	channel
moder*	stream	stream	friend	mod
commun*	mod	moder*	moder*	stream
stream	moder*	commun*	stream	streamer
mod	chat	mod	mod	chat
chat	help	chat	enjoy*	enjoy*
enjoy*	safe	enjoy*	chat	help
help*	peopl*	twitch	help	twitch
peopl*	commun*	help	love	viewer
twitch	person	person	peopl*	love
person	support	peopl*	viewer	peopl*
love	watch	support	twitch	content
support	close	grow	content	lot

Table 4. Most frequently occurring tokens (stemmed) in explanations of why Twitch users choose to moderate.

‘Extremely’ important). We summarize the findings for 918 respondents who provided a valid channel name in Table 3 below:

The responses show that when moderators receive structured information about how to moderate effectively, it is largely from streamers and other mods. 76% report that advice from streamers was “very/extremely important” and 68% point to advice or guidance from other moderators. Twitch-provided resources hold the next level of priority, with 49% of moderators stating that in-product info/tutorials, such as those in Mod View, were “very/extremely important” and 46% pointing to Twitch help articles. Finally, a smaller, but substantial fraction of mods indicate that third-party resources had helped them, with 42% referring to external resources and articles and 38% to experiencing moderating on other sites or services as very/extremely important to learning how to moderate on Twitch.

To evaluate how the importance of these categories of information varied with channel size, we performed ordinal regression to predict 5-point Likert responses using a single variable: $\log_{10}CCU$. We find that moderators in larger channels (T3+ 70%; T2: 74%) are more likely to learn from other mods ($\beta = 0.29, p < 0.001$) than those in smaller channels (T1: 66%, T0: 63%). In contrast, we find that mods in smaller channels (T0: 48%; T1: 54%) are more likely to learn from in-product information or tutorials ($\beta = -0.21, p < 0.01$), compared with mods in larger channels (T2: 48%, T3+ 41%). These results support findings from prior work that point to more informal social learning as the major avenue for moderator education, but they show that other more formal resources can also be important, particularly for moderators in smaller channels where peer support may be less readily available.

5 WHY DO INDIVIDUALS BECOME MODERATORS?

For the self-selected channel in which they had actively moderated in the past month, survey respondents were invited to explain in free-text what originally motivated them to start moderating on that channel. We present our findings below based on the 1,022 (97%) valid responses that we received to this prompt.

5.1 Motivations and channel size

In Table 4, we show the most frequently used terms appearing in these responses after tokenizing and stemming⁵ using the SnowballC package in R. The most frequently used tokens overall include **friend** (344 mentions), **streamer** (308 mentions), **channel** (260), **moder*** (226), **commun*** (222), **stream** (208), **mod** (169), **chat** (142), **enjoy*** (120), and **help** (105).

We note some clear patterns when we split responses by channel tier, for the subset of cases where moderators provided a valid channel name (N = 918). The token **friend** is the most commonly mentioned token among motivations of moderators of smaller streams, referenced by 45% of T0 moderators and 38% of T1 moderators, but is much less commonly referenced by moderators of larger streams (T2: 28%; T3+: 11%). Similarly, the term **close** (T0: 4%; T1: 5%; T2: 3%; T3+: 1%) is much more frequently referenced in motivations for moderators of smaller streams.

In contrast, as channels increase in size, moderators are much more likely to refer to the community (**commun***: T0: 6%; T1: 14%; T2: 25%; T3+: 32%). Finally, in the largest channels (T2/T3+), we see moderators starting to shift their language to talk more about **viewers** (T0: 2%; T1: 5%; T2: 8%; T3+: 13%) and **content** (T0: 4%; T1: 5%; T2: 6%; T3+: 11%).

5.2 Common themes in moderator motivations

We performed a high-level content analysis, aided by a Latent Dirichlet Allocation (LDA)-based topic model which identified seven common topics among the motivations described by moderators. LDA-based topic modeling is a statistical approach used to identify different “topics” by looking at relative frequencies of words across different “documents” – in this case, responses to open-ended questions. We selected the number of topics using the `ldatuning` package in R, finding that seven topics provided the best balance across multiple evaluation metrics [1, 4, 9, 17]. Based on an initial evaluation by the authors, four of the identified topics covered general or universal aspects of moderation, while three captured more distinct themes (see Appendix C for more detail). We summarize these three topics below, along with some anonymized representative examples:

5.2.1 Helping a Friend. One topic captured a motivation to support a close friend in their streaming goals. The terms most strongly associated with this topic include **friend**, **grow**, **close**, **mine**, and **continu***. Some examples which loaded strongly on this topic are presented below:

- *<streamer> is a close friend of mine and has performed as moderator on my channel in an exemplary manner. I want to return that kindness to him. (T1)*
- *He’s a close friend of mine and I want to help his stream grow on a healthy way (T1)*
- *This channel is run by a very close friend of mine, and we’re wanting to become popular streamers together. (T0)*
- *Streamer is a close friend of mine before her streaming career started. I wanted to help keep her same, and her channel enjoyable for everyone to come to. (T2)*

These examples highlight the role that personal connections can play in the decision to become a moderator. In addition, these illustrate some examples of ‘mutual moderation’ agreements, in which streamers help each other by moderating each others’ channels. While some individuals who gain moderation status through close friendships with the streamers have previously been characterized as ‘token mods’ [36], all of the participants in our survey had recently taken moderation actions within at least one channel, indicating that some close friendships can produce moderators who are actively engaged in managing the community.

⁵Word stems are indicated throughout using an asterisk (e.g. {indicate, indicating, indicated} → **indicat***).

5.2.2 *Supporting the community.* This topic captured a motivation to foster a positive and active community. The terms most strongly associated with this topic include **commun***, **twitch**, **love**, **person**, **posit***. Some examples which loaded strongly on this topic are presented below:

- *<streamer> has been a close friend for manys years, and has a good sense for building a diverse and warm community, both on Twitch, and in other medias. Keeping her channel a safe and welcoming place is extremely rewarding and important to me, and i derive much enjoyment out of being a member of her community. (T1)*
- *Being more involved in the community and helping foster a positive environment. (T1)*
- *As I just love the twitch experience in general, I'll keep moderating a channel as long as that place is full of friends and positive vibes [+] (T2)*
- *I appreciate the content she creates and the positive effect she has on her community. <streamer> is an amazing person and a great content creator and I am happy to assist her in continuing that path in any way possible and moderating for her Twitch, Discord, and in person events during my spare time is the easiest way for me to contribute to her success. (T3+)*

These examples illustrate that the desire to moderate can be a response to experiencing a positive sense of community, both within the channel and across the service, and a desire to sustain that environment for others. As we can see from the first example quote, many of the responses addressed more than one concern, such as motivations to support both a friend and a 'diverse and warm community'. This category of motivation aligns with the positive community orientation described by Wohn's 'Conversationalists' [36]. These may be moderators who would identify with the 'Facilitators' identified by Seering et al. [28], who act as 'hosts' to encourage socialization and conversation.

5.2.3 *Creating a safe space.* The third distinctive topic captured a desire to create a safe and welcoming environment for others. Top terms associated with this topic include **streamer**, **safe**, **support**, **feel**, **fun**, and some representative examples are shown below:

- *To ensure the streamer on the channel i cover continues to provide a welcome, fun and inviting place for everyone willing to follow rules of basic decency (T3+)*
- *I like helping out streamers build a safe community on Twitch. It's important to provide safe spaces for others while keeping the streamer safe and with the time to focus on their content. (T0)*
- *There's a lot of stuff that should be moderated in a chat room especially big streamers so I do my job making people feel safe and even include commands for my own page for everyone to feel safe. (T1)*
- *Building a community is incredibly important. It gives a place for people to come when feeling rough. Also the internet is terrible and constantly tries to bring people down. As a moderator it's your job to help foster the community you want. (Tier not available)*

These examples reference how creating a safe space requires active engagement as a moderator, including through enforcement actions, such as message deletions and bans. These motivation statements are evocative of Wohn's 'Justice Enforcers,' who approach moderation as a form of service to the community [36]. These statements also evoke aspects of the 'Protector' (creating a safe space) and 'Gardener' (planting seeds and pulling weeds) metaphors for moderation work outlined by Seering et al. [28] and Yu et al. [37].

6 WHAT WORK DO MODERATORS PERFORM?

In this section, we report findings about the tasks that moderators regularly perform and how these align with the various roles that moderators play within a channel.

	Overall	Channel Tier			
		T0	T1	T2	T3+
Welcoming and socializing newcomers to the community	75%	63%	77%	81%	72%
Contributing to or participating in discussions in chat **	74%	50%	76%	82%	86%
Managing irritating or disruptive behavior (e.g. spam, nonsense, malicious links) **	65%	55%	62%	70%	78%
Manually deleting chat messages *	63%	58%	62%	66%	71%
Taking actions against individual viewers (e.g. bans, timeouts) **	61%	52%	54%	71%	79%
Monitoring conversations to proactively avoid conflicts **	53%	34%	51%	64%	63%
Manage general incivility (e.g. rudeness, immature behavior, social faux pas) **	51%	28%	46%	64%	71%
Warning offenders that they could be punished **	38%	30%	29%	46%	60%
Managing or customizing chatbots or chat commands **	31%	21%	30%	36%	43%
Learning or gathering information about viewers to support engagement	31%	21%	34%	35%	27%
Monitoring channel moderation actions performed by other mods **	28%	17%	25%	30%	46%
Handling targeted attacks against the streamer or community members **	26%	17%	20%	33%	47%
Explaining or clarifying to viewers why they were punished **	25%	14%	17%	31%	52%
Discussing or documenting incidents with other moderators or the streamer **	21%	13%	15%	25%	39%
Production assistance (e.g. managing song lists, editor-only tasks, community challenges)	19%	14%	20%	19%	23%
Helping the streamer manage the moderator team or select mods	19%	19%	19%	18%	20%
Filing reports about users to Twitch	15%	14%	14%	17%	16%
Training or educating other mods *	12%	8%	9%	15%	18%

Table 5. Fraction of mods in channels of each tier who performed each task within the preceding month. Significance testing was performed using logistic regression with $\log_{10}CCU$ as a predictor (* : $p < 0.01$, ** : $p < 0.001$).

6.1 Common moderation tasks

Survey respondents indicated using checkboxes which of a set of 14 tasks they had personally performed within the past month in the channel they had reported being an active moderator for. The set of tasks included in the survey were drawn primarily from Seering et al. [30], either directly from item labels or indirectly from descriptions of moderator work. Two additional items were added, including ‘Filing reports about users to Twitch’ and ‘Production assistance (e.g. managing song lists, editor-only tasks, community challenges).’

Table 5 below summarizes responses to these measures for the 918 moderators who provided a valid channel name. We find that moderators are typically performing a wide variety of tasks, from positive community-building and socialization to managing moderation tooling and enforcing against community members. For each task, we analyze the relationship with channel size using a logistic regression model which predicts the likelihood of a moderator having performed a specific task in a channel, given the size of the channel ($\log_{10}CCU$). We address below some themes which emerged from our findings.

6.1.1 Common and uncommon moderation tasks. Across all respondents, the most common moderation tasks include welcoming newcomers (75% overall) and participating in chat (74%). Among enforcement-related actions, the most common tasks include managing disruptions (65%), deleting messages (64%), and banning/timing-out individual viewers (61%). In public discourse, moderation is typically described implicitly or explicitly as an activity primarily focused on the removal of content or users. Our results show that positive community-building and socialization is the work most commonly performed by moderators, across communities of all sizes.

The least frequently self-reported tasks include managing mod teams (19%), filing reports to Twitch (15%), and training other mods (11%). Given that 68% of moderators in our survey indicated that they worked in the channel with teams of 3 or more mods, this could reflect either a lack of desire to coordinate with or train other moderators, or a lack of tooling to support these activities. Prior work has identified that community moderators often prioritize enforcement of community-specific norms rather than sitewide community guidelines [7, 11]. Our findings show that moderators are four times as likely to have banned or timed-out users from the channel than to have reported users for sitewide enforcement, quantifying the gulf between these levels of enforcement.

6.1.2 Relationships with channel size. As channels increase in size, moderators are more likely to spend time on tasks associated with managing disruptions and enforcing against users. Some of these tasks for which the frequency was most strongly associated with channel size include explaining punishments to viewers ($\beta = 0.83, p < 0.001$), managing general incivility ($\beta = 0.75, p < 0.001$), and warning offenders ($\beta = 0.63, p < 0.001$). All of these tasks represent labor which can be associated with helping viewers understand the channel-specific standards that moderators are hoping to maintain within a channel.

Perhaps unsurprisingly, enforcement actions against individuals are also reported as occurring more frequently as channels increase in size, including banning/timing-out users ($\beta = 0.55, p < 0.001$) and deleting messages ($\beta = 0.24, p < 0.01$). Finally, we note a set of coordination-related tasks which increase in frequency along with the potential size of the moderation team, including discussing/documenting incidents with other mods or the streamer ($\beta = 0.63, p < 0.001$) and monitoring actions performed by other mods ($\beta = 0.45, p < 0.001$). These may capture how moderation teams need to professionalize and introduce operating processes in order to effectively manage larger channels.

For a subset of moderation tasks, there is no evidence of a relationship between the size of the channel and the frequency with which these tasks were performed. For some streamer-focused tasks, such as production assistance or helping the streamer manage mod teams, we might expect that the amount of work is constant or slow-growing as audiences increase in size. Certain conversation-focused tasks such as welcoming newcomers or gathering information about viewers may become untenable or lower in priority as communities increase in size. Similarly, while larger channels are certainly more likely to experience behavior that breaks sitewide rules, the fact that moderators in larger channels aren't filing reports about users to Twitch at higher frequencies may reflect a shift in prioritization as channels increase in size.

6.2 Moderator roles

We also asked moderators to self-identify with a set of roles drawn from the taxonomy established by Wohn [36] using a set of check-boxes, from which they could select any that characterized how they approached their work in the channel, or provide their own. Respondents chose approximately two roles on average. We summarize their choices (N = 1053) here:

- **Conversationalist (C):** “My role is to actively engage and socialize with viewers to build a community.” (57% of respondents)

- **Helping Hand (HH):** “My role is to help the streamer and the community accomplish its goals.” (56% of respondents)
- **Justice Enforcer (JE):** “My role is to maintain a safe environment and punish bad behavior.” (42% of respondents)
- **Neutral Observer (NO):** “My role is primarily to observe, and I only participate or step in when absolutely necessary.” (37% of respondents)

Our final research question considered whether differences in moderators’ self-identified roles mapped to meaningful differences in how they engaged with the channel through the sets of tasks that they performed. To answer this question, we built a linear regression model for each task, which predicts the change in probability of performing that task associated with identification across the four roles. Though our outcomes are binary, we follow guidance from Gomila [16] that linear regression can yield acceptable outcomes with more interpretable results. As respondents could select multiple roles, we have as predictors four binary dummy variables, with one for each role. We summarize the results of these models in Table 6 below; we show coefficients for which $p < 0.001$, omitting for clarity coefficients which don’t meet this criterion:

For each task, the intercept can be interpreted as the probability of performing that task for a moderator who doesn’t identify with any of these four roles. For each role, we show the change in probability of performing the task associated with identification with that role. For example, 79% of mods who identify with the Conversationalist (C) role indicated that they had recently welcomed and socialized newcomers, compared with 58% of mods who didn’t identify with this role.

We observe some clear patterns of association between moderator’s self-reported high-level role within a channel and the likelihood of performing various moderation tasks, based on moderators’ self reports. Identifying as a conversationalist, for instance, is associated with substantially higher rates of having performed tasks associated with communication, including welcoming and socializing newcomers (+21%), participating in discussions (+16%), gathering information about viewers (+16%), and explaining to viewers why they were punished (+9%). Moderators identifying as helping hands are significantly more likely to gather information about viewers (+11%), provide production assistance (+11%), and help the streamer manage the mod team (+12%).

Self-identified justice enforcers, in turn, perform a wide variety of enforcement tasks at higher rates, with the largest differences for explaining punishments (+22%), warning offenders (+21%), banning/timing-out individuals (+20%), and managing general incivility (+20%). Notably, self-identified justice enforcers are twice as likely to report users to Twitch. Finally, we note that identifying as a neutral observer is not associated with any meaningful pattern in the tasks performed, compared with moderators identifying with none of these roles.

7 DISCUSSION

Volunteer moderators’ work has social depth that extends far beyond removal practices. In this paper, we have presented evidence that supports findings from prior work on the breadth of these practices. We have shown that moderators’ mental models for their work — as exemplified through social role descriptions — match with discrete sets of activities they perform, and we show one model for ways in which services and researchers might explore these activities and roles in quantitative and survey work in the future.

Differentiating moderator roles. Broadly, this work highlights the presence of discrete social roles among moderators, correlating with different approaches and philosophies toward moderation. Beyond designing to highlight and support social processes in moderation, services should consider moving beyond designing “one-size-fits-all” moderation tools. For example, a toolset for a “Conversationalist” Twitch moderators would likely look very different from a toolset for “Justice

	Intercept	Moderator Role			
		C	HH	JE	NO
Welcoming and socializing newcomers to the community	58%	+21%			
Contributing to or participating in discussions in chat	58%	+16%			
Managing irritating or disruptive behavior (e.g. spam, nonsense, malicious links)	50%			+16%	
Manually deleting chat messages	50%			+18%	
Taking actions against individual viewers (e.g. bans, timeouts)	50%			+20%	
Monitoring conversations to proactively avoid conflicts	37%			+13%	
Manage general incivility (e.g. rudeness, immature behavior, social faux pas)	36%			+20%	
Warning offenders that they could be punished	25%			+21%	
Managing or customizing chatbots or chat commands	22%				
Learning or gathering information about viewers to support engagement	12%	+16%	+11%		
Monitoring channel moderation actions performed by other mods	22%				
Handling targeted attacks against the streamer or community members	12%			+18%	
Explaining or clarifying to viewers why they were punished	9%	+9%		+22%	
Discussing or documenting incidents with other moderators or the streamer	10%			+12%	
Production assistance (e.g. managing song lists, editor-only tasks, community challenges)	8%		+11%		
Helping the streamer manage the moderator team or select mods	9%		+9%		
Filing reports about users to Twitch	10%			+10%	
Training or educating other mods	6%				

Table 6. Fraction of mods identifying with each role who performed each task within the preceding month. Significance testing was performed using linear regression with four binary dummy variables as predictors, one for each role. Percentages in role columns indicate the change from baseline for moderators of that role. Coefficients are shown in cases where $p < 0.001$; otherwise, they are omitted for clarity.

Enforcer” moderators. Services could provide opportunities for moderators to choose their preferred approach and then provide tools as appropriate, perhaps allowing moderators to switch between different views or layouts for moderation tools depending on circumstances and preferences. Note that, while self-identified social role correlated with activities performed by respondents in this survey, the majority of respondents identified with more than one social role.

Though the Twitch-specific typology of social roles used here aligns well with moderation practices in a livestreaming context, other broader conceptual frameworks may also contribute to understanding of moderator roles on different services. Yu et al. [37] identify nurture-as-care practices within certain approaches to moderation on MetaFilter, and Seering, Kaufman, and Chancellor [28] provide a taxonomy of moderator roles on Twitch, Reddit, and Facebook via the use of metaphors. Matias [26] discusses moderation work on Reddit as “Civic labor,” a process that involves variable attitudes and competencies for interfacing with different audiences, and in particular highlights the responsibilities of Reddit moderators in interacting with internal employees.

Designing for varied moderation contexts. The implications of these findings for design align with suggestions from prior work in a number of ways. For example, in highlighting the breadth of moderator practices, Seering et al. argue that services should develop tools “that support, rather than supplant, the judgment of users” [30, p. 3]. Similarly, Dosono and Semaan [10] recommend designing for empowering moderators through visible social support and designing moderation tools with emotional affect in mind. Modern services, such as Reddit, are often designed to make certain types of moderator work visible (e.g., content removal) and others invisible (answering questions and maintaining community health) [14] but our work helps show the importance of the social aspects of moderation labor that are frequently less visible.

Community moderators on Twitch, as well as on services like Reddit, Discord, and Facebook, are embedded in the social fabric of their communities in a way that centralized, company-run moderation teams can never be. Not only does this allow them to make moderation decisions based on more nuanced understandings of specific social context, but it also positions them to be proactive in cultivating communities in a way that centralized models could never match. For example, corporate moderation teams cannot realistically foster positive cultures in all of the user-created communities on a service by actively welcoming and socializing newcomers into each community, but this is an activity that nearly three quarters of survey respondents engaged in. Focusing on developing volunteer moderators’ abilities to cultivate positive spaces through tool development and education can provide a valuable complement to the strengths of centralized moderation practices.

8 CONCLUSION

In this short paper, we have presented results from a large-scale survey of Twitch moderators showing the prevalence of processes and practices identified in prior work. In line with prior work, we find a heavy social component to both the selection and onboarding processes for moderators as well as for the activities they perform, and we identify how these processes vary across different community sizes. With this work, we hope to build on prior research on moderator practices by translating qualitative findings into quantitative evaluations in order to inform the design of moderation tools that are a better match for the social realities of community moderation.

ACKNOWLEDGMENTS

First and foremost, we wish to acknowledge and thank the moderators who shared their experiences with us as part of this research. We wish to thank Michael Bernstein, Jeanne Chinn, Amanda Cullen, Skyler Ferry, Trevor Fisher, Stephen Hicks, Dominic Nguyen, and Alison Huffman for their feedback and support in designing and deploying this study. We also thank our GROUP reviewers, whose thoughtful comments helped us to substantially improve this work.

REFERENCES

- [1] Rajkumar Arun, Venkatasubramaniyan Suresh, CE Veni Madhavan, and Narasimha Murthy. 2010. On finding the natural number of topics with latent dirichlet allocation: Some observations. In *Pacific-Asia conference on knowledge discovery and data mining*. Springer, 391–402.
- [2] United States Census Bureau. 2021. *U.S. Census Bureau QuickFacts: United States*. Retrieved May 15, 2022 from <https://www.census.gov/quickfacts/fact/table/US/PST045221>
- [3] Jie Cai, Donghee Yvette Wohn, and Mashaël Almoqbel. 2021. Moderation Visibility: Mapping the Strategies of Volunteer Moderators in Live Streaming Micro Communities. In *ACM International Conference on Interactive Media Experiences (Virtual Event, USA) (IMX '21)*. Association for Computing Machinery, New York, NY, USA, 61–72. <https://doi.org/10.1145/3452918.3458796>
- [4] Juan Cao, Tian Xia, Jintao Li, Yongdong Zhang, and Sheng Tang. 2009. A density-based method for adaptive LDA model selection. *Neurocomputing* 72, 7-9 (2009), 1775–1781.
- [5] Robyn Caplan. 2018. Content or context moderation? <https://apo.org.au/node/203666>

- [6] Eshwar Chandrasekharan, Chaitrali Gandhi, Matthew Wortley Mustelier, and Eric Gilbert. 2019. Crossmod: A Cross-Community Learning-Based System to Assist Reddit Moderators. *Proc. ACM Hum.-Comput. Interact.* 3, CSCW, Article 174 (Nov. 2019), 30 pages. <https://doi.org/10.1145/3359276>
- [7] Eshwar Chandrasekharan, Mattia Samory, Shagun Jhaver, Hunter Charvat, Amy Bruckman, Cliff Lampe, Jacob Eisenstein, and Eric Gilbert. 2018. The Internet’s Hidden Rules: An Empirical Study of Reddit Norm Violations at Micro, Meso, and Macro Scales. *Proc. ACM Hum.-Comput. Interact.* 2, CSCW, Article 32 (Nov. 2018), 25 pages. <https://doi.org/10.1145/3274301>
- [8] Amanda LL Cullen and Sanjay R Kairam. 2022. Practicing Moderation: Community Moderation as Reflective Practice. *Proceedings of the ACM on Human-Computer Interaction* 6, CSCW1 (2022), 1–32.
- [9] Romain Deveaud, Eric SanJuan, and Patrice Bellot. 2014. Accurate and effective latent concept modeling for ad hoc information retrieval. *Document numérique* 17, 1 (2014), 61–84.
- [10] Bryan Dosoно and Bryan Semaan. 2019. Moderation Practices as Emotional Labor in Sustaining Online Communities: The Case of AAPI Identity Work on Reddit. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland Uk) (CHI ’19). Association for Computing Machinery, New York, NY, USA, Article 142, 13 pages. <https://doi.org/10.1145/3290605.3300372>
- [11] Casey Fiesler, Jialun Jiang, Joshua McCann, Kyle Frye, and Jed Brubaker. 2018. Reddit Rules! Characterizing an Ecosystem of Governance. In *Proceedings of the Twelfth International AAAI Conference on Web and Social Media (ICWSM 2018)*. AAAI, Menlo Park, CA, USA, 72–81.
- [12] Claudia Flores-Saviaga, Jessica Hammer, Juan Pablo Flores, Joseph Seering, Stuart Reeves, and Saiph Savage. 2019. Audience and Streamer Participation at Scale on Twitch. In *Proceedings of the 30th ACM Conference on Hypertext and Social Media (Hof, Germany) (HT ’19)*. Association for Computing Machinery, New York, NY, USA, 277–278. <https://doi.org/10.1145/3342220.3344926>
- [13] Dean Gengle. 1981. *Communitree* (first ed.). The CommuniTree Group, San Francisco, CA, USA.
- [14] Sarah A. Gilbert. 2020. “I Run the World’s Largest Historical Outreach Project and It’s on a Cesspool of a Website.” Moderating a Public Scholarship Site on Reddit: A Case Study of r/AskHistorians. *Proc. ACM Hum.-Comput. Interact.* 4, CSCW1, Article 019 (May 2020), 27 pages. <https://doi.org/10.1145/3392822>
- [15] Tarleton Gillespie. 2018. *Custodians of the Internet: Platforms, Content Moderation, and the Hidden Decisions that Shape Social Media*. Yale University Press, New Haven, CT, USA.
- [16] Robin Gomila. 2021. Logistic or linear? Estimating causal effects of experimental treatments on binary outcomes using regression analysis. *Journal of Experimental Psychology: General* 150, 4 (2021), 700.
- [17] Thomas L Griffiths and Mark Steyvers. 2004. Finding scientific topics. *Proceedings of the National academy of Sciences* 101, suppl_1 (2004), 5228–5235.
- [18] Starr Roxanne Hiltz and Murray Turoff. 1978. *The Network Nation: Human Communication via Computer*. Addison-Wesley Publishing Company, Inc., Boston, MA.
- [19] Shagun Jhaver, Iris Birman, Eric Gilbert, and Amy Bruckman. 2019. Human-Machine Collaboration for Content Regulation: The Case of Reddit Automoderator. *ACM Trans. Comput.-Hum. Interact.* 26, 5, Article 31 (July 2019), 35 pages. <https://doi.org/10.1145/3338243>
- [20] Jialun Aaron Jiang, Charles Kiene, Skyler Middler, Jed R. Brubaker, and Casey Fiesler. 2019. Moderation Challenges in Voice-based Online Communities on Discord. *Proc. ACM Hum.-Comput. Interact.* 3, CSCW, Article 55 (Nov. 2019), 23 pages. <https://doi.org/10.1145/3359157>
- [21] Sanjay R Kairam, Melissa C Mercado, and Steven A Sumner. 2022. A Social-Ecological Approach to Modeling Sense of Virtual Community (SOVC) in Livestreaming Communities. *Proceedings of the ACM on Human-Computer Interaction* 6, CSCW2 (2022).
- [22] Charles Kiene, Jialun Aaron Jiang, and Benjamin Mako Hill. 2019. Technological Frames and User Innovation: Exploring Technological Change in Community Moderation Teams. *Proc. ACM Hum.-Comput. Interact.* 3, CSCW, Article 44 (Nov. 2019), 23 pages. <https://doi.org/10.1145/3359146>
- [23] Charles Kiene, Andrés Monroy-Hernández, and Benjamin Mako Hill. 2016. Surviving an “Eternal September”: How an Online Community Managed a Surge of Newcomers. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (San Jose, California, USA) (CHI ’16). ACM, New York, NY, USA, 1152–1156. <https://doi.org/10.1145/2858036.2858356>
- [24] Claudia Lo. 2018. *When All You Have is a Banhammer: The Social and Communicative Work of Volunteer Moderators*. Master’s thesis. Massachusetts Institute of Technology.
- [25] T. M. London, J. Crundwell, M. B. Eastley, N. Santiago, and J. Jenkins. 2020. Finding Effective Moderation Practices on Twitch. In *Digital Ethics: Rhetoric and Responsibility in Online Aggression*, Jessica Reyman and Erika Sparby (Eds.). Routledge, New York, USA, Chapter 4, 51–68.
- [26] J. Nathan Matias. 2019. The Civic Labor of Volunteer Moderators Online. *Social Media + Society* 4, 2 (2019), 12 pages. <https://doi.org/10.1177/2056305119836778>

- [27] Elizabeth Reid. 1999. Hierarchy and Power: Social Control in Cyberspace. In *Communities in Cyberspace* (1st ed.), Marc A. Smith and P. Kollock (Eds.). Routledge, New York, NY, USA, 107–134.
- [28] Joseph Seering, Geoff Kaufman, and Stevie Chancellor. 2022. Metaphors in moderation. *New Media & Society* 24 (2022), 621–640. Issue 3. <https://doi.org/10.1177/1461444820964968>
- [29] Joseph Seering, Robert Kraut, and Laura Dabbish. 2017. Shaping Pro and Anti-Social Behavior on Twitch Through Moderation and Example-Setting. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing* (Portland, Oregon, USA) (CSCW '17). ACM, New York, NY, USA, 111–125. <https://doi.org/10.1145/2998181.2998277>
- [30] Joseph Seering, Tony Wang, Jina Yoon, and Geoff Kaufman. 2019. Moderator engagement and community development in the age of algorithms. *New Media & Society* 21, 7 (2019), 1417–1443. <https://doi.org/10.1177/1461444818821316>
- [31] Jeff T. Sheng and Sanjay R. Kairam. 2020. From Virtual Strangers to IRL Friends: Relationship Development in Livestreaming Communities on Twitch. *Proc. ACM Hum.-Comput. Interact.* 4, CSCW2, Article 94 (oct 2020), 34 pages. <https://doi.org/10.1145/3415165>
- [32] Anna DuVal Smith. 1999. Problems of Conflict Management in Virtual Communities. In *Communities in Cyberspace* (1st ed.), Marc A. Smith and P. Kollock (Eds.). Routledge, New York, NY, USA, 135–166.
- [33] Janet Sternberg. 2012. *Misbehavior in cyber places: The regulation of online conduct in virtual communities on the Internet*. Rowman & Littlefield, Lanham, MD, USA.
- [34] Jirassaya Uttaraopong, Jie Cai, and Donghee Yvette Wohn. 2021. Harassment Experiences of Women and LGBTQ Live Streamers and How They Handled Negativity. In *ACM International Conference on Interactive Media Experiences* (Virtual Event, USA) (IMX '21). Association for Computing Machinery, New York, NY, USA, 7–19. <https://doi.org/10.1145/3452918.3458794>
- [35] Michael Waters. 2020. *How a 1980s AIDS Support Group Changed The Internet Forever*. <https://onezero.medium.com/the-long-forgotten-story-of-ben-gardiner-the-aids-activist-whose-network-transformed-the-internet-c14460a73165>
- [36] Donghee Yvette Wohn. 2019. Volunteer Moderators in Twitch Micro Communities: How They Get Involved, the Roles They Play, and the Emotional Labor They Experience. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland Uk) (CHI '19). ACM, New York, NY, USA, Article 160, 13 pages. <https://doi.org/10.1145/3290605.3300390>
- [37] Bingjie Yu, Joseph Seering, Katta Spiel, and Leon Watts. 2020. "Taking Care of a Fruit Tree": Nurturing as a Layer of Concern in Online Community Moderation. In *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI EA '20). Association for Computing Machinery, New York, NY, USA, 1–9. <https://doi.org/10.1145/3334480.3383009>

Self-Reported Ethnicity	% of Respondents
White or Caucasian	63%
Hispanic or Latinx	16%
Asian or Asian-American	13%
Black or African-American	9%
Native American, Alaskan Native, First Peoples, or Aboriginal	3%
Middle Eastern	1%
Native Hawaiian or Pacific Islander	1%
Other (please specify)	2%
Prefer not to Disclose	5%

Table 7. Respondent self-reported race/ethnicity. Respondents could select more than one category.

A ADDITIONAL SURVEY DEMOGRAPHIC INFORMATION

Survey respondents self-reported their race/ethnicity by selecting from a set of options (with the ability to select multiple options); these results are summarized in Table 7. 63% of respondents self-identified as *White or Caucasian*, 16% as *Hispanic or Latinx*, 13% as *Asian or Asian-American*, 9% as *Black or African-American*, and 3% as *Native American, Alaskan Native, First Peoples, or Aboriginal*; 1% or fewer of participants identified with any of the remaining categories. Participants self-reported their disability status using two questions. The first asked “Do you identify as disabled or as having a disability or other chronic condition (regardless of whether you have a formal diagnosis)?” The 14.2% of individuals who responded ‘yes’ to this question were prompted with the option to select options from the list shown in Table 8.

To the best of our knowledge, this is the first large-scale survey of the Twitch user population to ask about race/ethnicity or disability status, so we have no previous points of direct comparison. However, the overall distribution of self-identified race/ethnicity of US-based Twitch moderator respondents aligns roughly with estimates for overall US demographics [2] when taking into account possible variance due to sample size and potential response biases.

Twitch moderators identify as having a disability or chronic conditions at higher rates than government estimates for the overall US population (14.2% of survey respondents, compared with 8.7% of the population under age 65 years in US Census data [2]), but we note that our definition includes mental health conditions, which are classified separately from the US Census definition of disability as “serious difficulty with four basic areas of functioning – hearing, vision, cognition, and ambulation,” making direct comparisons impossible.

B ADDITIONAL INFORMATION ABOUT LDA ANALYSIS OF MOTIVATION TEXTS

In Section 5.2, we present a high-level content analysis of common themes in respondents’ statements about their initial motivations to participate in channels as moderators. This analysis was aided by a Latent Dirichlet Allocation (LDA)-based topic model, which was used to identify latent underlying themes represented across these motivations. LDA models assume that each document (a single free-text response) is generated based on a probability distribution over a latent set of topics, each of which is represented as a probability distribution over the space of possible words.

In Table 9, we list the top ten tokens associated with each of these seven latent topics. Topic numbers are used to reference each topic, but the ordering of these topics is arbitrary. Topic 2 corresponds to the ‘Helping a Friend’ topic discussed above. Topics 3 and 5 represent the ‘Supporting the Community’ and ‘Creating a safe space’ topics, respectively. Topics 1, 4, 6, and 7, which we have chosen not to discuss in detail in our results, capture a broader set of experiences related to

Self-Reported Disability or Condition	% of all respondents	% of those self-reporting as disabled
Mental health condition	7%	48%
Health-related disability	4%	31%
Attention deficit	3%	21%
Mobility-related disability	2%	17%
Autism	2%	16%
Blind or visually impaired	1%	10%
Learning disability	1%	10%
Speech-related disability	1%	5%
Deaf or hard-of-hearing	1%	5%
Other (please specify)	2%	13%
Prefer not to disclose	1%	8%

Table 8. Respondent self-reported disability or condition. Participants could select more than one disability or condition. Overall, 14.2% of respondents identified as disabled or as having a disability or other chronic condition.

Top Terms by Tier of Channel Moderated						
1	2	3	4	5	6	7
chat	friend	commun*	mod	streamer	moder*	channel
people*	grow	twitch	streamer	safe	enjoi*	help
stream	close	love	stream	support	time	stream
viewer	mine	person	channel	feel	commun*	content
time	continu*	posit*	continu*	fun	watch	lot
happi*	irl	fun	engag*	content	bot	support
friendli*	respect	motiv*	start	watch	care	nice
game	success	interact	experi*	enjoi*	stai*	welcom*
talk	person	enjoy	watch	stream	streamer	time
activ*	streamer	experi*	friendship	keep	focu*	it

Table 9. Top ten (stemmed) tokens most strongly associated with each of the seven topics identified through LDA topic modeling over free-text responses capturing respondents' initial motivations to moderate channels. LDA represents each topic as a probability distribution over tokens, indicating that words higher in the list for a topic are the most likely to occur in text about that topic. Topics 2, 3, and 5 correspond to the topics discussed in our results above. Topics 1, 4, 6, and 7 capture broader and less clearly-differentiated aspects of streaming and moderation (background topics). Topics are unordered, such that the topic numbers do not carry meaning beyond uniquely identifying the different topics.

'channels', 'moderating', 'people', and 'streaming', such that each of these topics did not lend itself to a clear interpretation about some specific aspect of moderation.

Received May 2022; revised August 2022; accepted September 2022