Title: LGBTQ Visibility Online Measured Consistently and Persistently from 2012 through 2023

Authors: Jason Jeffrey Jones[1]*, Irissa Cisternino[1]

Affiliations:

[1]Stony Brook University; Stony Brook, NY, USA.

*Corresponding author. Email: jason.j.jones@stonybrook.edu

Abstract: LGBTQ visibility is an often discussed but rarely quantified concept. Here we define online visibility as the prevalence of active social media accounts with an LGBTQ signifier in the profile bio and measure the prevalence of such accounts consistently and persistently at daily resolution over twelve years in the United States. We found that prevalence for all LGBTQ signifiers increased. The term 'gay' grew most rapidly. Accounts with LGBTQ signifiers were especially visible on days corresponding to political or violent events. The rainbow flag emoji also increased in prevalence, including a notable ratchet each June (Pride Month). Social scientists should embrace this new opportunity to observe how people describe their selves to a public audience.

One-Sentence Summary: The visibility of self-ascribed identities such as LGBTQ within online populations can be observed to change over time.

Main Text: Historically, those with LGBTQ identities have been shunned, persecuted, and marginalized (Noga-Styron, et al., 2012). In recent history, in the United States, however, the stigma has begun to dissipate. Public attitudes toward homosexuality have been decreasingly negative for decades (Gallup Inc., 2022). Broader recognition of civil rights for LGBTQ citizens has proceeded fitfully (Geidner, 2019).

As the stigma has lessened, the proportion of the US population that identify as LGBTQ has increased (Jones, 2022). Some in the LGBTQ community have advocated for increasing LGBTQ "visibility." At the individual level, visibility varies from keeping one's identity secret – being closeted – to identifying openly, frequently and publicly (Lasser & Tharinger, 2003). In a pamphlet titled "Transgender Visibility – A Guide to Being You," the Human Rights Campaign Foundation strongly encourages transgender individuals to express that identity (Human Rights Campaign Foundation, 2014). In the text, being visible is associated with "the joy of living an open, honest life and engaging in relationships as a whole and authentic person." The pamphlet concludes with a message from the President of the Foundation in which it is stated, "Progress toward equality is made when we choose to share our lives with others. This simple yet profound step is the greatest political action any of us can take."

Those outside the community have also affirmed the importance of visibility. On March 31, 2022, US President Biden stated, "To everyone celebrating Transgender Day of Visibility, I want you to know that your president sees you." (The White House, 2022). On that day, the Biden-Harris administration noted that over 14 percent of their appointees identified as LGBTQI+ and announced new efforts for "advancing visibility for transgender Americans." (Biden-Harris Administration, 2022)

In academic research, many have studied the visibility management of LGBTQ individuals (Dewaele et al., 2014; Lasser & Tharinger, 2003; Twist et al., 2017). As originally defined, visibility management is "the ongoing process by which GLB adolescents make careful, planned decisions about whether they will disclose their sexual orientation, and if they decide to disclose, to whom and how they disclose, and how they will monitor the presentation of their sexual orientation" (pg. 237, Lasser & Tharinger, 2003).

Here we present a method and results for measuring LGBTQ visibility online. The method has several desirable features. First, it provides a consistent and persistent measure. The method can be applied with no changes to past, present and future data. Second, it is unobtrusive and privacy-preserving. The method is applied only to publicly released data, and the measurements produced contain no individually identifying information. Third, the measurements produced have strong construct, external and temporal validity. We discuss each of these in turn.

Our measure of LGBTQ visibility online is deliberately simple: in the population of observed individuals, we estimate the prevalence of those using an LGBTQ signifier to describe themselves. The method to do so is straightforward; locating matching substrings and counting unique observations are well-defined problems with formal solutions. To measure the visibility of LGBTQ individuals online over time, it is hard to imagine any more face-valid measure than consistently and persistently estimating the proportion of individuals in the observed population who use an LGBTQ signifier in their self-description. Thus, we argue the current method has strong construct validity (i.e. it measures what it purports to – LGBTQ visibility online).

The method also provides external validity, i.e. generalization outside the research context. Here we apply the method to users of Twitter likely within the United States. Analysis of self-authored social media profile text has the advantage of observing actual, public self-descriptions rather than lab- or interview-acquired responses. No item needed to be contrived or implemented by the researchers. Rather, millions of individuals responded to the prompt: *Describe yourself in 160 characters or less*. The current work makes use of Twitter because of data availability, but the method could be generalized equally well to any online platform.

Temporal validity – despite its importance – is nearly universally neglected in social science (Munger, 2019). Research results have temporal validity if they generalize across time. Because people and societies change, any knowledge generated about them necessarily decays over time. This is an inconvenient truth for social science. Most studies are conducted in one limited moment, but the expectation imparted by training and culture in the social sciences is that one should claim universal and final knowledge. A popular strategy is to ignore time (or summarize over it) and not speak of temporal validity.

Here we take the opposite approach by placing time at the center of our analysis. Specifically, we examine temporal trends in the way active US Twitter users describe themselves. The general method is called Longitudinal Online Profile Sampling (LOPS; Jones 2021). With LOPS, one studies the choices people make when describing themselves with words. Call the output of this act: *personally expressed identity*. It is personal – the individual is describing themselves. It is expressed – these are words the individual emits, where others might see them. And it describes identity – the explicit purpose of the text is description of the author.

To gather observations of personally expressed identity with LOPS, one samples self-authored biographies from online profiles as often as possible. From such samples, compiled for many users (multiple millions of US Twitter users in the present work) one can draw inferences about temporal trends in aggregate personally expressed identity in a population.

LOPS has mostly seen use in the study of US politics. It has been observed that Americans increasingly use political words to describe themselves (Rogers & Jones, 2021). In response to the January 6, 2020 insurrection, Americans removed words associated with Donald Trump and the Republican party from their bios (Eady et al., 2022). *Any* set of signifiers can and should be studied, however.

Here we focus on signifiers of LGBTQ identities. LGBTQ signifiers were any of: *lesbian*, *gay*, *bisexual*, *trans*, *queer*. The prevalence of users with LGBTQ identity signifiers within their bio increased greater than fourfold from 2012 to 2023 (Figure 1). Specifically, per 10,000 unique, tweeting US Twitter users in 2012, we estimate 34 had a self-description containing an LGBTQ signifier. In 2023, we estimate the number as 158 per 10,000. These estimates are based on an average of 261,239 users per day in 2012 and 159,215 in 2023; a t-test indicated a significant difference, $p < 0.001$.

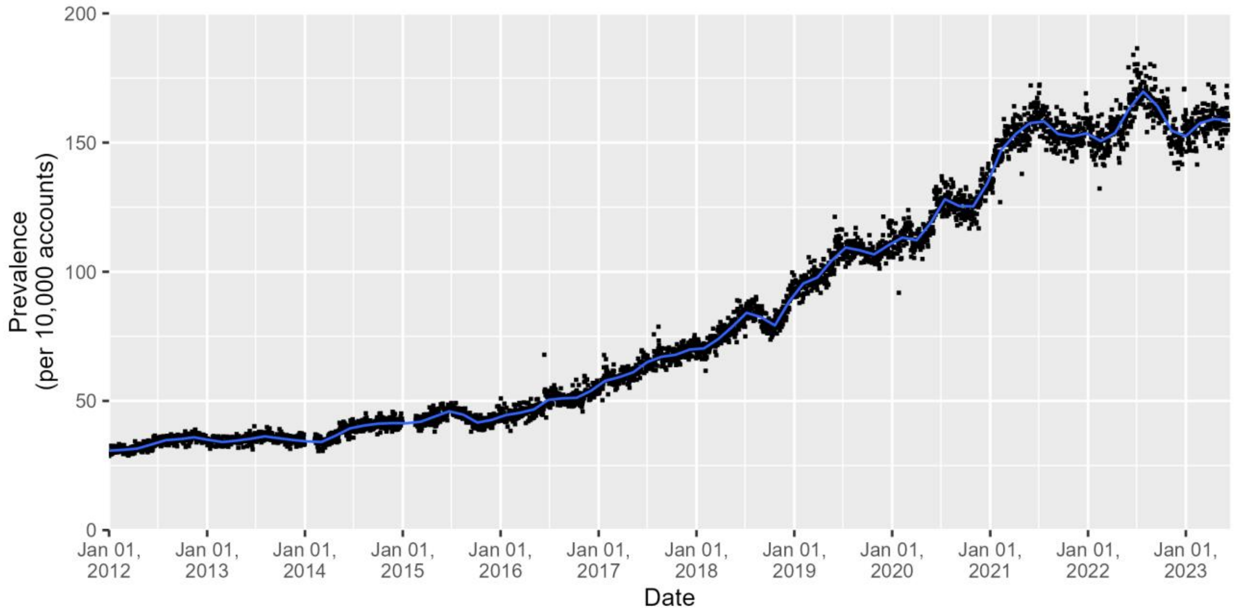

Figure 1. Daily estimates of the prevalence (per 10,000) of US Twitter users who tweeted that day and whose bio contained one or more LGBTQ identity signifiers. LGBTQ identity signifiers were *lesbian*, *gay*, *bisexual*, *trans*, *queer*. The blue line is a loess smoothed curve.

The signifiers did not rise in prevalence at equal rates (Figure 2). *Lesbian* was more prevalent than *queer* in 2012, but *queer* grew rapidly in visibility relative to *lesbian*. *Queer* increased in prevalence at a rate of 3.4 per 10,000 per year. For reference, per (Jones, 2021), tokens change in prevalence at a mean annual rate of zero with a standard deviation of +/-3 per 10,000. *Trans* also increased quickly relative to other signifiers in this period: +3.6 per 10,000 per year. Even faster growth is apparent for *gay* (+4.8 per year). Some individuals choose different words (e.g. *transgender* rather than *trans*) to express their identity and others express identities not included here (e.g. *asexual*). Figures S1, S2, S3, S4 and S5 in the supplementary material explore some alternatives.

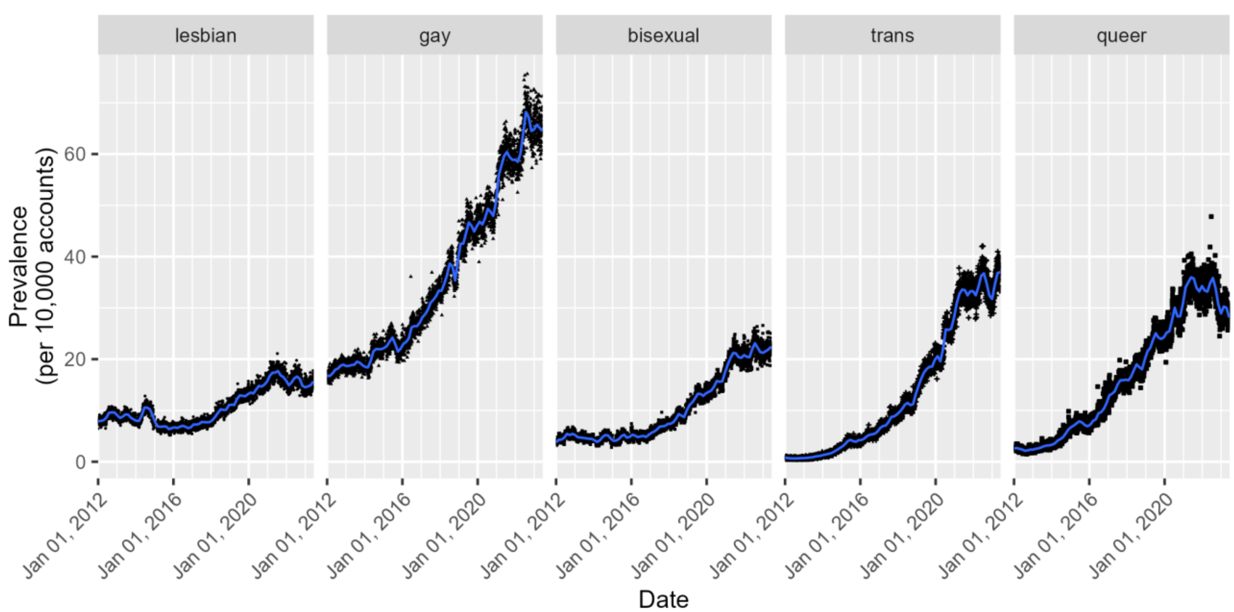

Figure 2. For each of the LGBTQ identity signifiers – *lesbian*, *gay*, *bisexual*, *trans*, *queer* – daily estimates of the prevalence (per 10,000) of US Twitter users who tweeted that day and whose bio contained that particular signifier. Each day, about 200,000 unique US Twitter users were observed tweeting.

Some daily prevalence estimates in Figure 1 are far from the smoothed trend curve. We investigated the 5 most unexpectedly high and low values manually. Table 1 lists these dates with speculative attributions for each. Every unexpectedly low estimate is temporally aligned with a high-profile professional sports event. One observes a *denominator effect* in each case. By denominator effect, we mean that the total number of active users was unusually large, and the number of LGBTQ active users did not increase by the same magnitude. For instance, on February 7, 2016, an unusually large number of users were observed as active: 285,610 as compared to 208,795 the day before and 231,600 the day after. The number of active users with LGBTQ signifiers in their bio increased as well, but not as much as the total.

Table 1. The five lowest and highest *relative to expectation* daily prevalence estimates.

Ratio of Observed Prevalence to Expected	Date	Attribution
0.81	2020-01-26	NBA Star Kobe Bryant dies in a helicopter crash.
0.86	2016-02-07	NFL Super Bowl 50
0.87	2015-04-30	First Day of NFL Draft, 2015
0.87	2021-02-07	NFL Super Bowl LV
0.87	2016-06-19	NBA Finals (Cleveland Cavaliers First Championship)
1.14	2019-06-01	Colorado bill banning conversion therapy for minors and first day of Pride Month
1.16	2016-01-02	Occupation of the Malheur National Wildlife Refuge
1.18	2017-08-12	Unite the Right march in Charlottesville, Virginia
1.18	2017-01-21	Womens' March following the Inauguration of US President Donald Trump
1.37	2016-06-12	Pulse Nightclub Shooting

Every one of the top five days with higher-than-expected LGBTQ prevalence coincided with a high-profile political or violent event. Users with an LGBTQ signifier were unusually likely to be observed tweeting on the day of the Pulse Nightclub Shooting, in which a gunman killed 49 people at a gay nightclub in Orlando. They were also more likely to be observed tweeting on the day of the 2017 Womens' March.

Figure 3 depicts daily prevalence estimates for several emojis: 🏳️‍🌈, 🌈, 🏳️‍⚧️, ⚧️. These emojis are frequently used either to identify as LGBTQ or to express support for others who do. Of the emojis, the Rainbow Flag has reached greatest prevalence. The prevalence appears to ratchet with each June (Pride Month). The Transgender Flag has increased in prevalence sixteen-fold from 2019 to 2023.

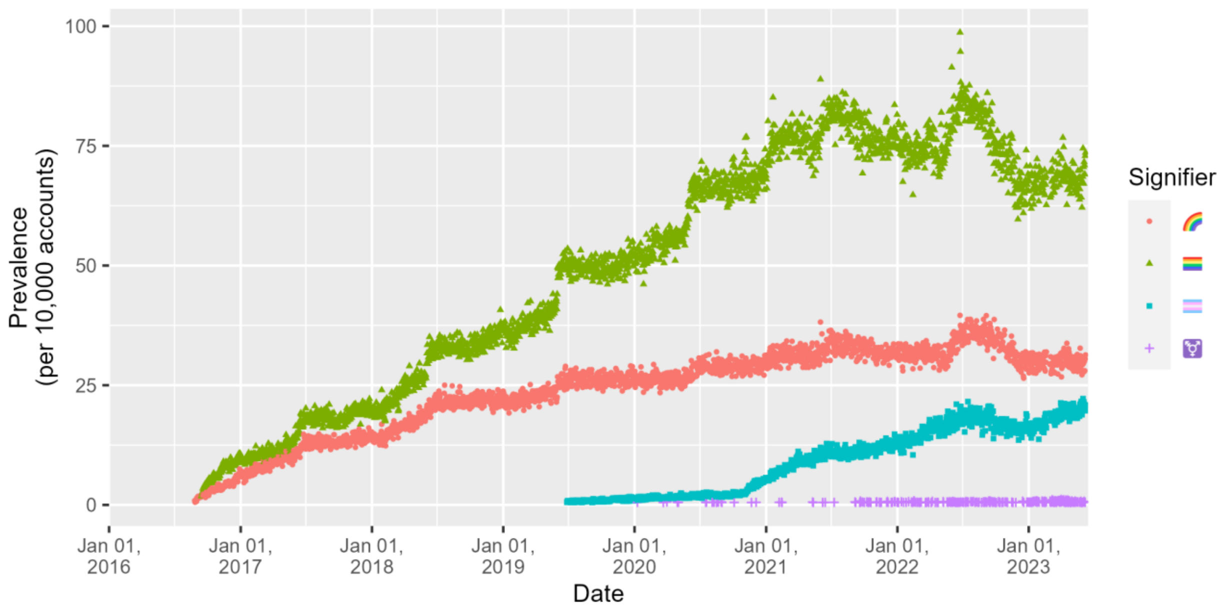

Figure 3. For each of the LGBTQ emojis – 🌈, 🏳️‍🌈, 🏳️‍⚧️, ⚧️ – daily estimates of the prevalence (per 10,000) of US Twitter users who tweeted that day and whose bio contained that particular signifier. Minor breaks correspond to July 1 each year; note that the rainbow flag emoji appears to ratchet up in June in many years.

Consistent, persistent measurement should be the standard method of social science now and in the future. Here we have demonstrated it can be done: LGBTQ visibility online was estimated from random samples of Twitter activity. We estimated prevalence at daily resolution with hundreds of thousands of observations per day and ultimately millions in total. Nowcasting (Bok et al., 2018) of *all* aspects of personally expressed identity could be achieved through modest investment in research infrastructure.

There are several points that should be made to avoid misinterpretation of these estimates. First, the current work is not an attempt to estimate how many Americans have LGBTQ identities. Nor is it to estimate that proportion within US Twitter users. Instead, we have estimated the prevalence of US Twitter accounts with a public expression of LGBTQ identity within the profile bio.

The authors are all too aware there is no one-to-one mapping of Twitter accounts to human beings. There exist "bot" accounts that are not human but rather automated scripts (Varol et al., 2017). There exist "sockpuppet" accounts (Kumar et al., 2017) – the same person controls multiple profiles – which would affect both numerator and denominator of prevalence estimates in unknown ways. "Inauthentic" accounts may present as individuals who do not exist (Gurajala et al., 2016). Each of these is a source of measurement error, but measurement error is an unavoidable cost of performing social science generally and big data analysis specifically (Salganik, 2019). The advantages – precision, temporal resolution and scale – vastly outweigh the disadvanatages.

We re-emphasize that we aim to measure visibility. Our daily samples of users derived from random samples of tweets. Thus, we are estimating prevalence within the daily active, tweeting subpopulation. We do this purposely, because the goal was measurement of visibility, and these accounts (tweet authors) were most visible. We measured explicit, self-provided signifiers in the bio, because these are unambiguous signals meant to be seen. Frequently, we are asked to "validate" this measure, but continue to resist these calls. It would be neither possible, nor appropriate to the aim of the endeavor to adjudicate which accounts are "real" or which accounts are properly LGBTQ. That is to say nothing of the ethics of deciding who counts and for what. Researchers who doubt the trends presented here are free to compare to other sources of data or to collect samples of authentic accounts (by their definition).

The current work is based on the assumption that language use reveals much about the minds of authors and the collective consciousness of the societies in which they are embedded. Many have adopted this idea and applied it to the language of published books (e.g. Michel et al., 2011). The surprisingly broad successes of word embeddings and large language models hint that analysis of text alone (at scale) can provide deep understanding.

Social scientists should embrace this new opportunity to observe how people describe their selves to a public audience (Jones, 2021). Further, they should challenge themselves to confront the issue of temporal validity. Instead of the historical custom of one-shot, small sample studies that purport universal, eternal discovery of truth, they should instead adopt consistent, persistent measurement. Astronomy slowly developed over millennia of watching the stars. So should a science of humans be built upon perduring observation of selves. With the current work, we present an early – perhaps crude, but effective – telescope for gazing upon the universe of personally expressed identity.

References and Notes

1. Biden-Harris Administration. (2022, March 31). FACT SHEET: Biden-Harris Administration Advances Equality and Visibility for Transgender Americans. The White House. https://www.whitehouse.gov/briefing-room/statements-releases/2022/03/31/fact-sheet-biden-harris-administration-advances-equality-and-visibility-for-transgender-americans/
2. Bok, B., Caratelli, D., Giannone, D., Sbordone, A. M., & Tambalotti, A. (2018). Macroeconomic nowcasting and forecasting with big data. Annual Review of Economics, 10, 615–643.
3. Dewaele, A., Van Houtte, M., & Vincke, J. (2014). Visibility and coping with minority stress: A gender-specific analysis among lesbians, gay men, and bisexuals in Flanders. Archives of Sexual Behavior, 43(8), 1601–1614.
4. Eady, G., Hjorth, F., & Dinesen, P. T. (2022). Do Violent Protests Affect Expressions of Party Identity? Evidence from the Capitol Insurrection. American Political Science Review, 1–7. https://doi.org/10.1017/S0003055422001058
5. Gallup Inc. (2022, May 2). LGBT Rights. Gallup.Com. https://news.gallup.com/poll/1651/Gay-Lesbian-Rights.aspx
6. Geidner, C. (2019, June 19). The Court Cases That Changed L.G.B.T.Q. Rights. The New York Times. https://www.nytimes.com/2019/06/19/us/legal-history-lgbtq-rights-timeline.html
7. Gurajala, S., White, J. S., Hudson, B., Voter, B. R., & Matthews, J. N. (2016). Profile characteristics of fake Twitter accounts. Big Data & Society, 3(2), 2053951716674236.

8. Human Rights Campaign Foundation. (2014). Transgender Visibility: A Guide to Being You. http://assets2.hrc.org/files/assets/resources/trans_guide_april_2014.pdf

9. Jones, J. (2022, February 17). LGBT Identification in U.S. Ticks Up to 7.1%. Gallup.Com. https://news.gallup.com/poll/389792/lgbt-identification-ticks-up.aspx

10. Jones, J. J. (2021). A dataset for the study of identity at scale: Annual Prevalence of American Twitter Users with specified Token in their Profile Bio 2015–2020. PLOS ONE, 16(11), e0260185. https://doi.org/10.1371/journal.pone.0260185

11. Jones, J. J. (2023). Ipseology—A new science of the self. Jason Jeffrey Jones Productions. https://jasonjones.ninja/ipseology-a-new-science-of-the-self-book/

12. Kuhn, M. H., & McPartland, T. S. (1954). An Empirical Investigation of Self-Attitudes. American Sociological Review, 19(1), 68–76. https://doi.org/10.2307/2088175

13. Kumar, S., Cheng, J., Leskovec, J., & Subrahmanian, V. (2017). An army of me: Sockpuppets in online discussion communities. Proceedings of the 26th International Conference on World Wide Web, 857–866.

14. Lasser, J., & Tharinger, D. (2003). Visibility management in school and beyond: A qualitative study of gay, lesbian, bisexual youth. Journal of Adolescence, 26(2), 233–244. https://doi.org/10.1016/S0140-1971(02)00132-X

15. Michel, J.-B., Shen, Y. K., Aiden, A. P., Veres, A., Gray, M. K., Team, T. G. B., Pickett, J. P., Hoiberg, D., Clancy, D., Norvig, P., Orwant, J., Pinker, S., Nowak, M. A., & Aiden, E. L. (2011). Quantitative Analysis of Culture Using Millions of Digitized Books. Science, 331(6014), 176–182. https://doi.org/10.1126/science.1199644

16. Munger, K. (2019). Knowledge decays: Temporal validity and social science in a changing world. Unpublished Manuscript. https://osf. io/4utsk

17. Noga-Styron, K. E., Reasons, C. E., & Peacock, D. (2012). The last acceptable prejudice: An overview of LGBT social and criminal injustice issues within the USA. Contemporary Justice Review, 15(4), 369–398.

18. Salganik, M. J. (2019). Bit by bit: Social research in the digital age. Princeton University Press.

19. The White House (Director). (2022, March 31). President Biden on Transgender Day of Visibility 2022. https://www.youtube.com/watch?v=8wCjz2SIYVo

20. Twist, M. L., Bergdall, M. K., Belous, C. K., & Maier, C. A. (2017). Electronic visibility management of lesbian, gay, and bisexual identities and relationships in young adulthood. Journal of Couple & Relationship Therapy, 16(4), 271–285.

21. Varol, O., Ferrara, E., Davis, C., Menczer, F., & Flammini, A. (2017). Online human-bot interactions: Detection, estimation, and characterization. Proceedings of the International AAAI Conference on Web and Social Media, 11(1), 280–289.

Acknowledgments: We thank the members of the Computational Social Science of Emerging Realities Group and members of the Institute for Advanced Computational Science for useful comments on these results.

Funding:

This material is based upon work supported by the National Science Foundation under grants IIS-1927227 (JJJ) and CCF-2208664 (JJJ). The Center for Advanced Internet Studies provided JJJ a fellowship in the Fall of 2023 that allowed the focus of time and attention on this work. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Author contributions:

Conceptualization: JJJ, IC

Methodology: JJJ

Investigation: JJJ, IC

Visualization: JJJ

Funding acquisition: JJJ

Project administration: JJJ

Supervision: JJJ

Writing – original draft: JJJ

Writing – review & editing: JJJ, IC

Competing interests: Authors declare that they have no competing interests.

Data and materials availability: All data, code, and materials used in the analysis are available at https://osf.io/f5xdw/.

Supplementary Materials

Materials and Methods

Supplementary Text

Figs. S1 to S5

Tables S1

Supplementary Materials for

LGBTQ Visibility Online Measured Consistently and Persistently from 2012 through 2023

Jason Jeffrey Jones, Irissa Cisternino

Correspondence to: jason.j.jones@stonybrook.edu

Materials and Methods

Observing Personally Expressed Identity Data with Online Profile Biographies

In this manuscript, we examined personally expressed identity text. Personally expressed identity text is self-authored, self-descriptive language that has been published where others might see. The source of this text was Twitter profile biographies. A bio appears on the user's profile page just under their photo and name. A bio is limited to 160 characters maximum. About 15% percent of US users leave their bios empty. We do not exclude empty bio users, because choosing to leave one's self-description blank is a valid choice.

As a source for bios, we chose to use the random sample of 1% of all tweets that was provided through the Twitter API. Although the tweet stream was the initial source, we wish to be clear that we always count unique users and never tweets. Here we use a daily-resolution, cross-sectional sample of US users filtered from the tweet stream.

First, tweets from users with non-US locations were filtered out. The geocoding function used to distinguish US locations from non-US is available at https://osf.io/472sf. Its aims are to mark US locations with a low false positive rate and to be understandable to anyone with Python and regular expression experience. The function takes as input the profile location text provided by the user. Based on this, it marks likely non-US locations (e.g. "Afghanistan", "Scotland") and likely US locations (e.g. "Clevleand, OH", "Staten Island"). The logic for defining US location patterns was developed by iteratively comparing the ways in which Twitter users described their locations and the output of the function. Supplementary Table S1 lists a sample of locations marked as US.

Second, one and only one profile observation per user was randomly selected per day. In the tweet stream, each tweet was accompanied by a snapshot of the author's profile at the moment of writing the tweet. In one day of tweets, when a user was observed tweeting more than once, we grouped their profiles into a set and chose exactly one observation at random. Thus, the denominator in every prevalence calculation contains the count of *unique* users observed tweeting that day – whether they were observed tweeting once or 100 times. Similarly, the numerator contains the count of unique users whose bio contained a signifier of interest.

<u>Operationalization of LGBTQ Visibility</u>

We defined LGBTQ visibility on the Twitter platform as the prevalence of unique tweet authors with an LGBTQ signifier within their bio. We used a short list of signifiers – namely, *lesbian*, *gay*, *bisexual*, *trans* and *queer*. In Figure 1 of the main text, we counted users whose bio contained any of these five signifiers. Specifically, we counted users whose bio matched the regular expression \blesbian\b|\bgay\b|\bbisexual\b|\btrans\b|\bqueer\b. This pattern matched text where any of the five signifiers was present as a whole word. The requirement to match the whole word increased specificity. For instance, "transportation" and "queers" did not count as matches.

For each individual signifier, we counted users whose bio contained the signifier as a token. We obtained tokens by splitting each bio on each instance of whitespace or word boundaries. It should be noted that we always counted unique users and not the tokens themselves. For instance, a bio that read "gay, gay, gay" would count as one match, not three. Emoji signifier prevalence was operationalized in the same manner.

<u>Operationalization of Visibility Relative to Expectation</u>

For each prevalence series, we fit a loess model to the daily estimates. The model predictions are plotted as a continuous blue line in each Figure in the manuscript. We used these model predictions as the "expected" value for each day. (Note that this is not the same usage for "expected value" one might see elsewhere.) In other words, we defined the model output for each day as the prevalence one would observe if all trends (local and longer-distance) continued without deviation.

We compared expected prevalence values to observed prevalence values by taking the ratio of observed to expected. If the value was 1.0, the observed prevalence matched exactly the model prediction. If the value was less than 1.0, we took this as an indication that the day's observed prevalence was below what one should expect. Values above 1.0 indicated an overabundance beyond expectation.

Ipseology

The methods described here are one manifestation of ipseology. Ipseology is the study of human identity using large datasets and computational methods (Jones, 2023). Ipseology refers to the study of ipseity - personal identity, selfhood and the essential elements of identity - at scale. Signifiers (i.e. the words, phrases, abbreviations and emojis) found within personally expressed identity text are the objects of study. Prevalence of use is the primary measure. Prevalence allows comparison across time and space. Using data at scale allows for precise, consistent and persistent measurement.

Supplementary Text

Twitter/X and Platform Risk

The source of data for this manuscript was the random sample stream of tweets from the Twitter API. This was available for free through Twitter's developer tools for many years. However, it was removed from the free tier after the platform was acquired by Elon Musk. The author's API keys continued returning data until June 9, 2023.

Two things remain true as of writing this manuscript: 1) The viability of using Twitter bios as a source for personally expressed identity text is severely limited compared to what it was. 2) More than a decade worth of personally expressed identity text observations continue to exist for hundreds of millions of individuals around the globe.

If nothing changes, the period of 2012-2022 may be the time when human identity is best understood. Data at high temporal resolution and wide geographic dispersion spilled out of Twitter's servers and onto hard drives across academia. However, the era came to end due to the whims of billionaires and markets. One can still pay Twitter for data, but it as a fraction of what it was for a steep price. At the $100 per month tier, we have been able to observe only 50,000 user profiles daily. That number is a global total, and pales compared to the 200,000 users per day observed from the US alone within the 1% tweet stream.

It is possible that the trend toward Internet platforms becoming more closed and more jealously hoarding their data will reverse. However, we believe researchers would be foolish to simply wait for that future. Below we propose new directions for the study of ipseology.

If one wishes to study Twitter, one needs Twitter data, of course. We urge Twitter (the company) to again make available a random sample of tweets and periodic random samples of user profiles.

New Directions for Ipseology

As computational social scientists and ipseologists, we need identity data at scale. We suggest two paths to new data streams. The first is traditional surveys. Scales that elicit personally expressed identity text are sometimes referred to as *Who am I?* instruments. Instruments like these can be delivered as short surveys. For a low investment of research expenditure ($1000) one could administer an instrument such as the Twenty Statements Test (Kuhn & McPartland, 1954) to a representative sample of a few hundred respondents.

Second, we urge the development of apps to collect self-authored, self-descriptive text from longitudinal panels. An app might simply collect periodic bios from volunteer citizen scientists. A more ambitious approach would be to build an app that provides value for users – perhaps feedback or accountability on goals for personal growth – in exchange for use of personally expressed identity text in research.

Notes on Development of this Manuscript

This manuscript began life in 2019 as *Do LGBTQ-related Events Drive Individual Online Disclosure Decisions?*. We pre-registered (https://osf.io/wvmfe) four hypotheses regarding the rate at which individuals would add and remove LGBTQ signifiers in response to events. Reviewers had difficulty understanding the methods, and the evidence was inconclusive due to our limited ability to construct *longitudinal* series at daily resolution.

We shifted focus and revised the manuscript. Find a preprint of *Societal Pressures, Safety, and Online Labeling-Investigating LGBTQ Self-identification in an Online Space* at https://doi.org/10.31235/osf.io/8yjcm. This manuscript further received reviews that seemed to miss the point and misunderstand the methods.

We believe the current version of the manuscript contains clearer claims regarding what the data demonstrate and simpler methods than our original ambitions. We hope this facilitates its review, but we invite anyone interested to explore the years of development this manuscript has borne.

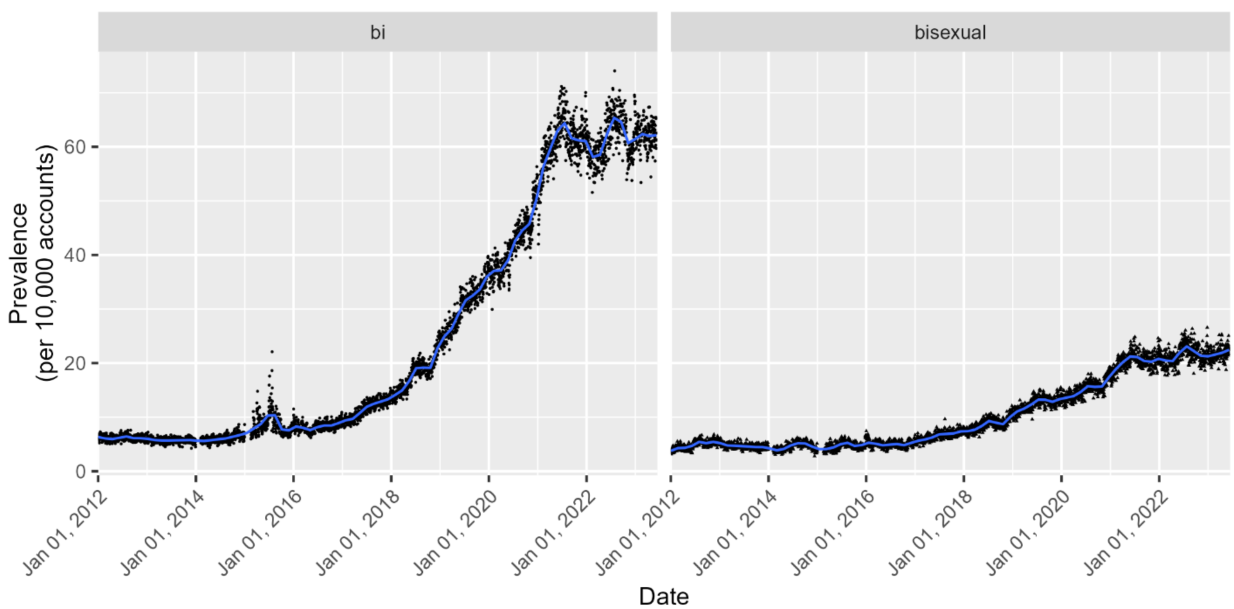

Fig. S1.

Prevalence within US Twitter users over time for two variants of bisexual signifier: *bi* and *bisexual*. Throughout the main text, we only include analysis for *bisexual*. The signifier *bi* has other, frequently used meanings such as an abbreviation for business intelligence and use in hyphenated terms. *Bisexual* leaves less room for ambiguity, but excluding *bi* may cause an undercount of LGBTQ visibility.

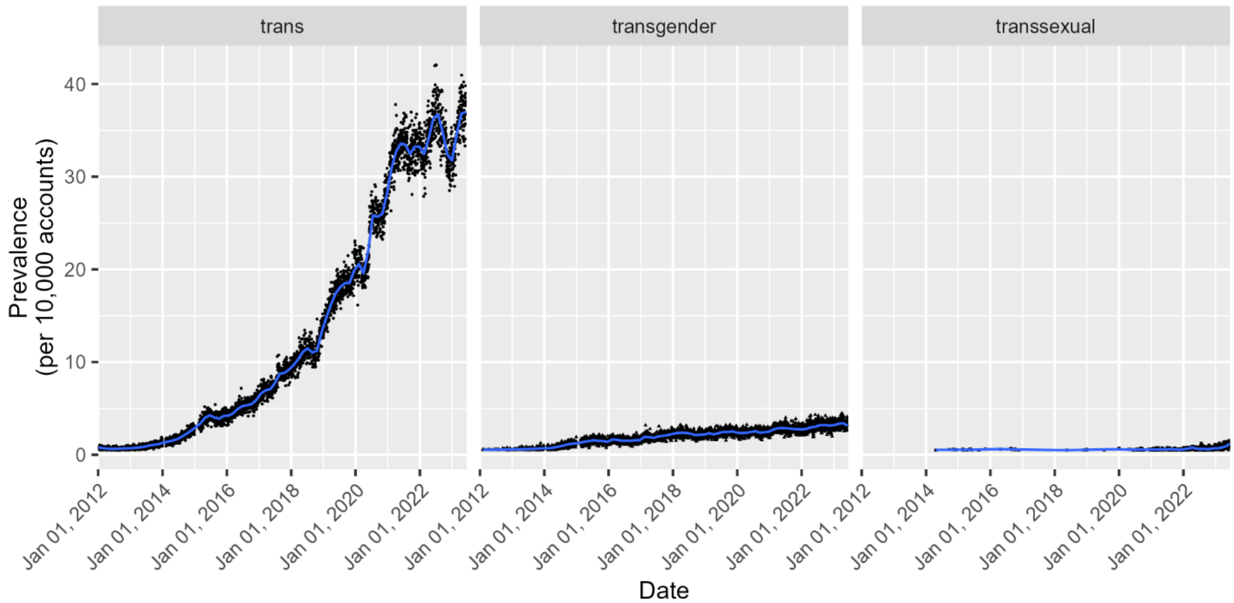

Fig. S2.

Prevalence within US Twitter users over time for three variants of trans signifier: *trans*, *transgender*, and *transsexual*. Throughout the main text, we only include analysis for *trans*. The signifier *trans* appeared to be users' preferred term, but may include use in other meanings. *Transgender and transsexual* leave less room for ambiguity, but do not show the pattern of rapidly increasing visibility other signifiers display.

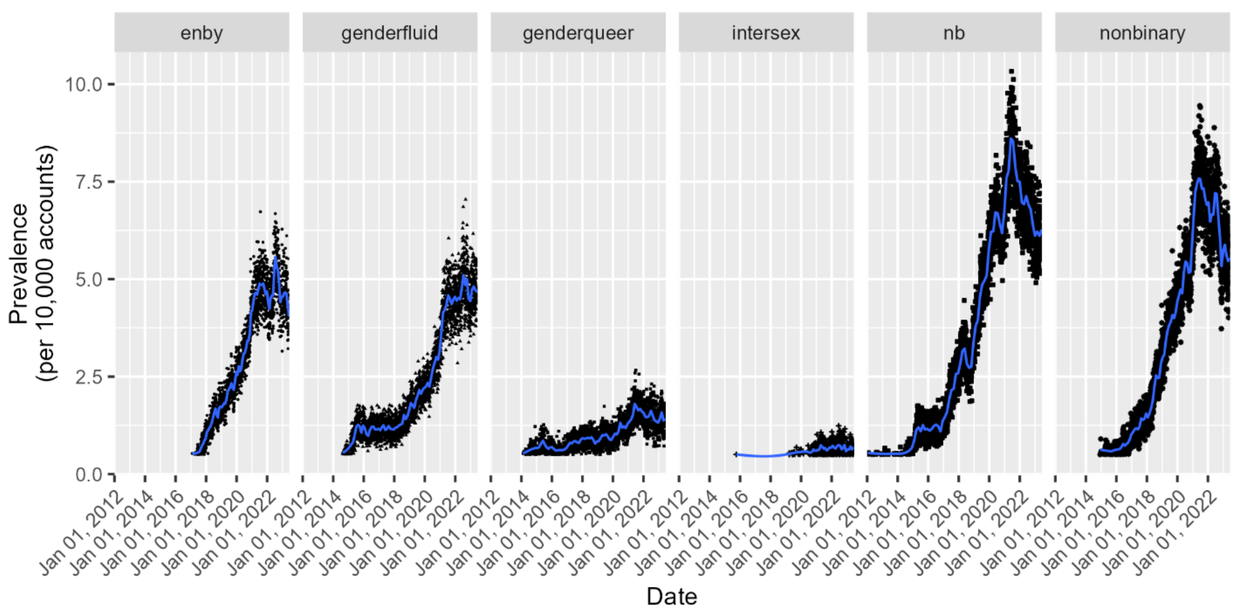

Fig. S3.

Prevalence within US Twitter users over time for three variants of gender non-conforming signifier: *enby*, *genderfluid*, *genderqueer*, *intersex*, *nb*, and *nonbinary*. We present these series here because they are related to LGBTQ visibility, but we do not interpret them in the main text.

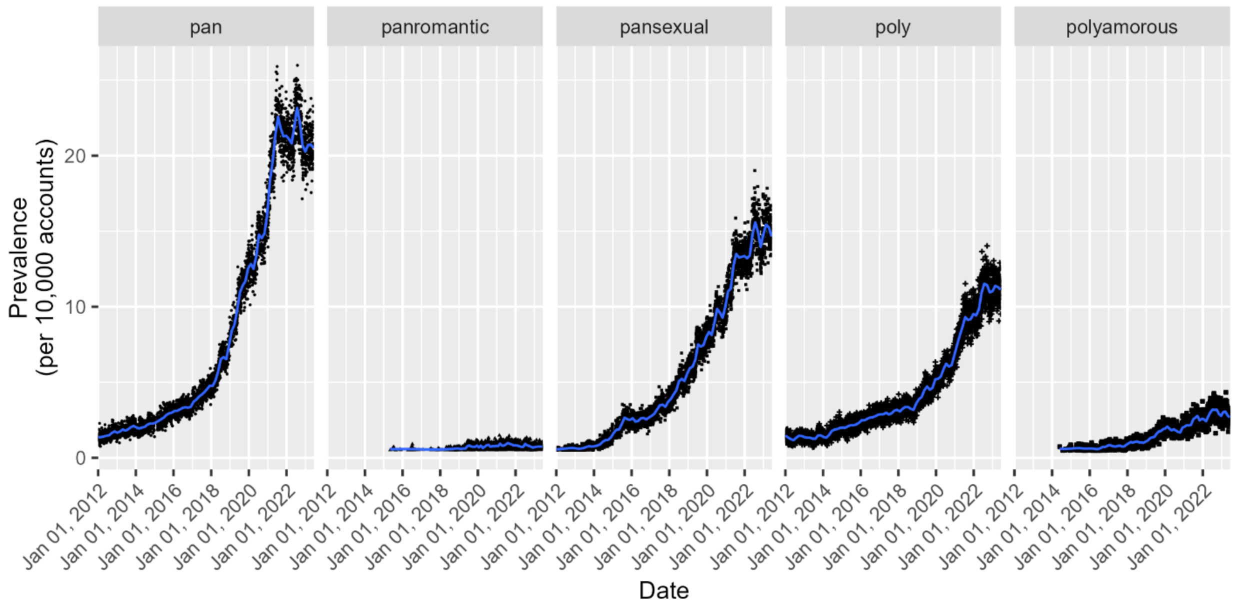

Fig. S4.

Prevalence within US Twitter users over time for five variants of pan and poly signifier: *pan*, *panromantic*, *pansexual*, *poly*, and *polyamorous*. We present these series here because they are related to LGBTQ visibility, but we do not interpret them in the main text.

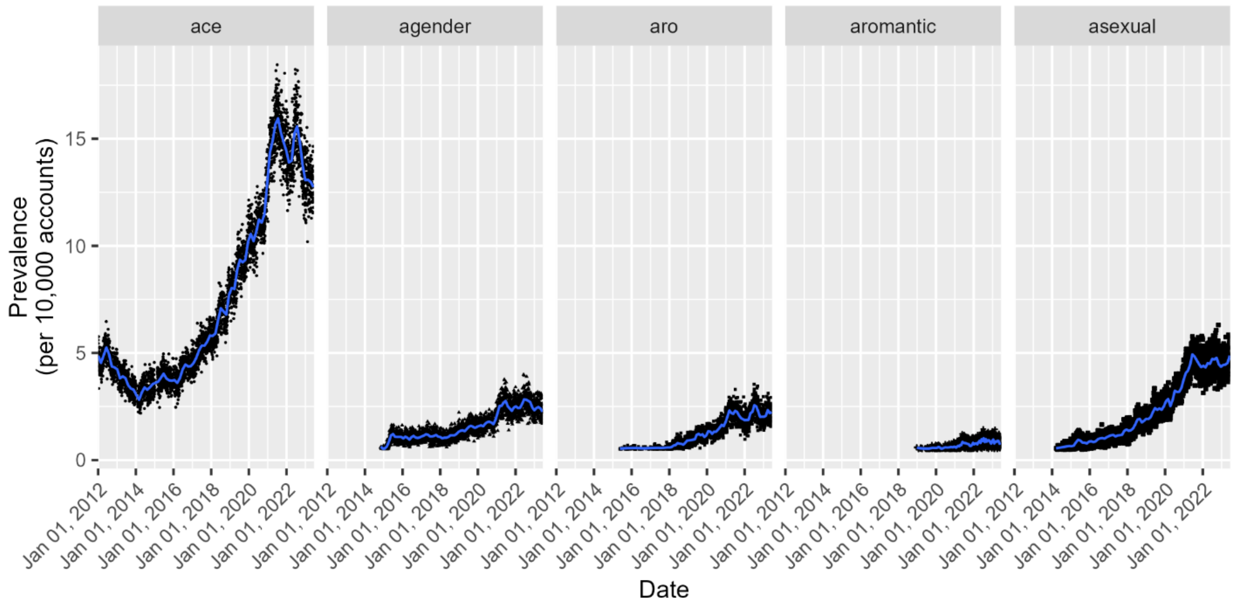

Fig. S5.

Prevalence within US Twitter users over time for five variants of a- signifier: *ace*, *agender*, *aro*, *aromantic*, and *asexual*. We present these series here because they are related to LGBTQ visibility, but we do not interpret them in the main text.

Table S1.

Location	Prevalence (mean)	Rank
united states	310	1
los angeles, ca	159	2
california, usa	124	4
atlanta, ga	86	8
new york, usa	66	16
ohio	34	32
washington, usa	20	64
detroit	12	128
u.s.a.	4	256
brownsville, tx	2	512

A sample of locations found within users' profile location field that our geocoding method marked as US. To compute Prevalence, we sampled one day per year, computed location prevalence (per 10,000 unique users) and then took the mean of those values. Rank is the ordinal rank of each location based on mean Prevalence (descending order).