# Successive Cohorts of Twitter Users
# Show Increasing Activity and Shrinking Content Horizons

FREDERIK WOLF

Potsdam Institute for Climate Impact Research (PIK), Germany

Humboldt University Berlin, Germany

SUNE LEHMANN

Technical University of Denmark, Denmark

Center for Social Data Science, University of Copenhagen, Denmark

PHILIPP LORENZ-SPREEN

Center for Adaptive Rationality,

Max Planck Institute for Human Development, Germany

The global public sphere has changed dramatically over the past decades: A significant part of public discourse now takes place on algorithmically driven platforms. Despite its growing importance, there is scant large-scale academic research on the long-term evolution of user behaviour on these platforms. Here, we evaluate the behaviour of 600,000 individual Twitter users between 2012 and 2019 and find empirical evidence for a cohort level acceleration of the way Twitter is used. Across time, we observe changing user-level behaviours: more tweets per time, denser interactions with others via retweets, and shorter content horizons, expressed as an individual's decaying autocorrelation of topics over time. We show that the change in usage patterns is not simply caused by a growing user base. While behaviour remains remarkably stable within each cohort over time, we relate these observations to changing compositions of new

Frederik Wolf: frederik.wolf@pik-potsdam.de
Sune Lehmann: sljo@dtu.dk
Philipp Lorenz-Spreen: lorenz-spreen@mpib-berlin.mpg.de
Date submitted: 2022-01-04

users with each new cohort containing increasingly active individuals. Our findings complement recent empirical work on social acceleration by tracking cohorts over time, controlling for cohort size, and analyzing their behavioural composition.

*Keywords: long-term data, cohort analysis, content horizon*

Year by year, the world is becoming more interconnected online (Hilbert and López, 2011), with news (Vosoughi et al., 2018), games (Morcos et al., 2021), and entertainment (Ribeiro et al., 2020) delivered to individuals via an increasing number of smartphones and computers worldwide (Taylor and Silver, 2019). The realization that this development may not be unequivocally beneficial for individuals and societies has spurred an active scientific and public debate (Zuboff, 2019; Aral, 2020; Hills, 2019; Bail et al., 2018; Allen et al., 2020; Allcott et al., 2020; Mosleh et al., 2021), while positive consequences of the growing connectivity can also be observed (Guess et al., 2019; Barberá et al., 2015; Yang et al., 2020; Boxell et al., 2017). Still, most aspects of the complex interplay between information technology, social interconnectedness, and human behaviour on the collective (Bak-Coleman et al., 2021), as well as on the individual level (Lorenz-Spreen et al., 2020) have yet to be empirically addressed through large-scale quantitative studies.

Here we want to shed light on one aspect of how social media can change people's within-platform activity over time. By studying individual users and looking beyond purely aggregate measures, our results complement recent empirical work on the acceleration of information behaviour. Our focus on individual-level statistics allows us to control for the number and type of users involved and to better understand both how much of the observed aggregate acceleration is accounted for by an average user and whether this acceleration is due to changes in the user base itself. These insights are important for designing online platforms that help users cope with the increasing overall speed and complexity of information dissemination.

The technological acceleration has been quantified in a variety of sectors, from ge-

nomic sequencing (Wetterstrand, 2013) to computing power (Moore, 1998) and transmission of information (Hilbert and López, 2011). The impact of such technological developments on the social sphere are more difficult to quantify (Wajcman, 2008). Gradually, however, more and more empirical hints are emerging. Recent work has provided evidence for acceleration of collective attention across various domains, including information search, communication, and entertainment (Lorenz-Spreen et al., 2019). These findings are supported by other empirical evidence for instances of acceleration of, for example, media consumption and production (Hutchins, 2011), the editing style of Hollywood movies (Cutting et al., 2011), the uptake of new concepts in books (Michel et al., 2011), the uptake of technological innovations (McGrath, 2013), and—particularly relevant to this study—information consumption on social media (Yang et al., 2020; Scharkow et al., 2020; Ford et al., 2021).

The majority of existing work, however, focuses on the aggregated level, leaving open the question of whether there is an accelerating trend when controlling for the number of participating individuals (with a few exceptions, e.g., surveys that point to a reduction in sustained attention in reading behaviour (Liu, 2005)). In other words, are the observed developments of acceleration driven simply by the fact that over time there are more and more people on social media participating in consumption or discourse—and their activity became more visible there? Or do people behave differently on these platforms now than they did just a few years ago? More specifically, are the observed aggregate changes caused by the same people who change their behaviour over time, or are the changes due to a distinct set of people on the platform now who are behaving differently than those who participated a few years ago?

We addressed the question from the perspective of changing behaviour on Twitter, using a longitudinal data set from which we sampled three times (March, April, and May 2019). Each sample consisted of 200,000 randomly chosen individual users and their activity (tweets and retweets). In the main text, we show the results for the sample of users based on tweets recorded in April 2019, but repeat our analysis on other samples to show the robustness of our findings. The choice of using the sample from April 2019 for the main manuscript is arbitrary and aside from slightly different user type compositions, all results from all samples agree quantitatively and qualitatively (for more details see section "Limitations and robustness test" and the Supplementary Materials). The data set enabled us to

measure user behaviour over 8 years (2012–2019), spanning a large fraction of the observable period of widespread social media adoption. Accordingly, we analyzed people's activity on Twitter from 2012 to 2019, aiming to understand differences between users who joined the platform at different points in time. We also explored the development of interactions with others and the amount of time users pay attention to topics.

Our work connects to recent efforts to describe and understand how the explosive and widespread uptake of social media changes long-term behaviour on these platforms, which has only recently become quantitatively accessible, despite the widespread use of social media over the last decade (Waller and Anderson, 2021; Alshaabi et al., 2021).

## Data set

We used a data set from Twitter's Decahose API that featured 10% of the Twitter traffic between January 1, 2012 and December 6, 2016 and 1% between December 7, 2016 and June 6, 2019. To perform a continuous analysis across the abrupt change in temporal resolution, we employed only 10% of the tweets in the first period, selected using random sampling (imitating the sampling mechanism of the Twitter API). Consequently, we considered 1% of all tweets in the period between January 2012 and June 2019. As this data was collected in real time, it included the activity of deleted user accounts; however, we expect that population to be very small due to our sampling technique.

The metadata related to each tweet consists of author ID, retweet ID, timestamp in seconds, and full tweet text. As the tweets and users were randomly sampled, there was no geographical or linguistic preference in the data set.

For our analysis, we set up three user samples, each containing the data from $200,000$ users. Users were selected randomly from all users who were active (at least one reported tweet) in March ($200,000$ users out of $28,883,037$), April ($200,000$ users out of $28,843,944$), and May 2019 ($200,000$ users out of $29,353,930$). Note that we first assembled all users who had tweeted and from there randomly selected users for our analysis in order to avoid biasing our data set towards the more active users in this period. The user samples were almost distinct; the pairwise overlap between the samples was less than 0.5%. Only two users appeared in all three data sets. This user selection allowed us to track individual user

activity over the whole time period (January 2012–June 2019) by considering all tweets from each selected user in the 1% subset. It also ensured that there was no drop-out of users in the studied period. This limited our analysis to users who remained on the platform until Spring 2019.

## Results

This type of data set and sampling, in combination with personal identifiers, ensures that active Twitter users can be recorded over the entire period and that the observed effects are not driven by people who had stopped using the platform. Our sampling strategy is biased towards more active users that have been using Twitter for longer, hence against more active users in younger cohorts (for an analysis of this bias, see Fig. S9 and the corresponding discussion). We deliberately chose a conservative sampling strategy for social acceleration. We did not use any other selection criteria, nor did we attempt to infer user characteristics; instead, we relied on the API to provide a sample that was randomized enough to observe general, long-term changes in behaviour[1]. Note that by considering a growing volume of overall tweets on Twitter, our sampling strategy underestimates the increase and can thus be understood as a lower limit for the actual acceleration that occurred (for details see section "Limitations and robustness tests" and the Supplementary Materials).

Given the relative sparseness of a typical individual's tweets (and that we are working with a 1% random sample of all tweets) and the complexity of individual trajectories, we could not meaningfully assess within-person trends over multiple years. However, we were able to associate people's average activity with their starting date on Twitter. Therefore, in this paper we focus on cohort-level developments and leave the study of within-person analysis to future research. For a first insight into the general trend we ran a simple linear regression between the individual starting date and the individual mean inter-event time (time between own tweets or retweets in the first active year). We found first indications of acceleration in the sample: The inter-event time decreased by approximately 6.5 hours per year between 2013 and 2018. For more detailed analyses, the heterogeneity of individual behaviour over time made it necessary to pool users in order to create meaningful average

---

[1]Decahose API documentation: https://developer.twitter.com/en/docs/twitter-api/enterprise/decahose-api/overview/decahose
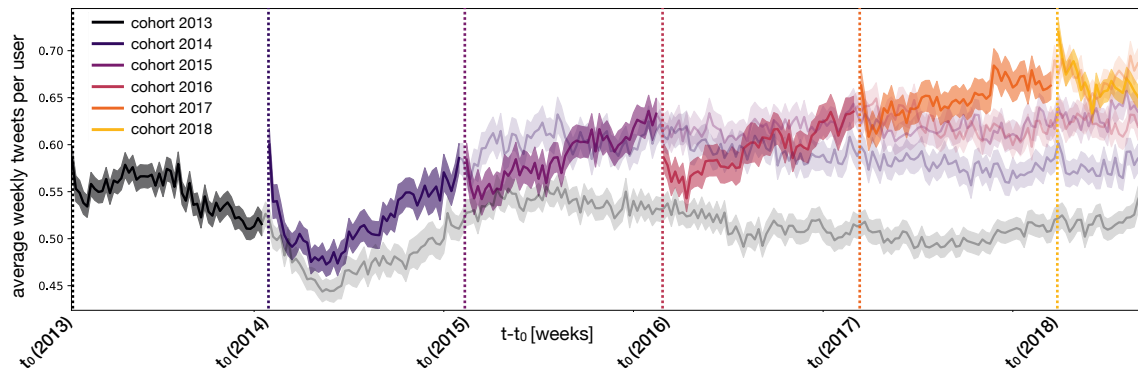
**Figure 1.** Average tweets per week and user within each cohort. The day of the first tweet is aligned per cohort, showing an average Twitter career starting at $t_0$ and lasting until 2019 for all users. First years are highlighted for comparison.

observational data at the systemic level. We chose two strategies for dividing users: by the year they began using Twitter and by user type. Forming groups of users based on the year they were first active on Twitter made it possible to divide them into eight *cohorts* that began using Twitter in the different years (2012, 2013, ...2019). Dividing users into groups of comparable user types according to their activity level made it possible to look beyond the simple cohort average and characterize the changing composition of behaviours across cohorts.

### *Users grouped by cohort*

We found that the tweeting behaviour of people who started using Twitter in 2013 was substantially different from that of people who started actively using Twitter more recently. To assess the change in user activity over time, we first identified the date $t_0$ of each user's first recorded tweet, then split the users into cohorts according to their starting year. To compare the behaviour of users who had spent the same amount of time on Twitter, we aligned all users from each cohort by setting their starting dates $t_0$ to an arbitrary but common date. This made it possible to approximately represent individual behaviour as an average Twitter experience within a cohort and to evaluate possible differences between cohorts.

We then calculated the mean number of tweets per week for all users at the same stage of Twitter use, for each week of the observation period. Aligning individual users created an offset of the trajectories within each cohort, which could be up to one year. To ensure that this offset did not extend beyond the active period of a user from our sample, we ignored the year 2019 for each cohort, leaving at least one year of buffer (and did not consider the 2019 cohort). Additionally, as we had no data from before 2012 and thus could not know whether users had joined Twitter before our observation period, we ignore the cohort from 2012. We interpret inactivity throughout 2012 as a proxy for not having joined Twitter earlier.

Fig. 1 shows the resulting average tweeting activity per user for all cohorts from 2013 to 2018—thereby going beyond merely illustrating growth in number of users. This visualization highlights both the dissimilarity between cohorts at the same stage of Twitter use and an offset in user activity (measured as average tweets per week) after multiple cohorts. We found a clear trend of increasing activity on the platform from one cohort to the next, whereas activity levels remained stable over long periods within each cohort of users (Fig. 1). These results show a lack of clear within-cohort development, but also growing activity from each cohort to subsequent cohorts. Users who were active on Twitter in 2013 continued using the platform in much the same way as when they started. They also connected preferentially to users who joined around the same time as they did. More generally, an average of 90% of all retweets occurred within a single cohort, despite increasing total interactions, indicating homophily among contacts of the same cohort. This finding mirrors previous findings of politically homophilous ties in social media networks (Mosleh et al., 2021) and the formation of topical groups therein (Cinelli et al., 2021).

Users' initial activity increased in each subsequent cohort and remained stable (and increasing) at a higher level of activity, especially for the cohorts after 2015. Furthermore, the cohorts demonstrated significantly different levels of activity at the end of our observation period.

What drives this change in behaviour, and what other dimensions are affected?
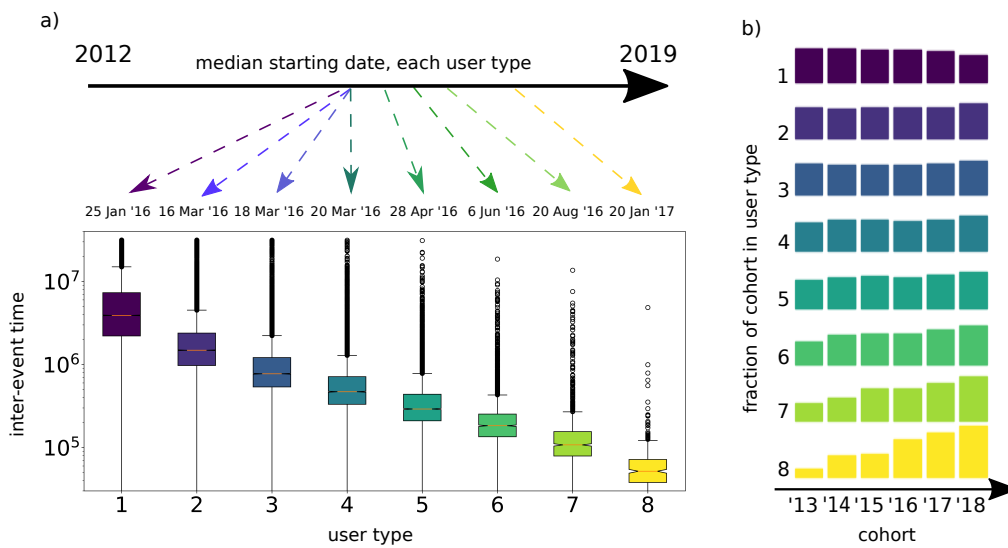
**Figure 2.** Composition and inter-event time distributions in the first year of user types. (a) User types are based on clustering along with the ratio of active days (at least one recorded tweet) versus total days on Twitter, labelled from 1 (least active) to 8 (most active). Top: Median starting dates (i.e., the first day of activity on Twitter $t_0$) for each user type. Bottom: Box plot representation of the individual inter-event time distribution for user types in seconds during their first year on Twitter. (b) Composition of user types per cohort.

### *Users grouped by activity*

To better understand the different roles users might play in this process, we grouped users according to activity level. Our analysis indicated that the growing activity on the cohort level stemmed from different compositions of user types, ranging from mostly passive spectators who tweeted only occasionally to extremely active users.

Activity was broadly distributed among users in each cohort (Fig. S1a, c, e). To disentangle the heterogeneity of users, we computed the ratio of active days (fraction of days with at least one recorded tweet) to total days (days between the individual starting date $t_0$ and May 30, 2019) for each user. We identified eight user types via a simple, unsupervised $k$-means-clustering (using the Python package *scikit-learn* (Pedregosa et al., 2011)) and labelled them from 1 (least active) to 8 (most active). This method allowed us to capture density variations in the data and thus set our bin edges in a data-driven way.

Figure 2a shows the resulting separation of the user types by their activity level (full distributions featured in Fig. S1b, d, f). Note that we here consider only the first year of activity (i.e., $t_0$) for each user in order to make a meaningful comparison. Therefore, the upper bound of the inter-event time equals $31,536,000$ seconds (one year). Additionally, due to the 1% random sample of tweets, the inter-event times reported here cannot be easily interpreted as such but rather serve as a proxy for individual activity.

Looking at user type and cohort, a trend becomes evident. Although we did not consider starting dates in the clustering that we used to define user types, we found a striking one-to-one correspondence between activity and the median starting date for individuals in each user type: The more recent the median starting date, the more active the user type (see Fig. 2a, top). In particular, less active user types were only marginally separated in terms of their median starting dates whereas there was a pronounced shift towards more recent median starting dates for the more active user types. This points to a change in individuals' user types over time. To investigate this finding and understand how user types are distributed within each cohort, we determined the proportions of user types in each cohort.

Figure 2b illustrates how the composition of user types changed over time. Perhaps the most striking observation is that the largest fractions of very active users are in later cohorts. The histograms show the proportions of user types in each cohort with respect to the absolute size of the cohort. Note that there are substantial differences between user types in terms of size (user type 1: $61,052$ users; user type 8: $445$ users; for all user types, see Fig. S1d) and slight differences in cohort size. Absolute numbers are not well represented in Fig. 2b because we show relative size increase.

The evolving composition of user cohorts on Twitter highlights that the most active user types grew, relative to the size of their group, more quickly than the less active user types did. Thus, the social acceleration observed on the collective level is likely to be driven by people who joined Twitter more recently and are using the platform differently than are users who have been on Twitter longer.

### *Activity relates to retweet interactions*

To unpack the observed behaviour of the increasingly large fractions of highly active users, we now turn to social interactions and how content changes for individuals. We found that highly active users were well connected with others both actively (retweeting) and passively (being retweeted).

To examine whether user types differed in respects other than their activity, we counted retweet interactions (aggregated over the full time range of $2012-2019$), consisting of both active (retweeting) and passive (being retweeted) retweets. We included retweet interactions outside of our random sample by using each retweet from the full Twitter data set whenever either the active or the passive user was from our sample. The analysis in Fig. 3 thereby consists of 648,880 unique users and a total of 5,963,284 interactions, resulting in 1,516,958 unique pairwise connections.

Figure 3 shows the distributions of the relative frequency of retweets per individual user $i$, for all eight user types. The increasing frequency of users with a high number of retweet interactions indicates that active user types were better connected on both measures. Hence, active users not only retweeted more actively (which can be expected due to their on average higher overall activity) but also elicited more activity from others. This is
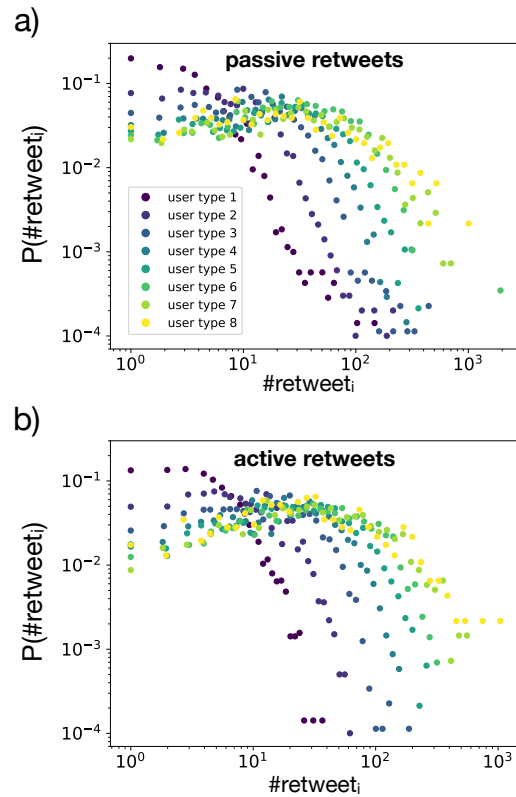
**Figure 3.** Distribution of the relative frequency $P(\#retweet_i)$ of the number of retweet interactions $\#retweet_i$ that one user $i$ has within each user type. (a) Distribution of the number of passive retweets (i.e., how often an individual user was retweeted by someone else) for each user type. (b) Distribution of the number of active retweets (i.e., how often an individual user retweeted someone else) for each user type.

also reflected in the relatively high reciprocity of the interactions (0.743). Reciprocity is the ratio of retweets that were followed by a reciprocal retweet at some later point. High reciprocity implies not only that activity increased over time (with more active users joining) but also that interactivity among users became more frequent. In other words, the observed trend towards higher activity did not occur in isolation but may be connected to a collective effect of mutual social acceleration and denser interaction among Twitter users—for example, when an elevated level of interaction leads to more content appearing in users' feeds. Individual activity can also speed up social acceleration: By virtue of their high activity, users become more central, filling each other's feeds and thereby collectively contributing to social acceleration.

### *Content horizons*

As a possible result of growing overall activity, the amount of time any individual topic appeared in people's tweets shrank over time. We call the amount of time that a topic tends to recur in a user's tweets that person's *content horizon.* We operationalized this notion as the autocorrelation of hashtags an individual uses over time.

To set tweets in the context of an ongoing discussion, users employ hashtags, a combination of the "#" symbol and keywords related to certain topics. Additionally, users can include URLs in their posts to link content to their activity. Did the way people interact with content change?

We found an increasing trend of sharing hashtags and URLs: Over our observation period, the numbers of hashtags per tweet and URLs per tweet almost doubled. While there was no clear difference among user types in terms of sharing URLs, more active user types tended to use more hashtags per tweet compared to less active user types (see Fig. S3). To understand the impact of the growing amount of interactions and content at an individual level, we measured the similarity of hashtags used over time and analyzed the content horizon, an adaptation of the concept of autocorrelation.

To do so, we pursued a nonstandard approach and employed Jaccard similarity to make it possible to compare the categorical hashtag data. Specifically, we defined the lagged

correlation between two tweets at $t$ and $t + \tau$ of an individual user $i$ as

$$A_i(\tau) = \frac{1}{N_{\text{match}}} \sum_{t=0}^{T-\tau} J(h_i(t), h_i(t + \tau)). \tag{1}$$

Here, $T - \tau$ determines the number of weeks that the user was active on Twitter and $J(A, B) = \frac{A \cap B}{A \cup B}$ represents the Jaccard similarity.

Zero correlations can be caused by different settings such as no activity, no used hashtags, and no common hashtags. To avoid an activity bias in the correlation measure, we excluded weeks in the computations during which one of the users $i$ and $j$ had been inactive and normalized only by the number of non-zero entries (i.e., $N_{\text{match}}$).

To estimate the individual content horizons of Twitter users, we computed the individual autocorrelation function $A_{i,j}(\tau)$ of hashtags used in tweets (Eq. 1) as a proxy for characteristic length of time users focused on a topic before moving on (see top panel of Fig. 4).

The decay of the autocorrelation was amplified in more recent cohorts (Fig. 4). This change indicates that users in more recent cohorts stopped tweeting about topics more quickly, potentially switching their focus to new topics. We call this development, which became stronger in successive cohorts of Twitter users, *shrinking content horizons.* Because we observed the qualitatively same results across user types (see Fig. S4), we assume that the shrinking content horizon is connected to rising activity on the part of individual users.

## Limitations and robustness tests

Here we summarize the limitations of our data and the tests conducted to ensure that the conclusions we draw in the Discussion section are indeed valid. For more extensive explanations, see the Supplementary Materials.

First, we selected three user samples, each containing $200,000$ users who were active in March, April, and May 2019, respectively. We established that our results based on the April 2019 sample are stable for other user samples obtained in a distinct period by obtaining
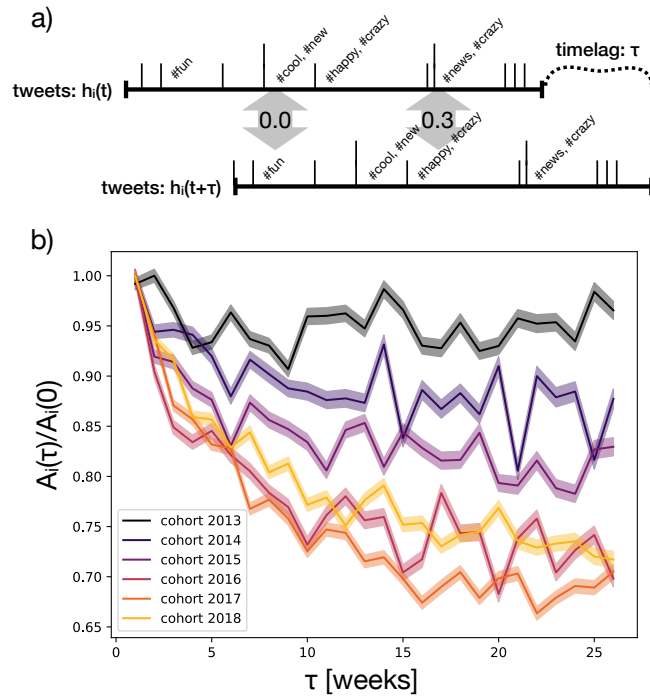
**Figure 4.** Autocorrelation of hashtags. (a) Illustration of the topical autocorrelation $A_i(\tau)$, via the relative overlap of hashtags between a shifted copy of one user's tweeting trajectory. (b) The resulting autocorrelation as a function of an increasing time lag for the cohorts of Twitter users. Values exclude the point without time lag $\tau = 0$ and are relative to the maximum autocorrelation value for better comparison.

the same results from analyzing the activity of users in the other two samples separately. We show the inter-event times distributions of user cohorts and user types for all three samples in Fig. S1 and present the cumulative age distributions (indicating that more recent users are more active; see Fig. 2) of the other two samples in Fig. S2. We illustrate the increasing usage of hashtags and URLs, which is qualitatively similar for all samples, in Fig. S3. To further check that we did not substantially bias our results by using the 1% stream of the Twitter data (Pfeffer et al., 2018), we briefly analyzed the full 10% sample. In Fig. S5 and Fig. S6 we present evidence that the results—increasing activity and more activity from recent users—are indeed qualitatively the same for the 10% sample from 2012–2016. The robustness of our observations across three randomly chosen samples and across user types make us confident in the findings, although we cannot fully exclude other overarching biases in Twitter's API. This issue highlights the problems of corporate data access (Pfeffer et al., 2018). Additionally, the absolute quantitative results should be treated with some caution as they depend heavily on the resampling. In particular, the inter-event time (and the tweets per week) would certainly be shorter (higher) given a full data set. Whereas some of these numbers can be at least compared to each other to estimate the relative increase (tweets per week per user have almost doubled), many absolute measures of user activity are difficult to estimate given our incomplete data set.

As a further proof of the significance of our results and in particular the differences between user cohorts and user types in the user sample we discuss in the main manuscript, we randomly selected $10 \times 10,000$ users of this sample and computed their inter-event time distribution (see Fig. S7). All 10 distributions collapsed on one common, indicating a significant difference between user cohorts and user types.

To confirm that our results were not driven by automated activity, we compared highly active users to a randomly chosen set of users. To measure repetitive postings we used Shannon information (Chu et al., 2012; Cresci, 2020) to quantify the complexity of the shared information. We found little difference in the distribution of complexity of tweets (on the word level) between very active user types and the random sample across all type (see Fig. S8a,b).

There are good reasons to believe that sampling users who were still active at the

end of the study period and tracking their activity back introduces selection bias. Although considering users who dropped out might influence the average activity of cohorts and user types, the sampling does not strongly affect our conclusion: Users who dropped out may have behaved systematically differently, but because they were more likely to be from early cohorts, and considering the finding that they were less active on Twitter before they dropped out, the implication is that differences in activity may be higher than reported here (see Fig. S9a,b and the corresponding explanations regarding selection and survivorship bias).

Another factor influencing the observed activity is the growing tweet volume on Twitter. As we based our analysis on a 1% random sample, we expected that given a constant tweeting rate of a single user, the probability of finding that user's tweets decreases over time. We sampled the user IDs randomly and independent of their individual activity during the observation period, but importantly all users were sampled at the same time of their activity, namely the month in 2019 at the end of the observation period. Consequently, the relatively constant tweeting activity of different cohorts could be interpreted as an increase in activity if the background activity had risen and thereby the probability of sampling a tweet had decreased.

## Discussion

We found that the average activity per user on Twitter increased year over year. We quantify this as a cohort effect and argue that this change can be explained by changing compositions of user types joining Twitter over time. Changes that are expressed in individual trajectories of Twitter user are more difficult to measure due to the diversity of user type patterns. Because there is an increase in not only the total number of activities on Twitter (Lorenz-Spreen et al., 2019), but also the activity per user, the general development is likely to be due to actual differences in people's behaviour, and not merely to a growing user base. The change in composition of user types is accompanied by growing connectivity among users through retweets, an increasing number of hashtags and URLs being shared, and shrinking content horizons (here operationalized as quickly decaying hashtag autocorrelations). Our findings can be related to an overarching sociological concept known as *social acceleration* (Rosa, 2013). Social acceleration is the interplay between technological

acceleration, acceleration of social change, and acceleration of the pace of life (Rosa, 2003).

Although we cannot draw causal conclusions from our analyses, the combination of observations allows us to offer an explanation of the factors that are likely to contribute to an individual's experience of social acceleration on social media platforms: The growing fraction of highly active users in each successive cohort combined with the pronounced tendency for tweets to be retweeted within a single cohort leads to more interactions with other users, likely resulting in more and more active users filling up their peers' feeds with content. This development is accompanied by an increasing proportion of hashtags and URLs in tweets over time. However, here we can only report observational evidence. In order to understand the underlying mechanisms behind our observations, researchers urgently need both better data access from platforms (including information about the platforms' sorting algorithms, design choices, and exposure, as well as about observed behaviour (Pasquetto et al., 2020)), and experimental control over some of the platform's features (e.g., the news feed). Only with variations between user groups can causal mechanisms be identified in the future—for example, how a news feed algorithm or retweet function may drive social acceleration, where it is possible to draw from methods of causal inference. Access to a complete data set of individual user behaviour would also make it possible to go beyond our cohort analysis and to examine how individuals change their platform use over time independent from overall tweet volume effects.

Although the amount of content on Twitter is increasing, the amount of information on a Twitter feed that a person can keep up with is finite—a simple consequence of the amount of information that fits on the screen in combination with users' attentional capacities (Hills, 2019). Users, therefore, typically encounter an overabundance of content (Bawden and Robinson, 2009). The inability to keep up with everything that appears on their Twitter feed may explain higher turnover rates of topics in each subsequent cohort. We quantified this development by measuring hashtag autocorrelation and confirmed the ever faster decay of individual content horizons. This points to a possible behavioural response to the growing information abundance: trading breadth for depth in their information behaviour (Carr, 2010).

Our data and methods come with several limitations. As the data only includes ac-

tive behaviour on Twitter it has little to say about the mechanisms at play that go beyond our analysis of interactions and content. For example, the data do not contain information about exposure (i.e., what people experienced before they tweeted) on and off Twitter, or events outside of Twitter that drive activity exogenously (Burton et al., 2021). We aimed to exclude the possibility of a predominant presence of automated accounts via a complexity analysis of content from very active users, showing that their posts are not overly repetitive. Although unobserved factors are important to better understand our findings, they do not trivialise our results; on the contrary, external factors are important to complement our findings and understand the factors that drive the broader development of social acceleration. For example, changes in the relative fraction of user types within each cohort over time are likely an indication that the platform itself is changing: Platform design choices may be altering how Twitter users are motivated to interact (Brady et al., 2020; Lorenz-Spreen et al., 2020; Kozyreva et al., 2020; Bak-Coleman et al., 2021). The interplay and potential mutual influence of user behaviour and platform design adds another layer of complexity on long-term field studies on social media. Other potential mechanisms outside the scope of our study include the increased professionalization and agenda-setting purposes of social media usage (Barberá et al., 2019) and an increasing migration of offline contacts to social media. Future research could also aim to connect the observed developments of acceleration with attentional bottlenecks and the success of the spread of negative, emotional, or hostile content on social media (Rathje et al., 2021; Acerbi, 2021; Brady et al., 2017; Alvarez et al., 2015). To determine the drivers behind these potentially unintended developments in public discourse—to distinguish between algorithmic curation, amplified human tendencies, and societal developments—future research on Twitter and other platforms is necessary.

Individuals interact with Twitter differently over time and influence each other in the process. Our work provides an empirical starting point highlighting the need to quantify these complex but important relationships further. Ultimately, the ability to quantify the directions in which the interplay of human behaviour, technological advancement, and corporate interests drive online behaviour and discourse would help society actively shape online discourse and would identify measures to promote a more deliberate online experience.

## Acknowledgments

## References

Acerbi, A. (2021). From storytelling to facebook. content biases when retelling or sharing a story.

Allcott, H., Braghieri, L., Eichmeyer, S., and Gentzkow, M. (2020). The welfare effects of social media. *American Economic Review*, 110(3):629–76.

Allen, J., Howland, B., Mobius, M., Rothschild, D., and Watts, D. J. (2020). Evaluating the fake news problem at the scale of the information ecosystem. *Science Advances*, 6(14):eaay3539.

Alshaabi, T., Dewhurst, D. R., Minot, J. R., Arnold, M. V., Adams, J. L., Danforth, C. M., and Dodds, P. S. (2021). The growing amplification of social media: Measuring temporal and social contagion dynamics for over 150 languages on twitter for 2009–2020. *EPJ data science*, 10(1):15.

Alvarez, R., Garcia, D., Moreno, Y., and Schweitzer, F. (2015). Sentiment cascades in the 15m movement. *EPJ Data Science*, 4(1):1–13.

Aral, S. (2020). *The Hype Machine: How Social Media Disrupts Our Elections, Our Economy, and Our Health—and How We Must Adapt.* Currency.

Bail, C. A., Argyle, L. P., Brown, T. W., Bumpus, J. P., Chen, H., Hunzaker, M. F., Lee, J., Mann, M., Merhout, F., and Volfovsky, A. (2018). Exposure to opposing views on social media can increase political polarization. *Proceedings of the National Academy of Sciences*, 115(37):9216–9221.

Bak-Coleman, J. B., Alfano, M., Barfuss, W., Bergstrom, C. T., Centeno, M. A., Couzin, I. D., Donges, J. F., Galesic, M., Gersick, A. S., Jacquet, J., et al. (2021). Stewardship of global collective behavior. *Proceedings of the National Academy of Sciences*, 118(27):e2025764118.

Barberá, P., Casas, A., Nagler, J., Egan, P. J., Bonneau, R., Jost, J. T., and Tucker, J. A. (2019). Who leads? who follows? measuring issue attention and agenda setting by legislators and the mass public using social media data. *American Political Science Review*, 113(4):883–901.

Barberá, P., Jost, J. T., Nagler, J., Tucker, J. A., and Bonneau, R. (2015). Tweeting from left to right: Is online political communication more than an echo chamber? *Psychological science*, 26(10):1531–1542.

Bawden, D. and Robinson, L. (2009). The dark side of information: overload, anxiety and other paradoxes and pathologies. *Journal of information science*, 35(2):180–191.

Boxell, L., Gentzkow, M., and Shapiro, J. M. (2017). Greater internet use is not associated with faster growth in political polarization among us demographic groups. *Proceedings of the National Academy of Sciences*, 114(40):10612–10617.

Brady, W. J., Crockett, M., and Van Bavel, J. J. (2020). The mad model of moral contagion: The role of motivation, attention, and design in the spread of moralized content online. *Perspectives on Psychological Science*, 15(4):978–1010.

Brady, W. J., Wills, J. A., Jost, J. T., Tucker, J. A., and Van Bavel, J. J. (2017). Emotion shapes the diffusion of moralized content in social networks. *Proceedings of the National Academy of Sciences*, 114(28):7313–7318.

Burton, J. W., Cruz, N., and Hahn, U. (2021). Reconsidering evidence of moral contagion in online social networks. *Nature Human Behaviour*, pages 1–7.

Carr, N. (2010). *The shallows: What the Internet is doing to our brains*. WW Norton & Company.

Chu, Z., Gianvecchio, S., Wang, H., and Jajodia, S. (2012). Detecting automation of twitter accounts: Are you a human, bot, or cyborg? *IEEE Transactions on dependable and secure computing*, 9(6):811–824.

Cinelli, M., De Francisci Morales, G., Galeazzi, A., Quattrociocchi, W., and Starnini, M. (2021). The echo chamber effect on social media. *Proceedings of the National Academy of Sciences*, 118(9):e2023301118.

Cresci, S. (2020). A decade of social bot detection. *Communications of the ACM*, 63(10):72–83.

Cutting, J. E., Brunick, K. L., DeLong, J. E., Iricinschi, C., and Candan, A. (2011). Quicker, faster, darker: Changes in Hollywood film over 75 years. *Iperception*, 2:569–576.

Ford, T., Krohn, R., and Weninger, T. (2021). Competition dynamics in the meme ecosystem. *arXiv preprint arXiv:2102.03952*.

Guess, A., Nagler, J., and Tucker, J. (2019). Less than you think: Prevalence and predictors of fake news dissemination on facebook. *Science advances*, 5(1):eaau4586.

Hilbert, M. and López, P. (2011). The world's technological capacity to store, communicate, and compute information. *science*, 332(6025):60–65.

Hills, T. T. (2019). The dark side of information proliferation. *Perspectives on Psychological Science*, 14(3):323–330.

Hutchins, B. (2011). The acceleration of media sport culture: Twitter, telepresence and online messaging. *Information, communication & society*, 14(2):237–257.

Kozyreva, A., Lewandowsky, S., and Hertwig, R. (2020). Citizens versus the internet: Confronting digital challenges with cognitive tools. *Psychological Science in the Public Interest*, 21(3):103–156.

Liu, Z. (2005). Reading behavior in the digital environment: Changes in reading behavior over the past ten years. *Journal of documentation*.

Lorenz-Spreen, P., Lewandowsky, S., Sunstein, C. R., and Hertwig, R. (2020). How behavioural sciences can promote truth, autonomy and democratic discourse online. *Nature human behaviour*, 4(11):1102–1109.

Lorenz-Spreen, P., Mønsted, B. M., Hövel, P., and Lehmann, S. (2019). Accelerating dynamics of collective attention. *Nature communications*, 10(1):1–9.

McGrath, R. (2013). The pace of technology adoption is speeding up. *Harvard Business Review*, 25.

Michel, J.-B., Shen, Y. K., Aiden, A. P., Veres, A., Gray, M. K., Pickett, J. P., Hoiberg, D., Clancy, D., Norvig, P., Orwant, J., et al. (2011). Quantitative analysis of culture using millions of digitized books. *Science*, 331:176–182.

Moore, G. E. (1998). Cramming more components onto integrated circuits. *Proceedings of the IEEE*, 86(1):82–85.

Morcos, M., Stavropoulos, V., Rennie, J. J., Clark, M., and Pontes, H. M. (2021). Internet gaming disorder: compensating as a draenei in world of warcraft. *International Journal of Mental Health and Addiction*, 19(3):669–685.

Mosleh, M., Martel, C., Eckles, D., and Rand, D. G. (2021). Shared partisanship dramatically increases social tie formation in a twitter field experiment. *Proceedings of the National Academy of Sciences*, 118(7).

Pasquetto, I. V., Swire-Thompson, B., Amazeen, M. A., Benevenuto, F., Brashier, N. M., Bond, R. M., Bozarth, L. C., Budak, C., Ecker, U. K., Fazio, L. K., et al. (2020). Tackling misinformation: What researchers could do with social media data. *The Harvard Kennedy School Misinformation Review*, 1.

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., et al. (2011). Scikit-learn: Machine learning in python. *the Journal of machine Learning research*, 12:2825–2830.

Pfeffer, J., Mayer, K., and Morstatter, F. (2018). Tampering with twitter's sample api. *EPJ Data Science*, 7(1):50.

Rathje, S., Van Bavel, J. J., and van der Linden, S. (2021). Out-group animosity drives engagement on social media. *Proceedings of the National Academy of Sciences*, 118(26).

Ribeiro, M. H., Ottoni, R., West, R., Almeida, V. A. F., and Meira, W. (2020). Auditing radicalization pathways on YouTube. In Hildebrandt, M. and Castillo, C., editors, *FAT\* 20: Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, pages 131–141, New York, NY. Association for Computing Machinery.

Rosa, H. (2003). Social acceleration: ethical and political consequences of a desynchronized high–speed society. *Constellations*, 10(1):3–33.

Rosa, H. (2013). *Beschleunigung: Die Veränderung der Zeitstrukturen in der Moderne [Social Acceleration: A New Theory of Modernity]*. Columbia University Press.

Scharkow, M., Mangold, F., Stier, S., and Breuer, J. (2020). How social network sites and other online intermediaries increase exposure to news. *Proceedings of the National Academy of Sciences*, 117(6):2761–2763.

Taylor, K. and Silver, L. (2019). Smartphone ownership is growing rapidly around the world, but not always equally. *Pew Research Center*.

Vosoughi, S., Roy, D., and Aral, S. (2018). The spread of true and false news online. *Science*, 359(6380):1146–1151.

Wajcman, J. (2008). Life in the fast lane? towards a sociology of technology and time. *The British journal of sociology*, 59(1):59–77.

Waller, I. and Anderson, A. (2021). Quantifying social organization and political polarization in online platforms. *Nature*, 600(7888):264–268.

Wetterstrand, K. A. (2013). Dna sequencing costs: data from the nhgri genome sequencing program (gsp).

Yang, T., Majó-Vázquez, S., Nielsen, R. K., and González-Bailón, S. (2020). Exposure to news grows less fragmented with an increase in mobile access. *Proceedings of the National Academy of Sciences*, 117(46):28678–28683.

Zuboff, S. (2019). *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power.* Profile Books.

## Online Appendix

In the main manuscript, we have based our analysis on the tweeting behavior of 200,000 users in who we randomly selected among all users who at least posted a single tweet on Twitter in April 2019. In this supporting information we do not only show figures regarding specific findings of this user group but also illustrate the robustness of the results by showing the raw data and comparing the outcomes with two other independently obtained data sets of the same size. These two other data sets are based on user activity (users randomly sampled from all users who tweeted) in March 2019 and May 2019, respectively. In the following we refer to the sample from March 2019 as sample 1, to the sample from April 2019 (used for the main analysis) as sample 2 and to the sample from May 2019 as sample 3.

## Inter-event times of user cohorts and user types

### *Cohorts*

In the main manuscript, we showed the rising activity of user cohorts in Figure 1. We based this measure on the cohort-mean number of tweets per week. A measure that is directly related to the user activity is the inter-event time. In Fig. S5a,c,e we show the inter-event times for all cohorts of the data sets based on user activity in March 2019 (a), April 2019 (c) and May 2019 (e). All three panels depict the increasing share of short (sub-daily) inter-event times and, thus, the rising activity. For all data sets there is an observable change towards shorter inter-event times for cohorts with users having joined Twitter more recently. This results do not only further confirm the findings stated in the main manuscript but underline its robustness as we find comparable results for the other two data sets.

### *User types*

In the main manuscript we spilt the users into user types of distinct activity by clustering the users by means of the fraction of active days. In Figure 2a of the main manuscript we confirmed, that the distributions of the inter-event times of the user types differ regarding their quantiles. In Fig. S5b,d,f we show the inter-event times of the distinct

user types for all three data sets (panel b shows the distribution from the data set studied in the main manuscript). The results confirm the robustness of the results as all three data sets show the same characteristic differences between the user types and agree regarding the overall distribution.
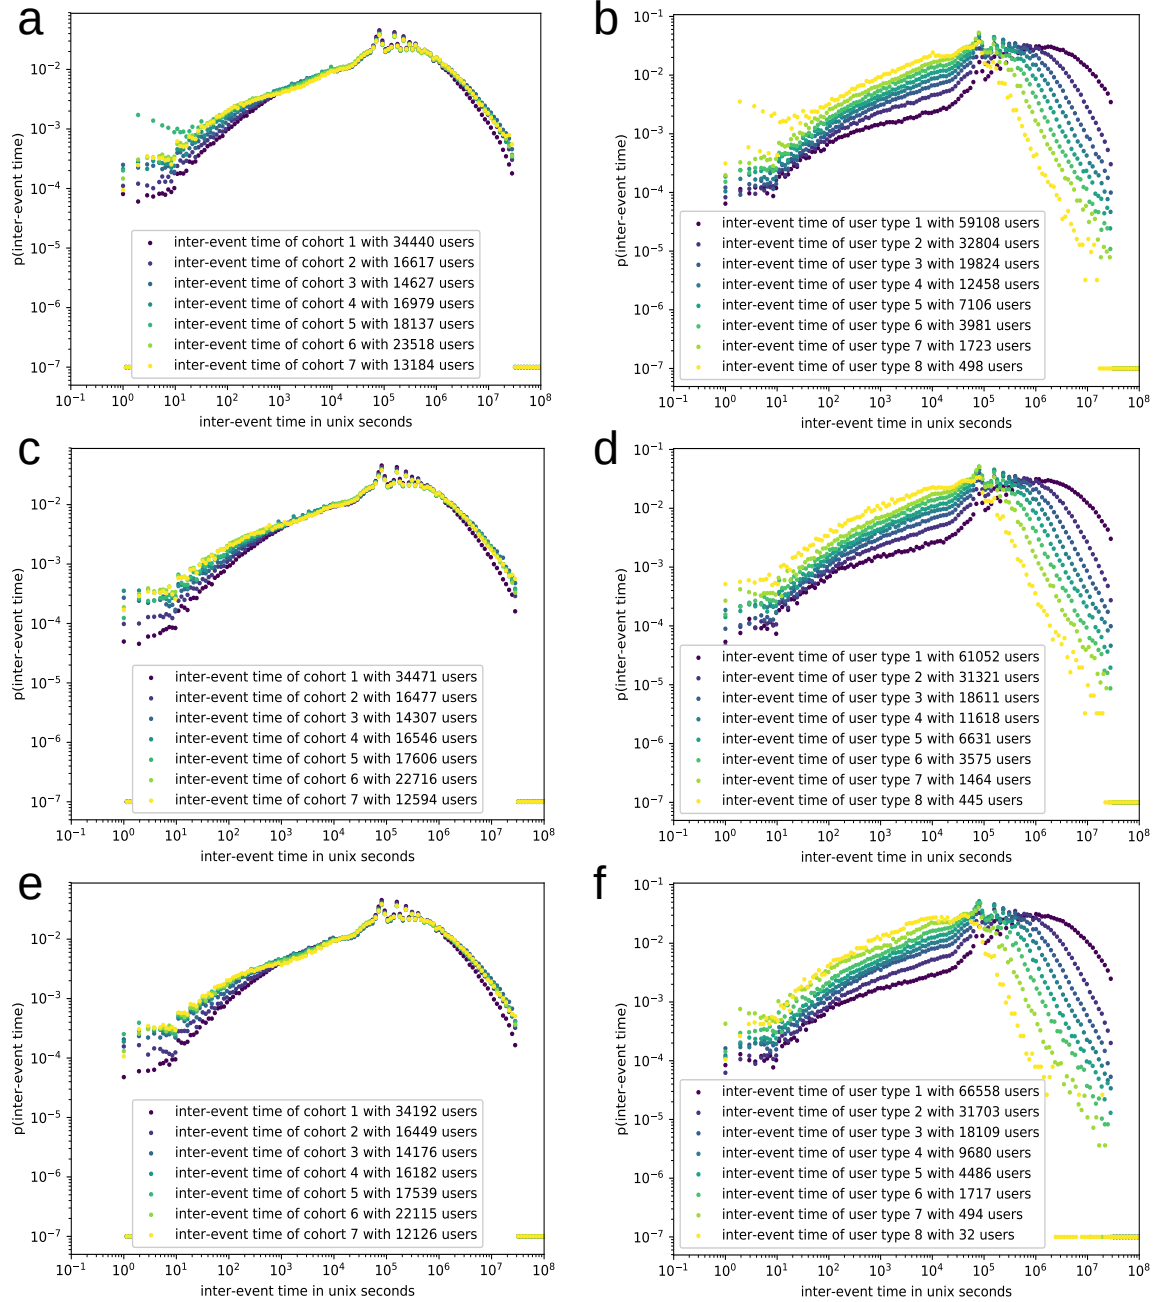
**Figure 5.** Inter-event times for cohorts (a,c,e) and user-types (b,d,f) for all three samples (top to bottom).
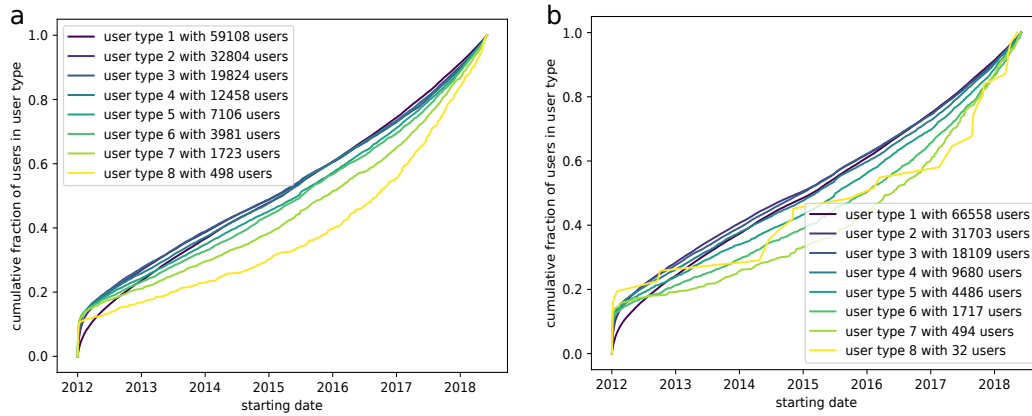
**Figure 6.** Cumulative age distribution of users in the different user-types in sample 1 (a) and sample (3). Note that we here include users from year 2012 (which we have not included in the calculations for Figure 2 in the main manuscript. Panels (c)-(f) show the average hashtags (c,d) and URLs (e,f) for sample 1 and sample 3.

### Median starting date and url/hashtag use

As an additional feature to the systematic offset between the user types, we studied the median starting date of the user types in Figure 2 of the main manuscript. To underline the robustness of our identified correspondence between user type activity and median starting date we show the cumulative starting date distribution for the two other data samples in Fig. S6a,b. For both samples we confirm that the user types exhibit a clear correlation between median starting date and activity as well as the increasing fraction of recent users in the more active user types.

The general increase of user activity goes in hand with a simultaneous increase of hashtag and URL usage. For the all three samples we show the average numbers of URLs and hashtags per tweet in Fig. S7.
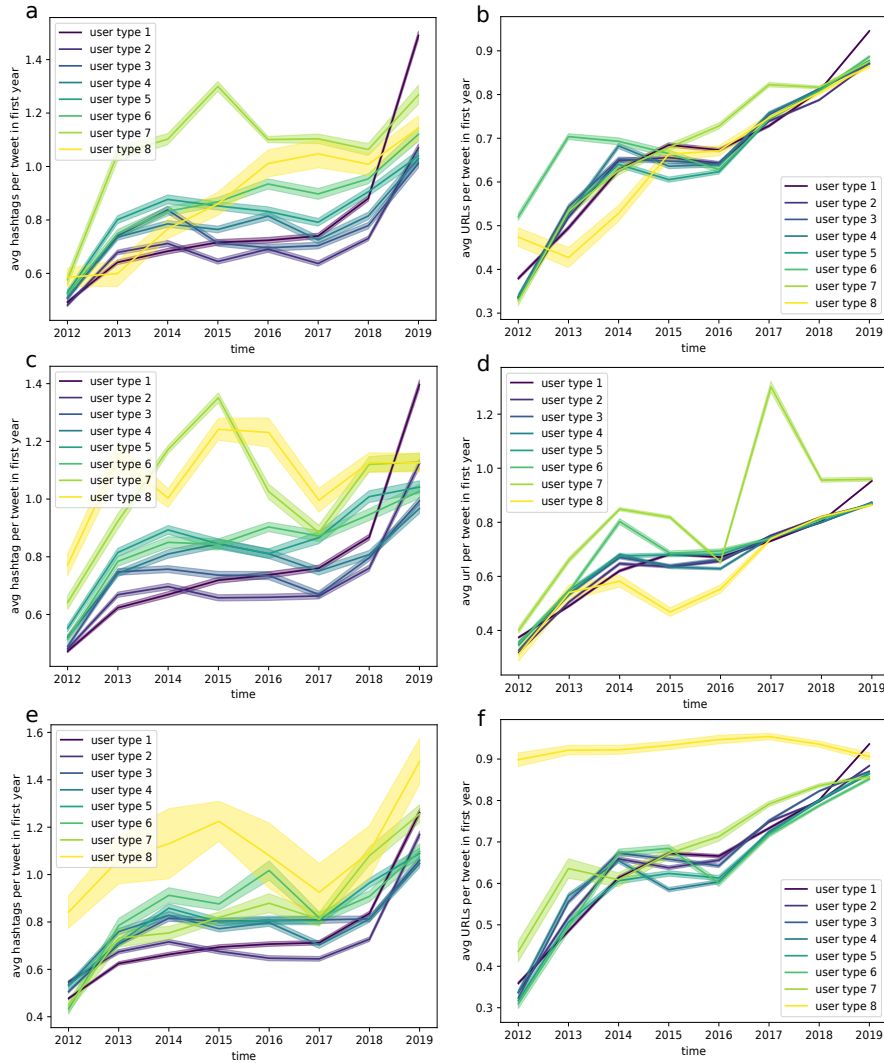
**Figure 7.** Average number of hashtags (a,c,e) and URLs (b,d,f) per user-type in sample 1 to 3 (top to bottom). Shaded areas indicate standard deviation.
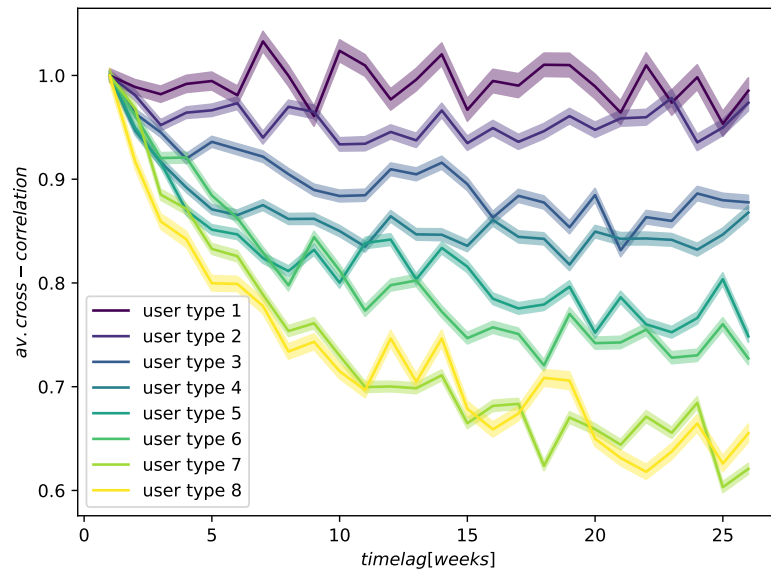
**Figure 8.** Content horizon by user types for user types in sample 2.

### Content horizons by user types

As mentioned in the main manuscript, we observe a shrinking content horizon not only for the cohorts but also for the user types (with a even stronger decaying autocorrelation function). To illustrate that, we show the content horizon for the user types of sample 2 in Fig. S8.
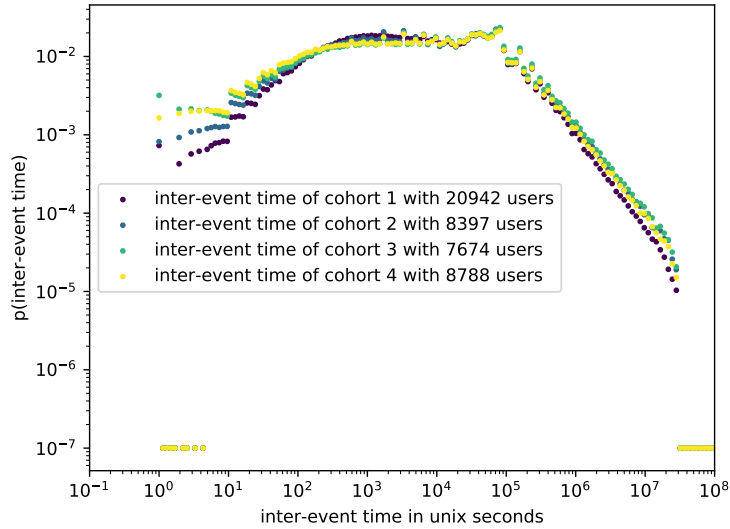
**Figure 9.** Inter-event time distribution using the full 10% data set for users in sample 2.

## Analysis of full 10% data set between 2012 and 2016

Additionally, we want to shed light on the downsampling procedure which enabled us to compare and contrast Twitter user behavior over the full study period. We are aware that reducing the density of the tweets can substantially affect our results. For assuring that the implications our our analysis would have been the same utilizing a data set with a larger fraction of tweets, we utilize the full 10% data set between 2012 and 2016. For this data set, we conducted the same analysis as for the whole data set. In Fig. S9, we show the inter-event time distribution of the four cohorts of this data set. Although the differences are not as striking as for the whole data set (which is caused by the shorter period covered), the same trend of an increasing share of shorter inter-event times can be observed. In addition, we obtained 4 user types of differing activity employing the same procedure as for the three other data sets (k-means clustering based on ratio of active days). Our results indicate the described observation: The relative fraction of the more recent users increases in when comparing less active and more active user types (see Fig. S10.
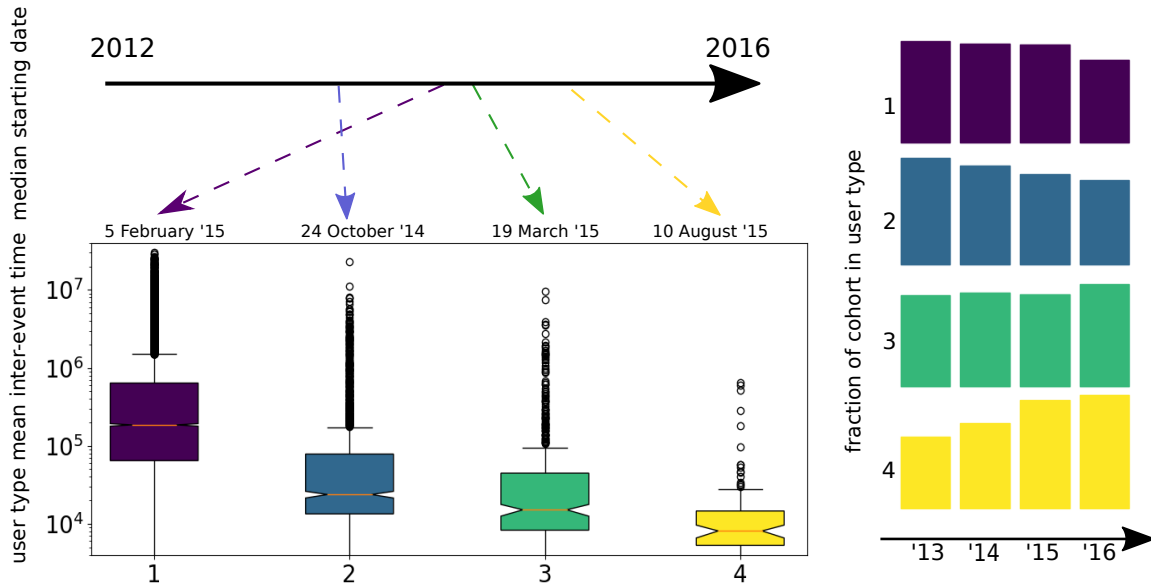
**Figure 10.** Composition and inter-event time of user types. a) User-mean inter-event time during the first year on Twitter and median starting date for each user type. b) Cohort compositions of user types over time. Both is shown using the full 10% data set for users in sample 2.

### Stability test with random users

To assure that the differences observed between the cohorts and user-types has not been caused by coincidence, we selected random samples of 10000 users from sample 2 and computed the inter-event times of each user to analyze the sample-wise inter-event time distribution. Note that we only include the inter-event times of all users which have tweeted for a full consecutive year and only take first-year inter-event times into account. Therefore, we exclude $\sim \frac{1}{3}$ of the random samples. As indicated in Fig. S11 all distributions collapse on one common and there is no systematic difference between the sub-samples. This is not only a confirmation that observing differences as shown in the previous figures is very unlikely. We also show that already a sample size of $\sim 7000$ users allows for obtaining the full typical distribution of inter-event times. We consider that as an additional validation of the significance of our results.
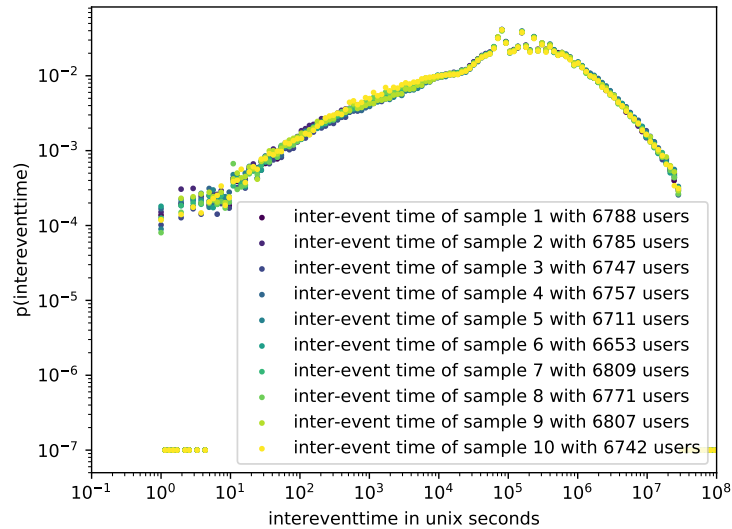
**Figure 11.** Inter-event time distribution of ten random user samples drawn from sample 2.

## Bot activity

In our analysis, we showed that the general increase of activity on Twitter is related to a growing share of the more active user types. On suspect might be that these highly active users are mostly bots tweeting nonsense at a high rate and, thus, contribute to the measured increase of activity. For comparing the content of highly active users with normal users, we randomly sampled 1000 users from the most active user types of sample 2 (200 of the most active user type, 400 of the second and third most active user types) and 1000 random users of sample 2. As we are only analyzing 1% of all tweets standard approaches like measuring the (unnatural) temporal distance between tweets does not easily work. As a first naive approach, we analyzed content repetition. Repeating content extensively can be associated with bot activity. 737 of 1000 users of the active user types have never repeated a post and 950 have more than 90% original posts. In the control group of 1000 random users we found that 847 users have never repeated a post and 985 users have more than 90% original posts. As the great majority of highly active users posts original content on Twitter, we confirm that there are no direct implications of bots significantly biasing the

results.

To complement this basic analysis with a more sophisticated approach, we decided to measure the information of words that each user has included in their tweets. Therefore, we first assembled all tweets of each individual user. To measure the complexity of the words this user has posted over the whole time this user has been active on Twitter, we have then quantified the information by considering each word as a single entry of a dictionary. Assuming that we have a dictionary of user $i$, $\mathcal{D}_i = \{\mathrm{word}_1, \mathrm{word}_2, ..., \mathrm{word}_n\}$, we quantify the information of this dictionary as

$$I_{total} = \sum_{i=1}^{n} m_i * I(\mathrm{word}_i)$$

with $I(\mathrm{word}_i) = -\ln p_{\mathrm{word}_i}$ and $m = 100 * p_{\mathrm{word}_i}$. $p_{\mathrm{word}_i}$ denotes the ratio of how often a particular word has been used to the total number of words that this user has tweeted.

Figure S12a shows the distribution of total information per dictionary for the active and the random user sample. Apparently, the more active users have more rich dictionaries. This is simply due to the fact that a dictionary with more words leads to a higher total information. But already from this, we can conclude that the majority of the highly active users tends to use more different words than the average user which is a strong indication of natural behavior which we do not expect from bots.

To further quantify how the users in the two selected groups use their own dictionary to formulate more or less complex tweets, we computed the information of each tweet of each user given the dictionary of this user. The information contained in tweet $T$ with $l$ words is computed as

$$I_T = \sum_{i=1}^{l} I(\mathrm{word}_i).$$

Figure S12b shows the distribution of the information per tweet for the different user groups. At the first sight, both distributions share similar features. There is a monotonous
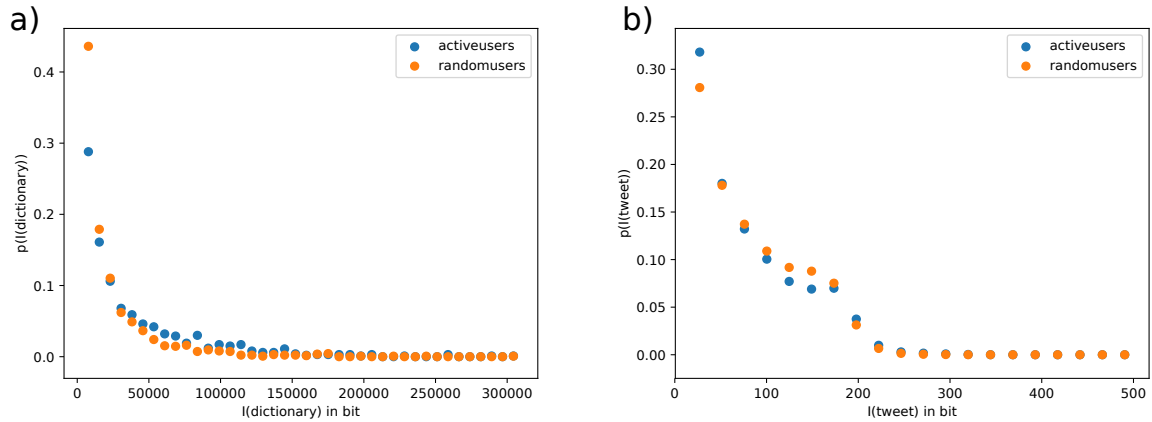
**Figure 12.** Inter-event time distribution of ten random user samples drawn from sample 2.

decrease towards high information per tweet and a small plateau for medium information. The only minor difference is the slightly larger share of tweets containing almost no information of the highly active users. We here want to emphasize that this does not necessarily mean that a user tweets nonsense or always the same; it rather refers to users having a large dictionary and only using the often occurring words for a majority of posts. This is exactly the behavior we would expect from highly active but human users.

Summarizing, the analysis of the semantics of the tweets revealed that there is no directly hint towards the suspicion that the highly active users are mostly bots. Of course, we cannot conclude that none of the users with a high activity is a bot (indeed we assume that there are bots active and included in our samples) but we believe that given the high similarity of the semantic information in tweets by random users and highly active users our results are not predominantly driven by bot behavior.

## Survivorship bias

In our analysis, we based our findings on active users sampled in three distinct periods in 2019. This inherits a substantial sampling bias: all users who we tracked have not dropped out. Therefore, older cohorts exclusively consist of users who have not left Twitter in seven years. On contrary, the younger cohorts also comprise users who might drop out in the following year and, thus, do not use Twitter for such a long time. We fully acknowledge this bias but believe that this does not undermine our conclusions.

All observations we have made are indicating that Twitter users have gotten more active over the years. More simply: younger cohorts are more active than older cohorts. To reject the hypothesis that this might be an artifact of the sampling strategy, we have conducted an additional analysis by selecting users who have been active in February 2013 and February 2015 and tracking their activity in the following years. The results are shown in Fig. 13.

For our analysis, we sampled $50,000$ users who have been active in the respective month and split them into cohorts that have left Twitter at different points in time. Figure 13 indicates the average number of tweets per user cohort for the years after the sampling period. We find that a higher initial activity strongly correlates with a slower drop-out of users. This observation is similar for both sampling periods.

Hence, by sampling users in 2019 we have obtained more active users, the older the cohort gets. Taking this into account, our results would even get amplified as the younger cohorts still comprise less active users who might drop out soon. to not bother with quantifying this effect we believe that the increase reported in the main manuscript can be evaluated as a lower limit of the actual acceleration.

Here, we want to note two important remarks regarding our analysis:

First, the comparison of Fig. 13a and Fig. 13b illustrates two snapshots of the analysis conducted in *Users grouped by cohort, see main MS*. We observe that the initial activity has risen from 2013 to 2015 from 25.02 posts per user to 26,98 posts per user in the first year. Second, for not considering the decrease of activity due to constant drop out
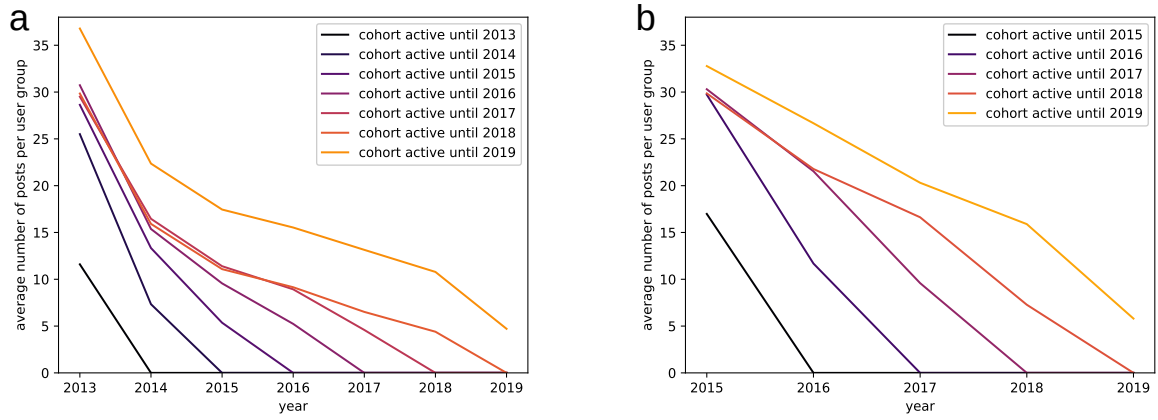
**Figure 13.** Average number of posts per Twitter user for cohorts over time for a) users sampled in February 2013 and b) users sampled in February 2015. Here, we also consider users which have dropped out. User groups differ in the year of their last recorded post as indicated by the colors.

of users which we observe here we have only included users tweeting for a consistent year and considering only this year in the part *users grouped by activity*.