

Proximate Social Factors in First-Time Contribution to Online Communities

Joseph Seering
Carnegie Mellon University
Pittsburgh, USA
jseering@andrew.cmu.edu

Jessica Hammer,
Geoff Kaufman
Carnegie Mellon University
Pittsburgh, USA
{hammerj,gfk}@andrew.cmu.edu

Diyi Yang
Georgia Institute of
Technology
Atlanta, USA
diyi.yang@cc.gatech.edu

ABSTRACT

In the course of every member’s integration into an online community, a decision must be made to participate for the first time. The challenges of effective recruitment, management, and retention of new users have been extensively explored in social computing research. However, little work has looked at in-the-moment factors that lead users to decide to participate instead of “lurk”, conditions which can be shaped to draw new users in at crucial moments. In this work we analyze 183 million messages scraped from chatrooms on the livestreaming platform Twitch in order to understand differences between first-time participants’ and regulars’ behaviors and to identify conditions that encourage first-time participation. We find that presence of diverse types of users increases likelihood of new participation, with effects depending on the size of the community. We also find that information-seeking behaviors in first-time participation are negatively associated with retention in the short and medium term.

Author Keywords

Newcomers; Twitch; Online communities; Social roles; Retention; Participation

CCS Concepts

•**Human-centered computing** → **Empirical studies in HCI**; *HCI theory, concepts and models*;

INTRODUCTION

The introduction of newcomers to online communities has been widely explored from a variety of perspectives in social computing research. From peer-production communities like Wikipedia and GitHub to creative fandom communities and channels on Twitch, the importance of recruiting and retaining newcomers has been thoroughly established [16, 26, 32]. New participants provide new perspectives and enthusiasm and help fill roles vacated through turnover, while affording established members the opportunity to deepen their engagement

with the community through teaching and mentorship. Moreover, research has found that users’ initial behaviors within a community can be strongly predictive of their long-term engagement [20], making these behaviors important to understand. Ultimately, a user needs to make an in-the-moment choice to participate for the first time, drawing from what they observe to decide if and when they would be welcomed, or whether they should join a particular community or move on to another. In this study we focus on these in-the-moment decisions where users decide to participate for the first time in communities on Twitch, identifying factors that lead them to participate and exploring which styles of initial participation are most associated with retention within the community.

Though this decision to participate is typically a user’s first *social* introduction to the community, the first-time participation that we focus on here is *not* always these users’ first experience within the community. Preece and Shneiderman’s Reader-to-Leader Framework [21] explains how users move from passive consumers to active participants, beginning by browsing content and observing communities often for lengthy periods of time prior to participating for the first time. These passive consumers, termed “lurkers” in popular online parlance and in previous work, can make up significant fractions of the population of online communities and social spaces [19]. The results we present here help to explain these lurkers’ transitions to first-time *participants*.

Though much other work has examined broader socio-cultural factors in participation in online communities, we focus here on the few moments immediately before a user decides to move from observer to participant. In synchronous environments like Twitch, users observe a flow of messages sent by other users, and these may help encourage them to decide whether to join the conversation for the first time. In these messages, they will observe signals about whether there is a place for users like them, whether they are likely to receive a response to their message or be welcomed by other users, and whether the community as a whole has a diverse set of members. In this work we analyze 183 million messages collected from Twitch channel chatrooms; we identify which messages were sent by new participants to each chatroom, where new participants are defined as users sending their *first* message to a specific channel; we compare the content in these first messages to content in messages posted by users who have previously participated in the channel; and we identify pat-

Permission to make digital or hard copies of part or all 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 third-party components of this work must be honored. For all other uses, contact the owner/author(s).

CHI’20, April 25–30, 2020, Honolulu, HI, USA

© 2020 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-6708-0/20/04.

DOI: <https://doi.org/10.1145/3313831.3376151>

terns of participation of other users immediately prior to new participants' first messages. We then identify which types of initial posting behaviors are most associated with retention. We conclude by proposing implications for future work based on these findings.

PRIOR WORK

Kraut et al. [16] identify five problems associated with newcomers in online communities: *Recruiting newcomers*, *Selecting the right newcomers*, *Keeping newcomers around*, *Teaching newcomers the ropes*, and *Protecting the Community from Newcomers*. The first of these has been a consistent challenge across a variety of online communities, where participation within the community is how the community self-sustains and generates value for members. For example, Wikipedia's ongoing challenges with recruiting and retaining new members have been well documented [7, 10]; while Wikipedia has no shortage of readers, it needs a steady flow of them to transition to editors in order to sustain the site. This same type of concern exists across a variety of other types of platforms, from fandom communities [6] to GitHub [18] to health support groups [30] and online games [5].

Prior work has approached these questions from a variety of directions. Preece and Shneiderman's "reader-to-leader" framework outlines different stages of participation [21]. "Readers" are users who explore a space via happening upon it as a result from a search term or a link from another site or a referral from a friend, and who browse the site and consume content without engaging or creating their own. The next phase of the process is "Contributor", which Preece and Shneiderman summarize as those who rate, tag, review, post, or upload; in this work we classify users who post comments as contributors. Preece and Shneiderman note a number of both usability and sociability factors that impact this transition [21, p. 21]. The former includes low technical thresholds for participation at small and large scales, visibility of contributions, and tools to manage misbehavior, and the latter includes social support for transitions between phases, a chance to build reputation and gain recognition, and policies for appropriate behavior.

Lurkers differ significantly from regular contributors in their relationship to the community. In work studying lurkers and regular participants across a range of forums and online bulletin boards, Ridings, Gefen, and Arinze found that lurkers trust other members significantly less than regular posters, and are significantly less motivated to get or give information or to exchange social support [22]. When users do transition from lurkers to occasional participants, their first contributions tend to be small, innocuous forms of participation. Ye and Fischer describe bug reporting as the first role users often hold in open-source software (OSS) spaces that actively contributes to the community [31], where bug reporting requires a much lower threshold to entry than actively reworking code. Similarly, Bryant, Forte, and Bruckman find that the first stages of engagement on Wikipedia include reading articles within interest areas and fixing small mistakes rather than making large changes or writing articles from scratch [2]. Yang, Kraut, and Levine found that common newcomer behaviors included asking questions of established members to better understand

the community and reduce uncertainty about how to engage [30]. As we note in our results, new participants' messages in our data matched these patterns; in this sample, new participants were much more likely than established members to ask questions and request information, though we find evidence that this was not necessarily a first step toward increased future participation. These cautious first steps are logical for new participants; the responses they receive to their first contributions may be read as major signals of whether they will be welcome on a platform [11].

Despite extensive exploration of new participants' first behaviors, little work has explored the *in-the-moment* factors influencing choice to transition from a lurker to a contributor from the perspective of the user [1]. Donath, in her work on online deception and social signaling theory [4] focusing on Usenet newsgroups, explores participation and discourse as a function of perceived social signals from other users. She splits identity signals into two groups. *Conventional signals* are cheap to send, as they do not require proof that a person actually has the trait they advertise, but are inherently unreliable. *Assessment signals* are those that are associated with clear proof, which can be costly for a person to produce but means that these signals are generally reliable.

Though most of the work within social computing research that builds on Donath's framing has focused on their application to asynchronous environments (e.g., [14, 17]), this framing can easily be applied to synchronous environments as well; as Donath notes, biology and game theory already study real-time co-located signaling in animals and humans. In this work we focus on social status signals, which on Twitch take the form of badges appearing next to usernames indicating a user's rank – moderator, subscriber, or regular user. These social status badges are assessment signals. They require either explicit appointment by an authority figure or a financial contribution. We work from hypotheses about these prior commenters' impact built from the above literature combined with an understanding of conventions of the Twitch platform. Prior work has shown significant impact of these specific status badges on the *content* of other users' subsequent posts [25], and we hypothesize that similar impact will be found on when users decide to post.

Prior work has identified a number of motivations for participating in Twitch streams. Gros et al. [8] find that "Information" and "Entertainment" are major motivations, both of which are tied significantly to the streamer and the content being streamed. However, Gros et al. also find a third motivation – "Socialization". Despite Twitch's focus on being a platform for built around interacting with streamers, Twitch users are also significantly motivated by the desire to socialize with other pseudonymous users in the associated chatrooms, with whom they may have no prior connections. Hilvert-Bruce et al. [13] and Sjöblom et al. [27] confirm and extend the importance of socialization to motivation to participation in Twitch chatrooms; Hilvert-Bruce et al. finds again that information-seeking and entertainment were important factors but also that social interaction, sense of community, and meeting new people all explained at least one dimension of measured en-

agement, as did a lack of external support in users’ personal lives.

In analyzing factors that may impact the decision to begin participation in Twitch chatrooms, we focus on the signals available to new participants. Whereas most previous work on newcomers has focused on forum-style interactions (e.g., [2, 3, 30]), Twitch is a synchronous chatroom environment, so a community is likely to have received a larger number of messages immediately prior to a new participant’s first post than on a space like Wikipedia. On Twitch, these interactions contain two types of information – signals about the social status of users as shown by status badges attached to their names, and the text content of their messages. Per Donath’s framing, conventional signals of users’ social status (i.e., low cost to produce and difficult to verify) abound on Twitch, notably through users’ own claims typed in chat. For example, a user could send a message to the chatroom claiming a personal connection to or friendship with a famous streamer, but proof for this assertion is unlikely to be able to be conveyed within the limits of a brief text message. In contrast, assessment signals about social status or identity, which are difficult to produce but can be easily verified, are less common in Twitch channels. The primary form that assessment signals take on Twitch are badges that appears next to users’ name in the chat: (1) green sword icons, which are platform-backed signals that identify moderators within the chat, and (2) variable but clearly identifiable icons that indicate users who pay a monthly fee (typically \$4.99) to *subscribe* to a particular streamer. These signals’ authenticity is platform-backed; a user can mouse over a badge to confirm its meaning. While other miscellaneous badges appear on Twitch (e.g., for special promotions or for donations of Twitch’s internal currency, “bits”), we focus on moderator and subscriber badges because of their ubiquity and the clarity of their meaning. We contrast these social status markers here in our models with “regular users”, users who have neither of these two badges, indicating a user who does not formally have status or authority in the channel.

As prior work has shown, a broad variety of factors contribute to how users engage with the Twitch platform, including on-stream content [13, 27], text of users’ messages [25], users’ pre-existing biases and motivations [29], and platform and channel-specific cultures [12]. We do not claim to model all of the factors impacting new participants’ decisions to contribute. Instead, we aim to show that recent presence of users with a diversity of roles is one significant contributing factor.

DATA COLLECTION

Twitch is a pseudonymous, public livestreaming platform where users stream a variety of content from playing games to making art to exploring foreign countries while other users watch and chat with each other and sometimes with the streamer. Twitch is composed of hundreds of thousands of *channels*, which are Twitch communities with stable URLs that are typically managed by a *streamer*. Each channel contains both a video *stream*, on which the streamer appears and through which they communicate, and an associated *chat*, the chatroom where users engage in synchronous, real-time conversations with each other and the streamer. Participation

is core to Twitch for several reasons. First, growth in participation can lead to increased revenue for streamers with communities above a size threshold, who are compensated for the number of views on advertisements shown on stream and for users’ paid subscriptions to the channel.¹ Second, participation is core to communities’ identities on Twitch [12]; though the streamer is a very visible figure in their channel, the participatory structure of Twitch works best when there are users for the streamer to engage with. Twitch reports more than 15 million daily active users and more than 140 million unique monthly users.²

In order to determine the effect of other users’ presence on users likelihood to participate for the first time, we gathered data from a subset of 1357 Twitch channels from June 3rd to August 18th, 2018. As channels can vary widely in number of active participants, we performed a pre-selection process to categorize channels before deciding which to scrape. Data on number of concurrent viewers in all active channels was collected via the Twitch API³ at regular intervals of ten minutes for one month. In order to better understand the interaction between channel size and social presence on new participation, we elected to divide channels into groups by median number of concurrent viewers as a proxy for size of community. We also elected to remove the largest and smallest streams from our pool. Though both very large and very small streams are important on Twitch, each was found less suitable for this analysis for different reasons. The largest channels in this sample had message volumes that were great enough such that a significant number of messages would be sent between the time a user decided to send their own message and the time they finished typing it, making it problematic to assess the impact of immediately recent behaviors. The smallest channels, at the time of data collection, were much less likely to be contractually “partnered” with Twitch, meaning that a large fraction of the channels did not have a subscription feature.

Prior work has divided Twitch channels by size in a number of different ways. Hamilton, Garretson, and Kerne primarily classify streams as “small” or “large” [12], while [24] divide channels into five size groups. We follow the latter approach, though in order to specify exactly five clusters we perform k-means clustering instead of the mean shift algorithm used in [24]. K-means clustering allows specification of number of clusters to identify. The center of each cluster and maximum size within the cluster in our dataset are shown in Table 1, along with the number of streams in each category in our final dataset. For example, the largest cluster had a center of 7085 concurrent viewers, and six channels in our final dataset fell within this cluster with a maximum size of 7730 median concurrent viewers. In order to test the hypotheses we present here, we focused on the three size categories shown in Table 1. Note that, while we use *Large*, *Medium*, and *Small* here for the sake of simplicity, these should be taken as relative terms.

¹Not every streamer actively seeks growth. Some are satisfied with the current size of their communities. However, recruitment of newcomers is still important, as maintaining size requires adding new members to replace those who leave.

²<http://twitichadvertising.tv/audience/>

³<https://dev.twitch.tv/docs/api/>

Table 1: Size categories of streams used in this analysis and associated viewer counts

Cluster	Center of cluster	Max size	Count in dataset
Large	3649	4901	13
Medium	1592	2534	31
Small	562	1122	113

We curated a final list of Twitch channels to scrape based on this size category classification, selecting more streams in each successively smaller size category with the goal of collecting a similar number of messages from each. Messages were collected over the course of 75 days via a custom-built Python script using the Socket library⁴. All messages sent to these channels during this time period were collected. The complete dataset contained 183 million messages.

Each message collected had several associated pieces of meta-data ranging from username to timestamp to status badges. Of these, we used the badges denoting whether the message was sent by a moderator, subscriber, or neither. To this, we appended a feature denoting whether each message was the user’s first message posted to the channel. We classified messages as a user’s first message if that user had not posted at all in the channel in the two months prior. As such, while data was scraped for a total of two and a half months, data from the first 60 days was used to establish a baseline for prior activity of users in the channel in order to be able to identify new participants; the models presented in subsequent sections are built based on the final 15 days of activity.⁵

Table 2 shows the primary features used in our analyses – moderators, subscribers, regular users, new participants, and bots, with bots being the held-out category in our analyses because they do not contribute to the theoretical goals of this work – along with the percentage of messages sent by each across each of the size categories. In order to manage covariates, each message was only labeled with one out of the first three tags; we assume that users perceived moderators primarily as moderators even if they were also subscribers. Table 3 shows the names of variables in this study as they appear in models.

Feature	Large	Medium	Small
Moderators	10%	5%	2%
Subscribers	44%	32%	32%
Regular users	40%	57%	61%
New participants	2%	3%	3%
Other (including bots)	4%	3%	2%

Table 2: Percentage of participant types across channel sizes by total message volume

In order to preserve privacy of users in the sample, no individual users’ behavior was aggregated or analyzed in depth in order to maintain user privacy. All quantitative analysis was performed using scripts to calculate statistics in aggregate,

⁴<https://docs.python.org/2/library/socket.html>

⁵Note that messages sent in “subscribers-only” mode, a chat feature that only allows paid users to talk, were removed from the dataset.

Variable name	Description
#modsinlast10	Number of messages in the previous ten that were sent by moderators
#subsinlast10	" " " subscribers
#reginlast10	" " " regular users
#newinlast10	" " " new participants in the channel

Table 3: Variables used in subsequent analyses and their descriptions

and qualitative analysis was performed on short chunks of text identified through searches for pre-determined properties. This study was approved by the Institutional Review Board at Carnegie Mellon University.

RESEARCH QUESTIONS AND HYPOTHESES

In this paper we focus specifically on the impact of participation by users with clear social roles of different types in the moments immediately leading up to first-time participation. Based on prior work (e.g., [24, 25, 26]), we expect these social factors to contribute to explaining users decisions to participate for the first time. There are three primary research questions that drive this work:

1. How is new participation different from established members’ participation?
2. How does presence of users of different social statuses impact likelihood for users to participate for the first time, and how do these effects vary across different sizes of communities?
3. How are these effects related to subsequent retention of new participants?

The following are hypotheses derived from prior work related to newcomers and work focused on the Twitch platform that address each of these questions. The first group focuses on characteristics of new participants’ messages:

H1.1: New participants’ messages will be shorter than established members’ messages.

H1.2: New participants will ask more questions than established members.

H1.3: New participants will use more bot commands than established members.

H1.4: Despite asking more questions, newcomers will engage less broadly with others of the community, focusing primarily on the most visible members - the streamer and the bots.

First-time contributors will be hesitant to make larger, more visible contributions. Prior work, e.g., [2, 31], has found that newcomers’ first contributions are small, simple, and come from areas of comfort or expertise. On Twitch, these will manifest as greetings, brief questions, and emotes. Common forms for first-time participation identified in prior work are asking questions and seeking information about how to behave [9, 30]. On Twitch, as on other platforms, we expect these questions to be both informational, e.g., “what mouse are you using?” and normative, e.g., “is it okay for me to post a link?”.

We predict that new participants will engage with automated agents on Twitch at higher rates than regular users due to information-seeking motivations. Seering et al. [24] identified chatbots as significant contributors to Twitch communities, with providing information being one of their major functions. A “command” is simply a message containing a specific keyword, often beginning with an “!”, which the chatbot identifies and responds to in a predetermined way. For example, in some channels users may type “!twitter”, which causes a bot to respond with the streamer’s Twitter account URL.

The second group of hypotheses focuses on the impact of other types of users’ presence in the chat in the recent past. We model the impact of the number of messages coming from each type of user - moderators, subscribers, and regular users - on new participation:

H2.1: A small volume of recent moderator participation will lead to increased rates of first-time participation.

H2.2: Recent participation by paid users will lead to decreased rates of first-time participation.

H2.3: Larger overall volume of participation from regular users will lead to greater participation from newcomer

As noted in prior research [25], moderators have significant influence over the behavior of other users in the chat. When they behave positively, other users follow suit, but too much moderator participation, which users refer to as “mod spam” or “mod walls”, is thought to discourage participation by other users [26]. Thus, we hypothesize that a small number of prior messages from moderators will encourage new participation. In contrast, active participation of subscribers, who have made a recurring financial commitment to the community, will slightly discourage new members from joining the community because they will feel that there is a higher barrier to legitimate participation within the community. Per Kraut et al. [16], social barriers to participation can increase the commitment of users who overcome those barriers, but drive away users who cannot or do not want to.

Finally, we anticipate that new participants will engage more when they observe a large volume of participation from users similar to them. A significant volume of participation from regular users will indicate that the community is not exclusive and that there is a place for newcomers to participate.

Our final hypotheses focus on factors affecting subsequent retention. Our goal is to understand whether different forms of initial participation are associated with different subsequent rates of subsequent participation. For example, are users who ask questions in their first post likely to post more or fewer subsequent messages than those who don’t?

H3.1: The presence of information-seeking behavior in newcomers’ first comments, including asking questions and using bot commands, will lead to a greater subsequent volume of participation.

H3.2: The presence of socialization behavior in newcomers’ first comments, including messages targeted toward another

user or the streamer, will lead to a greater subsequent volume of participation.

Here we draw from prior work that emphasizes the prominence of information-seeking as a common behavior of newcomers [30] and that information-seeking and social interactions are core motivations of a subset of Twitch users [13]. Hypotheses 3.1 and 3.2 are based on the idea that users whose initial behavior is connected to established motivations for participation may be more likely to continue participating in the future.

The following three sections each cover one of these categories of analyses in the order that they are presented above. Note that, due to the number of dimensions of the data, in some cases we do not include all tables and figures for the sake of brevity when findings are redundant.

ANALYSIS 1: OVERALL CHARACTERISTICS OF NEW PARTICIPANTS’ MESSAGES

Before modeling factors that contribute to first-time participation, we explore whether first-time messages are significantly different in content from regulars’ messages. We identify three categories for basic analysis of messages in order to highlight basic distinctions between first-time participation and regulars’ messages: message length, presence of questions, and messages directed toward bots. Per the above hypotheses **H1.1**, **H1.2**, and **H1.3**, we expect new participants’ messages to be shorter than regulars’, to contain more questions, and to be targeted more at the streamer and bots (who are the most obvious and visible members of the community) and less at other community members (who are less visible and likely will take longer to connect with).

In order to address these hypotheses, we selected a random subset of 10,000 new participants’ messages and 10,000 established members’ messages from across the full dataset and calculated values for the above five categories. We defined bot commands as messages starting with ‘!’, which is a standard convention on Twitch. Table 4 shows these comparisons, where length is shown as number of characters and all other values are proportion of messages containing each feature.

We find the differences predicted above to be accurate in all cases. New participants’ messages are approximately **25% shorter** than established members’ messages, confirming H1.1. New participants ask questions **125% more often** than established members in their first messages, with slightly over 15% of their first messages containing questions compared to 6.7% for established users’ average messages, confirming H1.2. New participants use bot commands **twice as often** in their first messages than established members do in their average messages, with 7.6% of new participants’ first messages containing a bot command, confirming H1.3. When directing their messages toward other users, newcomers targeted **half as many toward other users** as established members and **50% more toward the streamer**, confirming H1.4.

Though these specific ratios are likely situationally-dependent and may vary over time, their agreement with predictions based on prior work studying different contexts (e.g., [2, 9, 30, 31]), suggests certain generalizability across platforms.

	N_{new}	N_{reg}	μ_{new}	μ_{reg}	SD_{new}	SD_{reg}	t-score	p-value
Length	10,000	10,000	26.1	34.4	36.5	48.7	13.6	<0.0001***
Questions	10,000	10,000	0.151	0.067	0.358	0.112	22.2	<0.0001***
Bot commands	10,000	10,000	0.076	0.038	0.265	0.190	11.8	<0.0001***
@member	10,000	10,000	0.041	0.082	0.199	0.248	12.1	<0.0001***
@streamer	10,000	10,000	0.013	0.009	0.092	0.112	2.96	0.0031**

Table 4: Comparisons between features of new participants' messages and regulars' messages

ANALYSIS 2: MODELING LIKELIHOOD OF FIRST-TIME PARTICIPATION BASED ON CHARACTERISTICS OF RECENTLY ACTIVE USERS

The following section presents the core analysis of this work - models predicting likelihood of first-time participation based on social status of recently active users in the channel. We present these models, explain the underlying assumptions, and discuss key takeaways. We then visualize these models via use of odds-ratios, clearly showing the impact of presence of user types as they increase from zero to five out of the last ten messages.

Polynomial logistic regression models for likelihood of first-time participation

We begin by presenting models predicting probability of a subsequent message coming from a new participant (as opposed to a regular) based on social status of users who posted the previous ten messages in the chat. We elect to use multivariate logistic regression, which is more appropriate than linear regression in this case because the outcome is either zero (not posted by a new participant) or one (posted by a new participant). We also elect to use a second order polynomial model; in preliminary exploration of the data, we found that the proportions of new participant messages were distinctly non-monotonic in relation to input variables, meaning that in many cases the proportion of new participant messages increased as the number of messages from a particular group increased from, e.g., zero to one or two, but reached a peak and declined as the number of messages increased toward ten. The graphs presented later in this section clearly show these nonlinear relationships.

Before finalizing these models, we addressed several brief methodological questions in order to validate our approach. First, we tested to see whether factors were significantly correlated, which would impact how we interpret our results. We calculated correlations between each pair of independent variables. The results show expected correlations. The recent participation of regular users is moderately strongly negatively correlated with the recent participation of moderators and subscribers; if most of the previous messages were sent by subscribers, this means that, by definition, few of the previous messages will have been sent by regular users. Correlation matrices are shown in Appendix A. Next, in developing these models we also tested variants using both the prior ten and prior thirty messages as inputs, but found that messages further in the past had much smaller influence on subsequent probability of first-time participation. See Appendix B for supporting analysis. For the sake of brevity and clarity we present only

the models based on ten prior messages here, as the models based on thirty prior messages provide no additional insight. We also performed a robustness check on these models by excluding unusual participation patterns, i.e., instances with 8, 9, or 10 messages sent by moderators or subscribers, and found results consistent with those we present here.

The models presented in Tables 5, 6, and 7 show the signs and relative magnitudes of variables. In most cases the first and second order terms for single variables have opposite signs, indicating that the effect reverses direction after a certain value has been reached, supporting the non-monotonic assumptions inherent in **H2.1** but partially contradicting **H2.2** and **H2.3**. Note that, because these are second order polynomial models, the full models include six interaction terms covering all interactions between the three main variables. We omit these terms in Tables 5, 6, and 7 for simplicity as they fall beyond the scope of the hypotheses here and attempting to explain all of them would necessitate post-hoc speculation, but we include them in Appendix C.

	Coefficient	SE	z-score	p-value
Intercept	-4.548	0.089	-51.373	<0.001 ***
#modsinlast10 ²	-0.029	0.002	-11.903	<0.001 ***
#modsinlast10	0.302	0.028	10.777	<0.001 ***
#subsinlast10 ²	-0.015	0.001	-10.605	<0.001 ***
#subsinlast10	0.147	0.022	6.603	<0.001 ***
#reginlast10 ²	-0.017	0.001	-11.268	<0.001 ***
#reginlast10	0.270	0.023	11.834	<0.001 ***
#newinlast10 ²	-0.050	0.003	-15.898	<0.001 ***
#newinlast10	0.993	0.034	29.446	<0.001 ***

Table 5: Model of likelihood of new participant message based on prior number of messages by moderators, subscribers, bots, and other new participants in **small** streams. Nagelkerke's $R^2 = 0.038$

	Coefficient	SE	z-score	p-value
Intercept	-5.384	0.124	-43.282	<0.001 ***
#modsinlast10 ²	-0.054	0.004	-13.220	<0.001 ***
#modsinlast10	0.574	0.045	12.820	<0.001 ***
#subsinlast10 ²	-0.026	0.002	-12.166	<0.001 ***
#subsinlast10	0.392	0.032	12.239	<0.001 ***
#reginlast10 ²	-0.076	0.002	-39.514	<0.001 ***
#reginlast10	0.912	0.030	29.910	<0.001 ***
#newinlast10 ²	0.015	0.005	3.217	0.001 **
#newinlast10	0.470	0.050	9.419	<0.001 ***

Table 6: Model of likelihood of new participant message based on prior number of messages by moderators, subscribers, bots, and other new participants in **medium** streams. Nagelkerke's $R^2 = 0.028$




Figure 1: Probabilities of subsequent new message for each value of independent variables, relative to when those values are zero, in **small**, **medium**, and **large** streams respectively

	Coefficient	SE	z-score	p-value
Intercept	-2.734	0.104	-26.24	<0.001 ***
#modsinlast10 ²	-0.006	0.008	-0.72	0.474
#modsinlast10	-0.435	0.079	-5.52	<0.001 ***
#subsinlast10 ²	-0.040	0.002	-17.38	<0.001 ***
#subsinlast10	0.282	0.032	8.84	<0.001 ***
#reginlast10 ²	-0.017	0.002	-10.56	<0.001 ***
#reginlast10	0.101	0.025	4.12	<0.001 ***
#newinlast10 ²	0.051	0.005	10.22	<0.001 ***
#newinlast10	-0.130	0.051	-2.54	0.011 *

Table 7: Model of likelihood of new participant message based on prior number of messages by moderators, subscribers, bots, and other new participants in **large** streams. Nagelkerke’s $R^2 = 0.026$

The results shown in these tables provide mixed evidence for and against **H2.1**, **H2.2**, and **H2.3**. In the explanation we provided for our **H2.1**, we hypothesized that large volumes of messages posted by moderators (i.e., “mod walls”) could have a negative impact on newcomer participation compared to the impact of moderate volumes of messages, and we found this to be the case in all three size categories. However, our prediction that moderate volumes of moderator messages were only beneficial in small and large streams. Thus, **H2.1** was supported in most but not all cases. Our **H2.2** was half-supported: we had posited that subscriber participation would be a consistent detriment to newcomer participation, but found this only to be the case with large volumes of subscriber participation; the presence of subscriber messages had a generally positive impact with magnitude that increased up until a peak at between 3 and 8 recent subscriber messages, depending on size category. Finally, our **H2.3** was also partially supported in that a small to moderate volume of recent participation by regular users increased likelihood of newcomer participation, but that the magnitude of this effect began to decrease at higher volumes. We discuss possible explanations for all of these observed effects in our section on future work.

Broadly, the coefficients also show that a diversity of participation from different types of users led to increased participation of newcomers overall, lending support to the idea that it is beneficial for a variety of social roles to be represented.

Model visualization and interpretation

While the above models provide extensive statistical evidence for the impact of different factors, they are difficult to directly interpret in the context of a real situation in a Twitch channel. In order to provide more interpretable results, we visualize the above models in Figure 1 using odds-ratios. Odds-ratios are “the odds that an outcome will occur given a particular exposure, compared to the odds of the outcome occurring in the absence of that exposure” [28]. In a traditional logistic regression, the regression coefficient for a given variable is the estimate of the increase in the log odds of an occurrence given a one unit increase in the variable. The formula for odds-ratios in second order polynomial logistic regressions is somewhat more complicated, as the odds-ratio depends on what values of the variable one is starting and ending at rather than being constant. Given sample coefficients a_2 and a_1 of #modsinlast10² and #modsinlast10 at state x , the odds ratio between a state and the subsequent state is given by the following formula:

$$\frac{Odds_{x+1}}{Odds_x} = e^{a_1 + a_2 * (2x+1)}$$

The three graphs in Figure 1 show the relative probability of the next post coming from a new participant as we vary values for each of the major three variables - number of moderator, subscriber, and bot posts - from zero to five out of the last ten messages. We set the starting probability at one for the initial value of zero messages for each of these categories to show relative rather than absolute changes. For example, the probability of a post coming from a new participant is 1.41 times as large when two of the previous ten posts came from a moderators as it would be if zero of the ten previous posts came from a moderator in a small stream. Note that the plot for the mid-sized streams uses a different y-scale to accommodate the larger graph peaks. These plots clearly visualize the relationships shown in Tables 5, 6, and 7, which are discussed in the previous section.

	Period 1 (0-5 days)		Period 2 (6-10 days)		Period 3 (11-15 days)	
	Estimate	Std. Error	Estimate	Std. Error	Estimate	Std. Error
(Intercept)	4.414***	0.104	1.711***	0.117	1.325***	0.114
question	-1.045***	0.251	-0.484	0.282	-0.358	0.274
@streamer	-0.283	0.550	-0.109	0.619	-0.125	0.602
@member	-1.311	0.706	-1.080	0.795	-0.755	0.772
bot	-2.183***	0.317	0.255	0.357	0.439	0.347
question*@streamer	-0.328	0.957	-0.188	1.077	-0.165	1.046
question*@member	1.288	2.224	0.852	2.503	0.471	2.433

* = $p < 0.05$, ** = $p < 0.01$, *** = $p < 0.001$

Table 8: Factors predicting new participants' activity in small streams 0-5, 6-10, and 11-15 days into the analysis window

	Period 1 (0-5 days)		Period 2 (6-10 days)		Period 3 (11-15 days)	
	Estimate	Std. Error	Estimate	Std. Error	Estimate	Std. Error
(Intercept)	5.485***	0.164	2.022***	0.100	1.817***	0.097
question	-2.054***	0.416	-0.808**	0.254	-0.906***	0.245
@streamer	-1.957*	0.877	-0.600	0.535	-0.885	0.516
@member	-1.873	1.033	-1.173	0.630	-0.725	0.608
bot	-3.790***	0.602	-1.056**	0.367	-0.966**	0.354
question*@streamer	0.492	1.549	0.793	0.945	1.203	0.911
question*@member	5.054	3.892	2.482	2.374	2.546	2.291

* = $p < 0.05$, ** = $p < 0.01$, *** = $p < 0.001$

Table 9: Factors predicting new participants' activity in medium streams 0-5, 6-10, and 11-15 days into the analysis window

ANALYSIS 3: RELATIONSHIP BETWEEN TYPE OF FIRST-TIME PARTICIPATION AND SUBSEQUENT RETENTION

Though the above exploration has focused on factors that precede new participation, subsequent retention is also important to consider. In this section we examine the content of first messages sent by newcomers to determine whether factors used in Analysis 1 predict how much these users participated in subsequent weeks. We use four of the factors from Analysis 1: whether a new participant's first message contains a *question*, whether it contains a *bot command*, whether it includes a *tag of another user*, and whether it includes a *tag of the streamer*. The first two of these factors relate to the information-seeking motivation in **H3.1**, while the second two relate to the socialization motivation in **H3.2**. We also include two interaction terms—whether the message includes a question and also tags another user, and whether the message contains a question and also tags the streamer—because we predict that the impact of questions on newcomers' subsequent participation will depend upon who they direct their question toward. We do not include other interaction terms between the four main variables because they do not make sense in this context, e.g., bot commands are directed toward bots, not other users or the streamer, so a term indicating whether a message includes a bot command and also tags the streamer does not make logical sense. We also do not include the length variable from Analysis 1 because we do not find it to be theoretically meaningful.

Tables 8, 9, and 10 show regression coefficients for multiple linear regressions performed predicting number of posts made by new participants in the first 0-5, 6-10, and 11-15 days of the analysis period, from the start of day 61 of data collection through the end of day 75, across the size categories. We had to begin this analysis at day 61 rather than day 1 because, per our method, we could only identify newcomers starting on day 61. Across all three size categories, new participants posted a baseline of between 4 and 6 messages on average in the first period and between 1-2 on average in the next two periods.

The coefficient for each variable shows association between presence of that factor in the first message and subsequent number of messages. Across all three size categories, first messages with bot commands and first messages with questions were negatively associated with retention in the first time period. Asking questions in first messages was also negatively associated with retention in large and small streams in the first time period and medium streams across all three time periods, and use of bot commands in first messages was negatively associated with retention in all three size categories of streams across most time periods. The interaction effects were not significant in any cases, and tagging streamers or other community members in messages was not significantly associated with retention in eight out of nine cases.

These results show a strong negative connection between information-seeking behaviors, such as use of bots commands and asking questions, and subsequent retention. It would be plausible to hypothesize that users with information-seeking

	Period 1 (0-5 days)		Period 2 (6-10 days)		Period 3 (11-15 days)	
	Estimate	Std. Error	Estimate	Std. Error	Estimate	Std. Error
(Intercept)	4.656***	0.107	1.775***	0.069	1.154***	0.057
question	-0.995***	0.289	-0.064	0.187	-0.080	0.154
@streamer	-0.808	0.563	0.238	0.364	-0.113	0.300
@member	-1.012	0.722	-0.261	0.466	0.158	0.385
bot	-2.797***	0.296	-0.823***	0.191	-0.560***	0.158
question*@streamer	1.100	0.978	-0.860	0.632	0.580	0.521
question*@member	-0.184	2.524	-0.226	1.630	-0.267	1.344

* = $p < 0.05$, ** = $p < 0.01$, *** = $p < 0.001$

Table 10: Factors predicting new participants' activity in large streams 0-5, 6-10, and 11-15 days into the analysis window

motivations may engage when they perceive that information is readily available, as when bots are particularly visible, but that this may not lead to subsequent additional participation because these users' motivations for participation have been satisfied. Per the literature on Twitch users' motivations [8, 13, 27] and the behaviors of newcomers more broadly [30], information-seeking is an important motivation for early participation, but we suggest here that it may not in and of itself be a pathway from, e.g., a "Reader" to a consistent "Contributor", per Preece and Shneiderman [21].

This suggests that that there could be value in future work that looks in depth at the relationship between proximate social factors, newcomer motivations, and retention. It is likely that there are other factors that we could not measure here that explain these relationships. For example, the impact of questions here may be significantly mediated by whether or not they were answered - a direct, positive response from the streamer or another community member may increase likelihood to participate more, while a lack of a response or an unsatisfying response could significantly decrease likelihood, leading to a negative overall coefficient. We do not have data in this sample to determine at scale whether or how questions that were asked were subsequently answered, but we note this as a significant question that should be explored in future research. Similarly, our measures of socialization behaviors were fairly basic, and could be examined in more depth in the future in order to test multiple additional dimensions of the socialization as operationalized on platforms like Twitch.

DIRECTIONS FOR FUTURE WORK

This work aims to answer questions about the impact of presence of users of various levels of social status on users' likelihood to participate for the first time, and subsequent relationships to retention. First, we found clear differences between the types of first posts made by new participants and regular users' average posts. These differences matched existing literature about newcomers' behaviors. Next, we modeled the impact of three types of users' recent presence in the chat - moderators, subscribers, and regular users, and found a distinctly non-monotonic relationship between volume of recent participation for each group and impact on subsequent newcomer participation. Finally, we explored associations between initial behaviors and subsequent retention. Broadly,

this work contributes to a long line of prior work on recruiting and retaining newcomers and suggests directions for further exploration of the properties of opposing social factors and how they generalize to other contexts. In this section, we discuss possible directions for some of this future work.

Impacts of Moderator Participation

In this work, we found that moderators can increase newcomer participation through low to moderate engagement, particularly when those communities are smaller. Seering, Kraut, and Dabbish [25] have shown that moderators can significantly improve community behavior by using their authority to set examples of positive behaviors, but other work discusses negative impact from "mod walls" or "mod spam". The positive effects we observe could stem from moderators' *authority* but a corresponding negative effect could be due to *intimidation* of new users who aren't comfortable participating in a space full of so many high-status users. We hypothesize that moderators can make a difference in recruiting new participants by showing clearly that they will engage with regular users, helping new participants see them as less intimidating.

Other factors that may mediate the effects we observed in this work include how moderators' posts appear and which posts are shown publicly. On platforms like Reddit, moderators are allowed to decide which posts they make will be "distinguished", where posts that moderators choose to distinguish will be accompanied by a symbol announcing that the post was made by a moderator. Prior work showed that some Reddit moderators do not distinguish their posts when it isn't necessary to show their moderator status in order to avoid unduly influencing the conversation [26]. A future experiment could explore a scenario where Twitch moderators are only "distinguished" as moderators when taking an official action.

Impacts of Subscriber Participation

Subscribers' participation in these communities had a generally-positive effect on newcomer participation across all three community sizes. Prior work has discussed the impact of community identity on participation [15]. A moderately-sized subscriber base could indicate a community with more internal *cohesion* and a stronger *identity*. In particularly large streams, cohesion and identity are difficult to achieve. Though some level of subscriber participation was positive, a large

volume of recent subscriber participation had a less positive effect, which we suggest may be the result of a perceived *barrier-to-entry* effect, possibly resulting from the perception that, due to active presence of paid users, the community is more exclusive. Thus, perceived financial and social barriers to entry may discourage users from participating for the first time.

Future work could explore the impact of status signals that signal different levels of commitment; since this data was collected, Twitch added more badges that signal status, including a “VIP” badge typically given to users with a closer connection to the streamer or some external status. Also, in addition to the “bits” badge mentioned briefly earlier in this work, Twitch added a badge users can acquire by paying for other users to subscribe (“gifting” subscriptions). These each indicate a type of status different from just subscription, but each requires something that newcomers may not have (i.e., social influence or money). Status signals on other platforms that are worth future analysis include identity tags (flair) on Reddit and assigned “roles” on Discord.

Roles for Bots

This work has focused primarily on the impact of human social status signals, but the impact of bots on newcomer participation is also an important question. In prior work, Halfaker et al. [10] found that users’ whose first Wikipedia edits are reverted by algorithmic tools have much lower retention rates, indicating that having first experiences with humans may be better for retention than with bots, but Halfaker et al. only considered negative initial interactions (i.e., reverts). In this case we see that initial interactions with bots are associated with lower retention, but this finding could be explored across numerous other contexts.

Differentiating Information-seeking Behaviors

Multiple questions emerge from this work, most notably surrounding the relationship between users’ motivations, their early modes of engagement, and retention. In particular, though information-seeking has been identified as a common early motivation for users, prior work has not found as clear an association between certain types of information-seeking behaviors and lower retention as we present here. Further exploration of this relationship could lead to personalization of approaches to recruiting new participants based on their assessed motivations. This suggests broader questions about platforms’ roles in supporting or shifting users’ motivations - if users seek only to engage to gain information, should platforms be working to recruit them to participate more? Should recruitment focus on users who would like to participate but are uncertain whether or how they can, or should it also target users who don’t yet have a desire to participate? Future work could further explore a broader variety of initial participation types than the four studied here and could explore established users’ reactions to newcomer questions.

LIMITATIONS

As noted previously, there are several main limitations to this work. First, in terms of methodological limitations, this work focuses specifically on presence of users of different types

rather than the content of the messages they send, which we feel is an important topic to explore but too large a question to be explored fully here in addition to our current analyses. Second, we chose to focus only on factors occurring immediately prior to initial participation, and our work shows that these moments are important, but newcomers may have longer-term familiarity with the culture and norms within channels and on the platform as a whole that is also significant in their decision to participate. Third, though the overall volume of messages in our dataset was large, the final number of streams analyzed, particularly in the “Large” size category was relatively small. Because of greater overall volume of participation in larger streams as compared to smaller streams, it is not mathematically possible to balance both the number of streams in each size category and the volume of messages in each size category while still scraping each for the same overall time period. Future work could examine cross-channel variation in larger channels in more depth.

A final limitation results from the levels of collinearity that we found between our main variables in Analysis 2, which complicate our interpretation of those models. Multicollinearity is an inescapable structural feature of this dataset; because the three major categories are mutually-exclusive and the number of datapoints is fixed at ten, an increase in one of the “inlast10” variables is inevitably going to be associated with a decrease in the others on average. The medium size category plot in Fig 1, for example, should not be interpreted to mean that optimal newcomer participation happens at recent participation of six users of each user type, as it is impossible for 18 messages to have been posted in the last 10 messages. Multicollinearity does not mean that the model is predictively inaccurate; the most common problem associated with multicollinearity is an inflation of the variance statistics [23], which can cause important factors to be shown as not statistically significant. However, the size of our dataset means that nearly all of the factors we tested remained statistically significant despite potentially inflated variance. Though the predictive power of the model is not significantly impacted, multicollinearity can make it challenging to apply theoretical explanations to model results. As described in the previous section, future work could attempt to replicate these results in other contexts and distinguish between the theoretical explanations we propose.

Broadly, this work establishes constructs relevant to in-the-moment decisions made by users about whether to participate for the first time, and future research and design can expand upon these constructs and use them as starting points for considering new approaches to recruitment.

Acknowledgements

Diya Yang’s contributions to this work were supported by a grant from Google. We would also like to thank Adam Seering for general and technical support in creating the scripts used for analysis.

REFERENCES

- [1] Jonathan Bishop. 2007. Increasing participation in online communities: A framework for human-computer interaction. *Computers in Human Behavior* 23, 4 (2007),

- 1881–1893. DOI:
<http://dx.doi.org/10.1016/j.chb.2005.11.004>
- [2] Susan L. Bryant, Andrea Forte, and Amy Bruckman. 2005. Becoming Wikipedian: Transformation of Participation in a Collaborative Online Encyclopedia. In *Proceedings of the 2005 International ACM SIGGROUP Conference on Supporting Group Work (GROUP '05)*. Association for Computing Machinery, New York, NY, USA, 1–10. DOI:
<http://dx.doi.org/10.1145/1099203.1099205>
- [3] Moira Burke, Elisabeth Joyce, Tackjin Kim, Vivek Anand, and Robert Kraut. 2007. Introductions and requests: Rhetorical strategies that elicit response in online communities. In *Communities and Technologies 2007*. Springer, New York, NY, USA, 21–39.
- [4] Judith Donath. 1999. Identity and Deception in the Virtual Community. In *Communities in Cyberspace* (1st ed.), Marc A Smith and Peter Kollock (Eds.). Routledge, London, UK, 27–58. DOI:
<http://dx.doi.org/10.1519/JSC.0b013e3181e4f7a9>
- [5] Rosta Farzan, Laura A. Dabbish, Robert E. Kraut, and Tom Postmes. 2011. Increasing Commitment to Online Communities by Designing for Social Presence. In *Proceedings of the ACM 2011 Conference on Computer Supported Cooperative Work (CSCW '11)*. ACM, New York, NY, USA, 321–330. DOI:
<http://dx.doi.org/10.1145/1958824.1958874>
- [6] Casey Fiesler, Shannon Morrison, R. Benjamin Shapiro, and Amy S. Bruckman. 2017. Growing Their Own: Legitimate Peripheral Participation for Computational Learning in an Online Fandom Community. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing (CSCW '17)*. Association for Computing Machinery, New York, NY, USA, 1375–1386. DOI:
<http://dx.doi.org/10.1145/2998181.2998210>
- [7] R Stuart Geiger. 2017. Beyond opening up the black box: Investigating the role of algorithmic systems in Wikipedian organizational culture. *Big Data & Society* 4, 2 (2017), 1–14.
- [8] Daniel Gros, Brigitta Wanner, Anna Hackenholt, Piotr Zawadzki, and Kathrin Knautz. 2017. World of Streaming. Motivation and Gratification on Twitch. In *Social Computing and Social Media. Human Behavior*, Gabriele Meiselwitz (Ed.). Springer International Publishing, Cham, 44–57.
- [9] Jamie A Gruman, Alan M Saks, and David I Zweig. 2006. Organizational socialization tactics and newcomer proactive behaviors: An integrative study. *Journal of vocational behavior* 69, 1 (2006), 90–104.
- [10] Aaron Halfaker, R Stuart Geiger, Jonathan T Morgan, and John Riedl. 2013. The rise and decline of an open collaboration system: How Wikipedia’s reaction to popularity is causing its decline. *American Behavioral Scientist* 57, 5 (2013), 664–688.
- [11] Aaron Halfaker, Aniket Kittur, and John Riedl. 2011. Don’t Bite the Newbies: How Reverts Affect the Quantity and Quality of Wikipedia Work. In *Proceedings of the 7th International Symposium on Wikis and Open Collaboration (WikiSym '11)*. Association for Computing Machinery, New York, NY, USA, 163–172. DOI:
<http://dx.doi.org/10.1145/2038558.2038585>
- [12] William A. Hamilton, Oliver Garretson, and Andruid Kerne. 2014. Streaming on Twitch: Fostering Participatory Communities of Play within Live Mixed Media. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '14)*. Association for Computing Machinery, New York, NY, USA, 1315–1324. DOI:
<http://dx.doi.org/10.1145/2556288.2557048>
- [13] Zorah Hilvert-Bruce, James T Neill, Max Sjöblom, and Juho Hamari. 2018. Social motivations of live-streaming viewer engagement on Twitch. *Computers in Human Behavior* 84 (2018), 58–67. DOI:
<http://dx.doi.org/10.1016/j.chb.2018.02.013>
- [14] Ruogu Kang, Laura Dabbish, and Katherine Sutton. 2016. Strangers on Your Phone: Why People Use Anonymous Communication Applications. In *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing (CSCW '16)*. Association for Computing Machinery, New York, NY, USA, 359–370. DOI:
<http://dx.doi.org/10.1145/2818048.2820081>
- [15] Joon Koh, Young-Gul Kim, Brian Butler, and Gee-Woo Bock. 2007. Encouraging participation in virtual communities. *Commun. ACM* 50 (2007), 68–73. DOI:
<http://dx.doi.org/10.1145/1216016.1216023>
- [16] Robert Kraut, Moira Burke, John Riedl, and Paul Resnick. 2011. The challenges of dealing with newcomers. In *Building Successful Online Communities: Evidence-Based Social Design*, Robert Kraut and Paul Resnick (Eds.). MIT Press, Cambridge, MA, USA, 179–230.
- [17] Xiao Ma, Nazanin Andalibi, Louise Barkhuus, and Mor Naaman. 2017. “People Are Either Too Fake or Too Real”: Opportunities and Challenges in Tie-Based Anonymity. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI '17)*. Association for Computing Machinery, New York, NY, USA, 1781–1793. DOI:
<http://dx.doi.org/10.1145/3025453.3025956>
- [18] Jennifer Marlow, Laura Dabbish, and Jim Herbsleb. 2013. Impression Formation in Online Peer Production: Activity Traces and Personal Profiles in Github. In *Proceedings of the 2013 Conference on Computer Supported Cooperative Work (CSCW '13)*. ACM, New York, NY, USA, 117–128. DOI:
<http://dx.doi.org/10.1145/2441776.2441792>

- [19] Blair Nonnecke and Jenny Preece. 2000. Lurker Demographics: Counting the Silent. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '00)*. Association for Computing Machinery, New York, NY, USA, 73–80. DOI : <http://dx.doi.org/10.1145/332040.332409>
- [20] Katherine Panciera, Aaron Halfaker, and Loren Terveen. 2009. Wikipedians Are Born, Not Made: A Study of Power Editors on Wikipedia. In *Proceedings of the ACM 2009 International Conference on Supporting Group Work (GROUP '09)*. ACM, New York, NY, USA, 51–60. DOI : <http://dx.doi.org/10.1145/1531674.1531682>
- [21] Jennifer Preece and Ben Shneiderman. 2009. The reader-to-leader framework: Motivating technology-mediated social participation. *AIS transactions on human-computer interaction* 1, 1 (2009), 13–32.
- [22] Catherine Ridings, David Gefen, and Bay Arinze. 2006. Psychological barriers: Lurker and poster motivation and behavior in online communities. *Communications of the association for Information Systems* 18, 1 (2006), 16.
- [23] Colleen Schwarz, Andrew Schwarz, and William C. Black. 2014. Examining the Impact of Multicollinearity in Discovering Higher-Order Factor Models. *CAIS* 34 (2014), 62.
- [24] Joseph Seering, Juan Pablo Flores, Saiph Savage, and Jessica Hammer. 2018. The Social Roles of Bots: Evaluating Impact of Bots on Discussions in Online Communities. *Proc. ACM Hum.-Comput. Interact.* 2, CSCW, Article Article 157 (Nov. 2018), 29 pages. DOI : <http://dx.doi.org/10.1145/3274426>
- [25] 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 (CSCW '17)*. ACM, New York, NY, USA, 111–125. DOI : <http://dx.doi.org/10.1145/2998181.2998277>
- [26] 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. DOI : <http://dx.doi.org/10.1177/1461444818821316>
- [27] Max Sjöblom, Maria Törhönen, Juho Hamari, and Joseph Macey. 2017. Content structure is king: An empirical study on gratifications, game genres and content type on Twitch. *Computers in Human Behavior* 73 (2017), 161–171.
- [28] Magdalena Szumilas. 2010. Explaining odds ratios. *Journal of the Canadian academy of child and adolescent psychiatry* 19, 3 (2010), 227.
- [29] Donghee Yvette Wohn, Guo Freeman, and Caitlin McLaughlin. 2018. Explaining Viewers' Emotional, Instrumental, and Financial Support Provision for Live Streamers. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18)*. Association for Computing Machinery, New York, NY, USA, Article Paper 474, 13 pages. DOI : <http://dx.doi.org/10.1145/3173574.3174048>
- [30] Diyi Yang, Robert Kraut, and John M. Levine. 2017. Commitment of Newcomers and Old-timers to Online Health Support Communities. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI '17)*. ACM, New York, NY, USA, 6363–6375. DOI : <http://dx.doi.org/10.1145/3025453.3026008>
- [31] Yunwen Ye and Gerhard Fischer. 2007. Designing for participation in socio-technical software systems. In *International Conference on Universal Access in Human-Computer Interaction*. Springer, New York, NY, USA, 312–321.
- [32] Bowen Yu, Xinyi Wang, Allen Yilun Lin, Yuqing Ren, Loren Terveen, and Haiyi Zhu. 2017. Out With The Old, In With The New?: Unpacking Member Turnover in Online Production Groups. *Proc. ACM Hum.-Comput. Interact.* 1, CSCW, Article 117 (Dec. 2017), 19 pages. DOI : <http://dx.doi.org/10.1145/3134752>

Table 14: Model showing decaying influence of previous messages as newer messages are posted.

	Coefficient	SE	z-score	p-value
Intercept	-3.83	0.006	-604.85	<0.001 ***
#subsinlast10	-0.04	0.001	-95.86	<0.001 ***
#botsinlast10	-0.01	0.001	-2.68	0.007 **
#newinlast10	0.36	0.002	221.98	<0.001 ***
1st prev	-0.14	0.010	-13.14	<0.001 ***
2nd prev	-0.11	0.010	-10.40	<0.001 ***
3rd prev	-0.08	0.010	-7.94	<0.001 ***
4th prev	-0.07	0.010	-6.50	<0.001 ***
5th prev	-0.07	0.010	-6.87	<0.001 ***
6th prev	-0.05	0.010	-5.31	<0.001 ***
7th prev	-0.03	0.010	-2.69	0.007 **
8th prev	-0.04	0.010	-4.42	<0.001 ***
9th prev	-0.05	0.010	-4.67	<0.001 ***
10th prev	-0.03	0.010	-3.16	0.001 **

APPENDIX A: CORRELATION MATRICES

The following are matrices showing correlations between variables used in the predictive models used in this work. All correlations between variables were small to insignificant.

Table 11: Correlation matrix for inter-variable correlations for small streams

	#regin last10	#modsin last10	#subsin last10	#newin last10
#reginlast10	1	-0.32	-0.82	0.03
#modsinlast10		1	-0.14	-0.08
#subsinlast10			1	-0.20
#newinlast10				1

Table 12: Correlation matrix for inter-variable correlations for medium streams

	#regin last10	#modsin last10	#subsin last10	#newin last10
#reginlast10	1	-0.42	-0.85	-0.12
#modsinlast10		1	0.03	-0.03
#subsinlast10			1	-0.11
#newinlast10				1

Table 13: Correlation matrix for inter-variable correlations for large streams

	#regin last10	#modsin last10	#subsin last10	#newin last10
#reginlast10	1	-0.33	-0.86	-0.11
#modsinlast10		1	-0.02	-0.03
#subsinlast10			1	-0.16
#newinlast10				1

APPENDIX B: DEMONSTRATING EXPECTED TIME SERIES BEHAVIOR

Prior to developing full models, we first performed analyses to show that observed statistical behaviors matched with expectations, namely that messages decay in influence over time as new messages are posted. If this were the case, we would expect the most recent message to be the most influential on subsequent behavior, and each message further in the past would be less influential.

Table 14 shows a basic example logistic regression model with the dependent variable being whether a subsequent message is posted by a first-time participant. In this case, the model shows decaying influence of messages sent by moderators in the small stream size category as new messages are posted, with the 1st prev through 10th prev variables indicating the most recent to least recent messages' influence and the "prior" variables indicating controls for other user types. The coefficient of the most recent message's user status (1st prev) has approximately four times the magnitude of the coefficient for a message ten messages prior (10th prev), showing the expected pattern.

Though we did perform these analyses for each user status type (moderators, subscribers, and bots) and each stream size category, we present only this model as an example for the sake of brevity. Each of the eight other models showed the same pattern with the notable exception that in models for the largest stream size category the second most recent message was more influential than the most recent. This is likely the result of the faster rate of message posting in these channels, where the most recent message may have been sent after the user decided to participate and began typing their comment.

APPENDIX C: FULL MODELS FOR LIKELIHOOD OF NEW PARTICIPANT MESSAGE

For the sake of simplicity, we omit interaction terms in Tables 5, 6, and 7 above. The following three tables show the full models.

Table 15: Model of likelihood of new participant message based on prior number of messages by moderators, subscribers, bots, and other new participants in small streams. Nagelkerke's $R^2 = 0.038$

	Coefficient	SE	z-score	p-value
Intercept	-4.548	0.089	-51.373	<0.001 ***
#modsinlast10 ²	-0.029	0.002	-11.903	<0.001 ***
#modsinlast10	0.302	0.028	10.777	<0.001 ***
#subsinlast10 ²	-0.015	0.001	-10.605	<0.001 ***
#subsinlast10	0.147	0.022	6.603	4.0 e-11
#reginlast10 ²	-0.017	0.001	-11.268	<0.001 ***
#reginlast10	0.270	0.023	11.834	<0.001 ***
#newinlast10 ²	-0.050	0.003	-15.898	<0.001 ***
#newinlast10	0.993	0.034	29.446	<0.001 ***
#modsinlast10*#subsinlast10	-0.034	0.003	-11.109	<0.001 ***
#modsinlast10*#reginlast10	-0.052	0.003	-15.935	<0.001 ***
#modsinlast10*#newinlast10	-0.048	0.005	-9.925	<0.001 ***
#subsinlast10*#reginlast10	-0.023	0.003	-8.793	<0.001 ***
#subsinlast10*#newinlast10	-0.045	0.004	-11.813	<0.001 ***
#reginlast10*#newinlast10	-0.044	0.004	-11.930	<0.001 ***

Table 16: Model of likelihood of new participant message based on prior number of messages by moderators, subscribers, bots, and other new participants in medium streams. Nagelkerke's $R^2 = 0.028$

	Coefficient	SE	z-score	p-value
Intercept	-3.889	0.007	-518.79	<0.001 ***
#modsinlast10 ²	-0.036	0.002	-24.05	<0.001 ***
#modsinlast10	0.250	0.009	26.90	<0.001 ***
#subsinlast10 ²	-0.004	0.000	-7.52	<0.001 ***
#subsinlast10	0.017	0.004	4.04	<0.001 ***
#reginlast10 ²	-0.076	0.002	-39.44	<0.001 ***
#reginlast10	0.610	0.011	54.41	<0.001 ***
#newinlast10 ²	-0.024	0.002	-14.50	<0.001 ***
#newinlast10	0.709	0.008	84.64	<0.001 ***
#modsinlast10*#subsinlast10	-0.027	0.002	-17.38	<0.001 ***
#modsinlast10*#reginlast10	-0.058	0.004	-14.39	<0.001 ***
#modsinlast10*#newinlast10	-0.082	0.004	-19.39	<0.001 ***
#subsinlast10*#reginlast10	-0.053	0.002	-26.88	<0.001 ***
#subsinlast10*#newinlast10	-0.025	0.002	-13.69	<0.001 ***
#reginlast10*#newinlast10	-0.114	0.005	-24.99	<0.001 ***

Table 17: Model of likelihood of new participant message based on prior number of messages by moderators, subscribers, bots, and other new participants in large streams. Nagelkerke's $R^2 = 0.026$

	Coefficient	SE	z-score	p-value
Intercept	-2.734	0.104	-26.24	<0.001 ***
#modsinlast10 ²	-0.006	0.008	-0.72	0.474
#modsinlast10	-0.435	0.079	-5.52	<0.001 ***
#subsinlast10 ²	-0.040	0.002	-17.38	<0.001 ***
#subsinlast10	0.282	0.032	8.84	<0.001 ***
#reginlast10 ²	-0.017	0.002	-10.56	<0.001 ***
#reginlast10	0.101	0.025	4.12	<0.001 ***
#newinlast10 ²	0.051	0.005	10.22	<0.001 ***
#newinlast10	-0.130	0.051	-2.54	0.011 *
#modsinlast10*#subsinlast10	-0.009	0.009	-1.03	0.303
#modsinlast10*#reginlast10	0.044	0.008	5.58	<0.001 ***
#modsinlast10*#newinlast10	0.076	0.012	6.15	<0.001 ***
#subsinlast10*#reginlast10	-0.073	0.003	-21.70	<0.001 ***
#subsinlast10*#newinlast10	0.018	0.006	2.98	0.003
#reginlast10*#newinlast10	0.028	0.005	5.33	<0.001 ***