Journal of Quantitative Description: Digital Media 1(2021), 1-36 10.51685/jqd.2021.017

The Times They Are Rarely A-Changin': Circadian Regularities in Social Media Use

SEAN KATES

University of Pennsylvania, United States

JONATHAN NAGLER

JOSHUA A. TUCKER

RICHARD BONNEAU

New York University, United States

This paper uses geolocated Twitter histories from approximately 25,000 individuals in 6 different time zones and 3 different countries to construct a proper time-zone dependent hourly baseline for social media activity studies. We establish that, across multiple regions and time periods, interaction with social media is strongly conditioned by traditional bio-rhythmic or "Circadian" patterns, and that in the United States, this pattern is itself further conditioned by the ideological bent of the user. Using a time series of these histories around the 2016 US Presidential election, we show that external events of great significance can disrupt traditional social media activity patterns, and that this disruption can be significant (in some cases doubling the amplitude and shifting the phase of activity up to an hour). We find that the disruption of use patterns can last an extended period of time, and in many cases, aspects of this disruption would not be detected without a circadian baseline.

Kates: skates@sas.upenn.edu
Nagler: jonathan.nagler@nyu.edu
Tucker: joshua.tucker@nyu.edu
Bonneau: bonneau@nyu.edu
Date submitted: 2021-01-05

Keywords: social media, biorhythms, ideological conditionality

There is a quickly growing literature on the role played by social media and social media users in politics and political outcomes. However, there is very little work on what the baseline use of social media looks like, how it is conditioned by attributes of the user, and how quickly disruptions to the pattern return to normal.

This paper begins to address these issues, presenting evidence from Twitter users in three different countries (France, England, and the United States) and across six different time zones that social media use is heavily controlled by traditional Circadian patterns, as well as social norms. We dive deeper into the American data to show that usage patterns can be conditioned by attributes of the users, in particular their estimated ideology.

Finally, we use social media usage histories from over 15,000 individuals located in the United States around the time of the 2016 election to show that, while most breaks in the established pattern are transient blips and usage quickly returns to expected levels, some events are so disruptive that they upset the balance for multiple days and can even shift key features of the use pattern by multiple hours. In the process of identifying these baselines and stylized facts, we also make methodological advances in establishing the strongly Circadian pattern of social media use at the aggregate level. Accounting for this baseline in a time-zone dependent manner can increase the resolution and sensitivity of social media based studies of response to large scale social events.

The Importance of Proper Baselines

Studies of social media routinely consider changes in social media use and the structure of social media networks for their effect on political outcomes.

These include, among many others, social media's capacity to allow individuals to organize for political actions like protests ((Gonzàlez-Bailòn et al., 2011), (Barberà et al., 2015b), (Gonzàlez-Bailòn and Wang, 2016)), the way in which individuals connected in a social media network may influence the political behavior of others in the same network ((Aral and

Walker, 2012), (Bond et al., 2012)), and how social media can be used as a tool for purposes both democratic and illiberal ((Tucker et al., 2017)).

However, these analyses tend to focus on contexts and activity that we expect deviate significantly from the norm of social media behavior. There is limited work on what social media use looks like in the population on an every day, hourly, or even more granular basis, the degree to which it follows predictable patterns, and the differences in patterns among different types of people.

We believe these are important questions, both methodologically, and substantively. In the subsections below, we cover first why proper baselines are important for broad methodological reasons, and then why they are important to answer the substantive questions of interest we spend the rest of the article pursuing.

Methodological Importance

Social media use (measured as total volume, an action/inaction dummy, or shifts in content or linking dynamics) can act as either a dependent or an independent variable of interest. One can ask how shifts in social media use affect a desired outcome, or alternatively how particular actions or changes shift social media use. In both cases, it is imperative for the researcher to understand the underlying baseline of use, in order to provide either the proper dosage (independent variable) or proper effect (dependent variable) in analysis.

Granularity in measurement allows for more complex, and more complete, analyses. Consider a researcher analyzing how a focused, and localized shutdown of state-provided internet or phone access affected the capacity of citizens of that state to organize resistance to some policy. While the researcher would likely desire to compare the volume of social media usage during the shutdown to some other period (establishing that the shutdown had the expected intermediate effect leading to the ultimate outcome of reduced organization), doing so without a fairly granular baseline for social media use in that state would be fairly capricious. As we show later, the combination of time of day and day of the week could vastly change the expected level of participation during the shutdown period.

While researchers could and should establish proper baselines on a case-by-case basis, informed by their subject matter expertise, it is also important to establish certain common expectations on baseline use in social media more generally. Doing so opens up avenues for additional research on the resilience of the patterns themselves, as well as the long-term shifts in behavior from equilibrium-altering events. While a simple point, it should be apparent that if a researcher wishes to analyze changes in social media behavior over time or in response to a specific event, it is important to establish a baseline for social media use *before* the event or time break in question.

Substantive Importance

Investigating why and how patterns of social media use change over time is a valuable (and relatively unexplored) area of future research. One particular area where such analyses are likely to be impactful is the study of political participation. In this paper, we focus on one particular aspect of political participation: the level to which highly salient political outcomes can mobilize or demobilize individuals from a particular form of political participation - in this case, social media use.

Previous work has demonstrated how election results can have vastly different effects on the way winners and losers perceive democracy's legitimacy,[1] and how increased dissatisfaction with democracy may affect turnout in elections (see, e.g. (Flickinger and Studlar, 2007) and (Karp and Milazzo, 2015), but also see (Kostelka and Blais, 2018) for suggestion that this relationship may be backwards). More generally, high-salience elections are likely to affect not only the psyches of winners and losers, but also how those individuals act politically.

Social media has provided an increasingly larger forum for political participation, and we know that the makeup of individuals participating in social media political discourse is an important shaper of its underlying substance (see, e.g. (Feezell, 2018)). If we expect that the dynamics we observe in other forms of participation after election outcomes are mirrored on social media, particularly among the most ideologically driven, it is important to study how different people or groups of people enter, exit, and change their social media behavior after

[1] (Anderson et al., 2005) presents an important discussion of how democracy naturally creates electoral losers whose support remains vital for a well-functioning democracy, and (Pierce et al., 2016) and (Daniller, 2016) discuss how these effects are exacerbated by how invested individuals are in the election.

high-salience events.

From the broader literature, we should expect that elections of high importance will split the winning and losing users in some fashion. "Winners" are likely to be satisfied, and if anything increase their participation (or volume) in response. However, there are multiple directions in which losing users might move. Following the literature on legitimacy, we might expect that social media users who experience a sharp loss would reject further political participation, and perhaps withdraw from particular forms, either in an attempt to avoid painful mentions of the outcome or perhaps out of shame and responsibility for claims made leading up to the event. Alternatively, participation literature that emphasizes the focusing power of political loss might suggest that social media "losers" actually increase their social media use in the aftermath of a tough loss, either as a source of comfort, or as a place to make plans on how to respond to the loss.[2]

These issues are important enough at the individual behavior level, but they take on additional significance when we consider the downstream network effects of large numbers of individuals shifting their social media use, particularly in cases where exit is a frequently-chosen strategy.

Changes in the volume or tenor of the discussion on social media, even if isolated to the most ideologically invested of the platform's users, will affect what is seen by those connected to the individuals, as well as by the broader network.[3] Thus, the question of how social media use changes over time is ultimately also an indirect way of measuring the seeding of political information throughout the social media platform.

Our paper is descriptive in nature, and we focus on showing only evidence that, around a highly-salient election, there can and do exist large shifts in social media use from well-established and otherwise stable baselines. We leave it to further research to determine the regularity of this occurrence, the contexts in which elections are likely to have this effect, as well

[2] Ongoing events in the United States at the time of this writing suggest that in the future it might be worth examining whether concerns regarding electoral fraud – be they legitimate or entirely illegitimate – might also exacerbate post-election social media usage among election losers.

[3] Consider (Feezell, 2018), for instance, on how exposure to different types of messages can alter the salience recipients attach to the underlying policy issue.

as many other interesting and important questions about participation effects at the individual level.

All of this is in service of showing how important it is to establish a proper baseline for social media use, as none of the ultimate evidence we present about shifts in social media use would be possible without one.

There are, of course, other areas of interest where studying baseline use is important. Ideological asymmetry in social media use is one such area. We have strong evidence that "leftists" and "rightists" have significant psychological and dispositional differences in areas as diverse as uncertainty management, tolerance of ambiguity, integrative complexity, and self-deception (for meta-analyses, see (Jost et al., 2017) and (Jost, 2017)). However, little work has been done on establishing the different ways in which ideologically distinct sub-groups interact with social media platforms on a daily basis. Establishing baselines for subgroups of the population allows us to distinguish how the same event can affect different groups in different directions.

If we repeated the analyses below across elections that varied which ideological group was the winner and which the loser, establishing baselines for ideological subgroups would allow us to judge how each subgroup responds to an important electoral outcome that casts them in their roles. In doing so, we could provide some evidence of an "elective affinity" between strong or weak reactions to victory or loss and ideological camp.[4]

For now, however, we move on to discuss the main goals of this paper.

Analytical Goals

Our primary goal in this paper is to identify patterns in social media use, both in large populations, and in subsets of the populations created based on the attributes of the subgroup members. From a top down structure, this allows us first to establish that there is some basic

[4]One might think of this in a 2X2 structure where an individual can either react to a win or a loss and the reaction could be relatively strong or weak. The proposed study, possible only with the construction of proper baselines of behavior, would allow us to speak to whether an individual's ideological type conditions the likelihoods across outcomes.

underlying rhythm to social media use across social media users writ large, including populations in different countries and cultures. We then break a specific population (American social media users) into smaller subgroups based on ideology to analyze the differences in baseline use, to the extent there is any, over time.

Our secondary goal is to investigate how different types of shocks upset the pattern of use, and how quickly the pattern returns to normal after said shocks. We combine this investigation with the subgroup description above to analyze whether shocks can also affect subgroups differently, particularly in cases where we expect the shocks to be more or less meaningful to different groups. Here, we specifically look at how the 2016 American presidential election, held on November 8th, 2016, affected the separate ideological subgroups of users in the days following the outcome. We compare not only the mobilization effects between subgroups, but also the degree to which the whole population's pattern of use shifted in comparison to previous shocks to the same users.

Before we move to exactly how we carry out these goals, we first discuss why we might expect patterns in social media use, and how we expect it to be structured.

The Rhythms of Social Media

Much of human behavior is cyclical in nature. These cycles are driven by biological patterns (responding to hormone release, as well as light and temperature shifts) and societal constructs like "work weeks" and meal schedules. Measuring deviations in behavior of any type, then, requires establishing baselines that take these cycles into account where appropriate. Recognizing the cyclicality inherent in behavior (including on social media) is particularly important when expected behavior differs wildly depending on the position in the cycle one is analyzing. In this paper, we extend this proposition directly to social media use. We propose that social media use is extremely cyclical, and that we should account for this in the establishment of baselines for use.

Social media use is a voluntary activity in which it is costly for an individual to participate. In particular, there are opportunity costs to the time spent reading and responding to messages. These costs are predictably greater during specific parts of an individual's daily life. Perhaps the most obvious example is during traditional sleeping hours. We expect the vast majority of our individuals to prefer sleep to social media use, and thus for social media use to be predictably lower during this time.

While nowhere near as stark, we also expect social media participation to be conditioned by time of day (more likely in the periods after cortisol release in the morning, less so after the body's release of melatonin), day of the week (individuals may be more free to use social media on weekends), or some combination of the two (individuals may be less likely to use social media during hours and days normally reserved for entertainment, like Friday or Saturday nights).

Additionally, we expect the very nature of social media use itself to have cycle-reinforcing effects. Individuals tend to use social media in an effort to increase social interaction - to the extent that this is a shared goal of individuals on the medium , there is likely to be an equilibrium schedule found when more or most individuals choose to engage in order to have more potential interaction. That is, individuals looking to engage with others will tend to drift towards a common schedule when attempts to engage during fallow times are met with non-response, and those during heavy times receive quick positive reinforcement.

Finally, social media platforms exist in a larger world of communication and media, containing its own patterns that may shape those on the platform via interactions between the two. If people reliabily react to stimuli from the external world on social media, the regularity of that stimuli could also produce, or heighten usage patterns. One particularly good example of this is so-called "second screening," where individuals watch news media or other television programming while also participating on social media, reacting to what they see (for a much more thorough discussion of second screening and its ramifications for political behavior more generally, see (Gil de Zúñiga et al., 2015)). With both news and entertainment programs generally on set schedules, interactions on social media that are driven by these events are likely to create and reinforce their own patterns in social media use.

We are not the first researchers to theorize about potential patterns in social media use, and others have found similar regularities to the type we demonstrate below, albeit generally for different purposes. Scholars of neuroscience, for instance, have used the patterns in social media use in an UK sample to measure how positive and negative mood shifts during the day tend to themselves be cyclical (though differently) and potentially usable in diagnosing sleep and mental disorders ((Dzogang et al., 2017)). Biologists and biophysicists have identified and measured what they call “Twitter social jet lag” - the lull in activity when most users are asleep, and noted how it shifts depending on the time of year, the day of the week, and even the east-west geography of its U.S. sample ((Leypunskiy et al., 2018)). Even earlier work exists that visualized some of the same descriptive patterns in social media use we produce below, aggregated at the city-level ((Rios and Lin, 2013)).

Our focus in this paper is to build on this work, making it simultaneously both more general and more specific. In terms of generality, we attempt to establish that some basic patterns in social media use are common across multiple geographies, and multiple sub-populations within the same population. This basic descriptive finding is an important step in organizing the way researchers across a wide range of fields and areas think about establishing appropriate baselines for social media use.

However, we also wish to make a specific contribution to the literature, showing that disaggregation of samples (in our case, by political ideology) can yield important findings when we adjust for small changes in baseline behaviors that are similar, but not identical.

The next section describes our data, as well as the methodological approach we use to properly account for each of the concerns noted in this section.

Data and Methodology

In order to fully analyze the questions presented above, we require rich social media data for a variety of subsamples of the chosen population. In this section, we detail how the data was collected and the methodology we used to run our analyses.

The ability to observe and measure patterns in social media activity at an aggregated level requires data of a specific form. First, the social media activity must be tied to an individual for whom we can trace social media usage over time. Second, that individual must be reasonably locatable in a specific temporal milieu (in our case, a time zone). Finally, we must be able to identify the individual in some manner as a member or non-member of the sub-group in which we are interested.

In this paper, we focus on the social media platform Twitter as the source of such data. Twitter represents a specific kind of social media, where individuals interact in shorter stints with larger potential audiences than on platforms that are focused on more immediate acquaintances (compare, for instance, Facebook, where most individuals largely interact with people they know or to whom they can identify clear linkages, or messaging platforms such as What's App, where all communication is limited to small opt-in groups). However, the *timing* of interaction, overall, is not likely to be considerably different across platforms, and we remain confident that the patterns discussed in this paper are roughly similar across a wide variety of social media platforms.[5]

Twitter is also an important platform for political participation, which we leverage in the analyses below. It is a platform used by political candidates, commentators, and the public to exchange a wide range of political communication ((Jungherr, 2016) presents an extensive review of research conducted on political communication on Twitter, and the volume of such research has only grown since), and presents perhaps uniquely among social media platforms an avenue for otherwise weakly tied individuals to organize and communicate effectively (see, e.g. (Larson et al., 2019) and (Valenzuela et al., 2018)). Thus, it presents a good opportunity to analyze political activity at the platform level.

In order to meet the requirements discussed above, the pattern of data collection for each of our subject areas of interest followed roughly the same pattern. For each time zone or geographic area of interest, we used the geoboxing feature in the Twitter API which allowed us to collect a random sample of tweets from a specific area.[6] We then sampled a group of users from these tweets.

To build our time series of tweets, we then set up a collection of the sampled users going forward, as well as collecting their available history.[7] From this pool of users and user histories, we select all users who whose tweet histories extend at least as far back as our specified time

[5]Of course, verifying that these patterns sustain across platform is also a worthy research endeavor, as differences would be very interesting.

[6]Effectively, we drew a box whose dimensions were longitude and latitude combinations around a particular time zone.

[7]For most users, this covers the last 3200 tweets from the time of the API request. If the user has not tweeted 3200 times, it is their full history.

period.[8] This eliminates users that tweet so frequently that they only appear in a portion of our "learning" stage, and disrupt estimation of the baseline rhythm of the group.[9]

It is this data set that we use to produce the results for our primary focus: the establishment of appropriate baselines for social media use. Generally, it might be seen as a weakness to collect samples from geographic areas likely to be much different in social media pickup, as well as composition, as we do here. Indeed, past work has found there to be significant differences not only in the rate of adoption of various social media platforms across the world (see, e.g. (Mocanu et al., 2013) and (Jungherr, 2016)), but differences within each country or locality as to who opts into social media platforms (see, e.g. (Hargittai and Litt, 2012) and (Hargittai, 2020) for discussion of how the American Twitter population differs from the more general population, as well as (Obholzer and Daniel, 2016) and (Daniel and Obholzer, 2020) for discussions of how the structure of Twitter and changing political incentives also shift which elites participate in social media). Thus, we recognize there are likely to be differences between our three major samples in both the type of individual found in each, as well as the purposes for their participation.[10]

However, we consider this a boon for our study, instead of a limitation. We are attempting to establish that certain regularities of social media use bridge across geographies, while also establishing that differences in sub-populations within a geography can be meaningfully exploited to study differences in those sub-populations. In order to test this most robustly, we *require* differences in our populations, as well as at the sub-population level.[11]

[8]For each geography (US, UK, France), we center our investigation on a national election of some importance just to maintain some consistency. For the US, this was the 2016 presidential election that we focus more on below. For the UK, this was the 2017 Parliamentary elections; for the French case, the second stage of the 2017 Presidential Election. In each case, we focus on the period approximately 12 weeks prior to the electoral event.

[9]This reduction in our sample is relatively small for all three samples. In each, fewer than 8% of the users sampled from the geobox were excluded from our analysis. We return to these highly active users in Appendix A, to show that their general behavior roughly tracks that of their less voluminous sample mates.

[10]In the most basic example, adoption of Twitter in the UK and France is far below the level of adoption in the United States, and more likely to be heavily weighted towards the highly educated and the wealthy.

[11]There may be an additional concern about the generalizability of our findings to broader Twitter

In this paper, we establish a sub-population within the American data set along "winning" and "losing" a highly important election. We proxy this with the estimated ideology of the user. Our expectation is roughly that, in the aggregate, more liberal Twitter users will identify as "losers" after Hillary Clinton lost the election, and more conservative users will identify as "winners" after Donald Trump's victory.[12]

In order to code our users ideologically, we rely on previous research and statistical packages by Pablo Barberà and co-authors (Barberà, 2015; Barberà et al., 2015a). This process of classification uses a Bayesian ideal-point method to estimate the ideology of our users based on the accounts they follow. Users are coded as more liberal (conservative) as they follow political and media accounts that have been estimated themselves to be more liberal (conservative) as well as non-political accounts that are predominately followed by liberals (conservatives).[13] This form of estimation is limited to individuals who follow at least one political or non-political elite; those who do not cannot be estimated and fall out of our sample.[14] Our final sample of users with both full tweet histories and ideological scores are then divided into three equal-sized groups, which we roughly consider the "Liberal", "Moderate" and "Conservative" groups.[15]

populations due to the focus on geoboxed tweets, which we know from past work are likely to be from areas that are wealthier, more urban, and younger, among other attributes (see, e.g. (Malik et al., 2015)). We address this concern directly in Appendix B by replicating our analysis in non-geoboxed tweets.

[12]These are not infallible assumptions at the individual level, but measurement error introduced into the sampling in this way should only serve to add noise to our estimations, rather than bias.

[13]The Bayesian ideal-point estimation of ideology is originally introduced in Barberà (2015), while the details behind the use of correspondence analysis to project Twitter users into the latent space that we use in this article can be found in Barberà et al. (2015a).

[14]This reduction in our sample is even smaller than that associated with history length. Here, only approximately 2% of the individuals for whom we have full histories do not qualify to be sorted ideologically and thus, fall out of our sample. From the group originally sampled from the geobox, over 91% ultimately end up in our analysis.

[15]In the sample, these groupings correspond with individuals who are more than 0.8 standard deviations away from the mean liberally, those from .8 standard deviations liberal to 0.67 standard deviations conservative, and those 0.67 standard deviations more conservative than the mean.

Methodology

Above, we hypothesized that the aggregated social media activity of a group should follow normal human circadian rhythms, with fluctuations in use that reflect hormone release[16] and the resultant sleep patterns, as well as traditional schedules that include set times for both employment and entertainment. We also hypothesized that time trends in volume across the Twitter universe may plausibly trickle down to sub-populations of Twitter users, and thus must be accounted for when modeling sub-group activity.

Researchers in biology have a long history of modeling similar patterns and we turn to that field to guide our modeling process. Specifically, we follow the advice in (Klerman and Hilaire, 2007) (referencing in turn (Brown and Luithardt, 1999)) and adopt a modeling procedure that first identifies a general form of the model before specifying possible model equations and parameters. We then solve for these parameters using available data, and perform diagnostic checking on the model's goodness of fit. We then iterate over these steps to produce the best possible model.

We broadly consider the problem one of modeling a time series of behavior with strong associations between observations and their lagged values, and so adopt an approach focused on autoregressive integrated moving average (ARIMA) modeling. In our specific case, we expect there to be patterns in our data at the 24 hour mark associated with a daily rhythm, as well as at the 168 hour mark - associated with particular times of the week being more commonly assigned to "work" vs. "play."[17] We can capture these effects by allowing for seasonality on a 168-hour basis. We also allow for a time trend that effectively "lifts all boats." This accounts for increased Twitter usage over time that we suspect is present not only at the overall level, but in subgroups within the Twitter population.

There are multiple approaches to fitting an ARIMA model, but ultimately the choice of the correct model balances the number of explanatory terms and the fit of the model to

[16] Specifically here we consider melatonin and cortisol.

[17] As an example, we expect social media usage by our subgroup at Thursday, 3 p.m. to be strongly associated with social media usage by our subgroup at Friday, 3 p.m. However, we might expect there to be a "weekend" effect that needs to be modeled, to explain why the relationship between 3 p.m. Sunday and 3 p.m. Monday is different.

the underlying data. We use the tools provided in the forecast package for the R statistical programming language (see (Hyndman et al., 2017)) to automatically search over the possible specifications of a seasonal ARIMA model that optimizes the corrected Akaike information criterion (AICc). We train the model on the 12 week period leading up to each election, and bin each group's number of Tweets by hour.[18] This model helps us make predictions about what behavior in the weeks following the election *would have looked like*, had there been no election and the social media user groups followed their typical pattern. In the next section, we describe the appropriate models for our populations, as well as each of our subgroups, and then move on to discussing how the 2016 American presidential election altered social media use in our three ideological subgroups.

Results

It is perhaps useful to start our analysis by simply presenting the raw data visualized as a time series. In Figure 1, we map the total volume of tweets in our American population on a quarter-hourly basis for the five weeks before the 2016 Presidential election.[19] There is a clear pattern across days (marked by the deep troughs that represent nighttime hours), but with spikes on particular days of high volume.

[18]When we visualize and discuss phase shifts, we require more granularity and restructure the bins to capture 15-minute time blocks.

[19]As in all of the analyses, we adjust for the time zone at the individual level. Thus, tweets in one bin are not actually occurring in the same time period. Rather, a bin represents all the tweets in the Eastern Time Zone at, for instance 3pm EST, all the tweets in the Central Time Zone at 3pm CST, etc., all the tweets at 3pm MST in the Mountain Time Zone, and all the Pacific Time Zone tweets at 3pm PST.

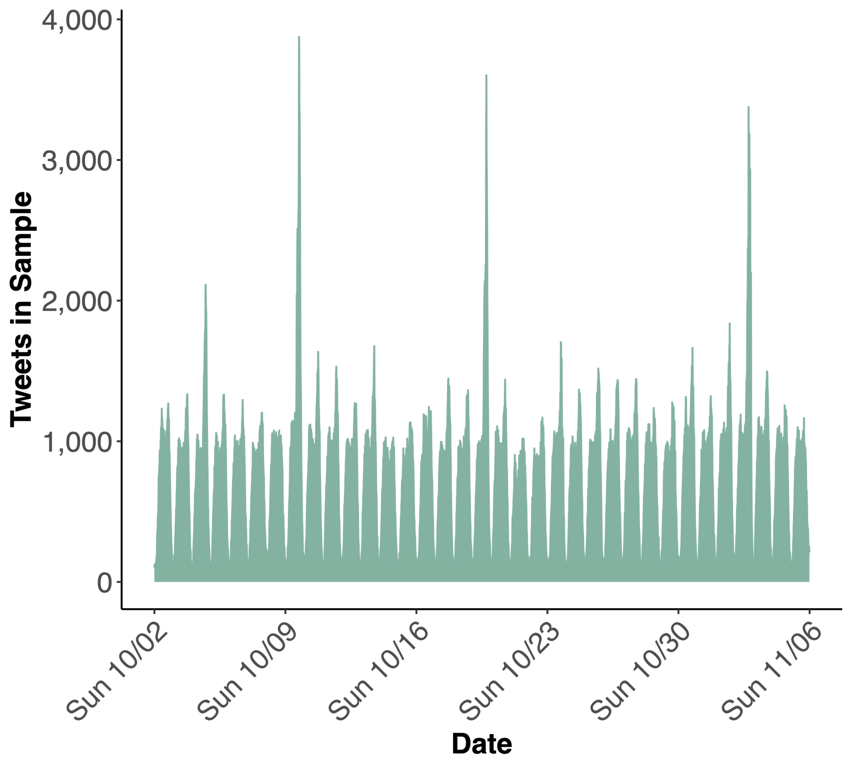

Figure 1. Timeline of tweets leading up to 2016 American election (15m intervals).

Note here that most days do indeed look the same, with dual peaks in the mid morning and early evening. Friday and Saturday evenings, however, differ from the rest of the week, sitting lower in total volume, and with lower evening peaks. This pattern suggests that individuals may not, despite common belief, use social media as a distraction device when enjoying recreation time with friends on the weekends.[20]

We see the same patterns in our international populations. In Figure 2, we provide an alternative display of this cyclicality using our French data. Here, we superimpose the 10

[20]It is worth noting that our analysis precedes the onset of the global Covid-19 pandemic in early 2020, and it would be worth exploring in the future whether or not this observed pattern changed at that time.

weeks prior to the 2017 election onto a single week.[21] As can be seen, the daily pattern remains strong, and the weekend evening reduction in use sustains across the Atlantic.

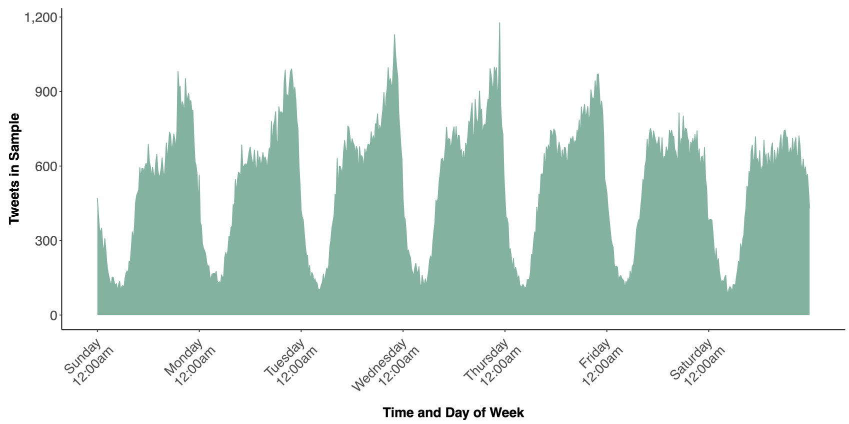

Figure 2. Superimposed weekly timeline (15m intervals), ten weeks to French election.

Finally, in Figure 3, we compare all three countries' 10-week super-impositions on the same graph, with the y-axis now measuring tweets per user in each sample. Note that the *shape* of the activity is nearly identical across all three countries, but that the *volume* per user is vastly different in our American population, compared to the European samples.

[21] So we aggregate all ten 10:00-10:15 Monday morning bins into one bin, all ten 10:16-10:30 Monday morning bins into another, etc.

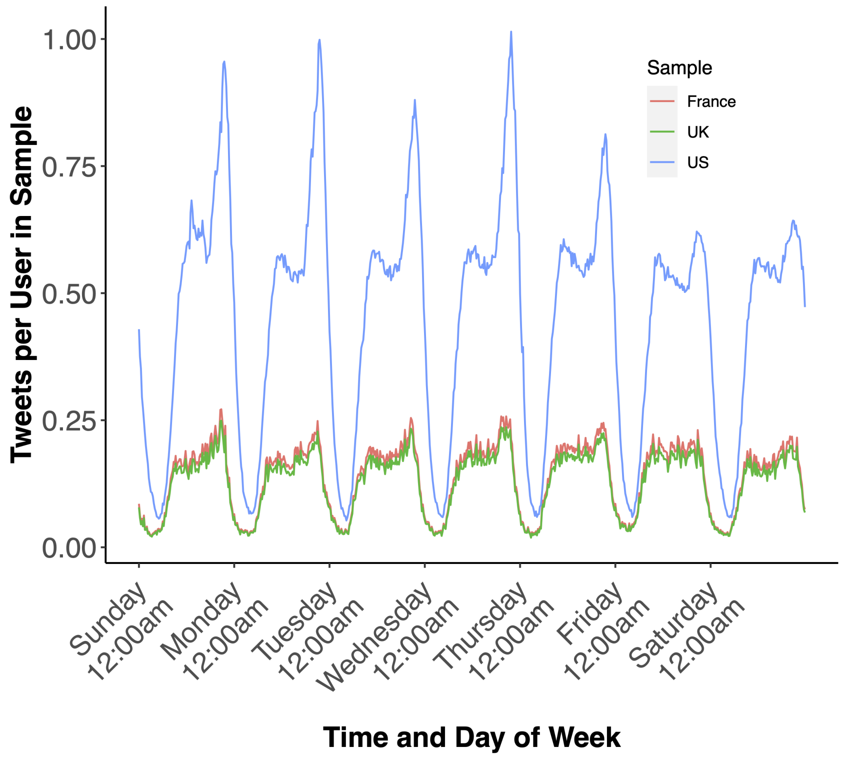

Figure 3. Superimposed timelines (15m intervals), ten weeks to each election.

The users in our American sample tweet at roughly 4-6 times the rate as individuals in our European samples, but they do so *in the same pattern.* This simultaneously validates our expectations of a singular underlying baseline for social media use, while also reinforcing the need to think about how different samples could vary from each other around the analytical edges. Here, there is a huge volume difference between geographic samples, pressed onto the same roughly circadian rhythm.

When we extend our analysis and begin to model both the larger samples and ideological sub-samples of the American case, we do see small differences. In Table 1, we list the best-fitting ARIMA models to the pre-election periods in each of our samples. As discussed previously above in the Methodology section, we search over a broad range of possible models, choosing the model that maximizes the AICc. Thus, each specification listed in Table 1 represents the model that best approximates the underlying distribution of volume of tweets in each population.

Table 1: ARIMA Models with Best Fit for each Population or Subgroup.

Population or Subgroup	Model Specification
French Pop.	$ARIMA(1,0,1)(0,1,0)_{168}$ with drift
UK Pop.	$ARIMA(2,0,1)(0,1,0)_{168}$
American Pop.	$ARIMA(3,0,4)(0,1,0)_{168}$
Liberal Subgroup - US	$ARIMA(3,0,3)(0,1,0)_{168}$
Moderate Subgroup - US	$ARIMA(1,0,0)(0,1,0)_{168}$
Conservative Subgroup - US	$ARIMA(1,0,3)(0,1,0)_{168}$ with drift

To best interpret these models, we focus on the model of best fit for the Conservative subgroup. This is the one of only two models we observe with "drift," which means that there is in fact a statistically significant upward trend in the mean of the series over time. This suggests conservatives in our sample tweeted at an increasing frequency over our pre-election period, even after accounting for the seasonality and auto-correlation in the data.[22]

The rest of the model is characterized by a single first differencing in the seasonal part of the model,[23] as well as a first order autoregressive term and a third order of the moving average part of the non-seasonal aspect of the model.[24] The "168" included in the subscript of the model specification merely designates the period of seasonality. Here, we are measuring observations at the hour, so we use a period of one week (168 hours) in our seasonal model.

While each of the models contains a slightly different weighting of autoregressive and moving average terms (the first and third numbers in the first grouping of three are different), each is heavily seasonal (driven by the weekly timeline), and strongly influenced by auto-regressive terms, as well as (with the exception of the "Moderate" subgroup) the shocks that

[22]Likewise, users in our French sample also tweeted consistently more in the weeks leading up to the election.

[23]This is the "1" in the (0, 1, 0) part of the model.

[24]These are the "1" and "3," respectively, of the (1, 0, 3) part of the model.

propagate throughout the timeline. We leave deeper explanations as to why this may be the case to future research, and now turn to what one can learn from this modeling process.

These models reflect the baseline usage of each of our groups and are used to generate predictions for the post-election period, which for this paper extends to the two weeks after the election. In this way, we can measure exactly how impactful a shock may be to the underlying pattern of usage. More specifically, by comparing the actual number of tweets of our groups to the predictions, we can gauge how the pattern has shifted from a shock, and whether those shifts are different for different subgroups.

Consider first shifts in volume during the immediate period after the shock. We use the 2016 election because it represents both a surprising event (most polls and pundits expected a Clinton victory) that could be expected to affect different subsets of the population differently. And when we dive deeper into our data, we find just this. First, in Figure 4, we present quantile-quantile ("QQ") plots for each of the subgroups, comparing predicted and actual tweets from the post election period. Here, the diagonal lines represent a situation where actual tweets are equal to the number of predicted tweets. While these plots serve mostly as a descriptive look at our data, we can see that in all three groups, the distribution of actual tweets is more dispersed than that of the predicted tweets, and largely above the line. Thus, we should roughly expect to find that each subgroup tweeted more than was predicted from their baselines.

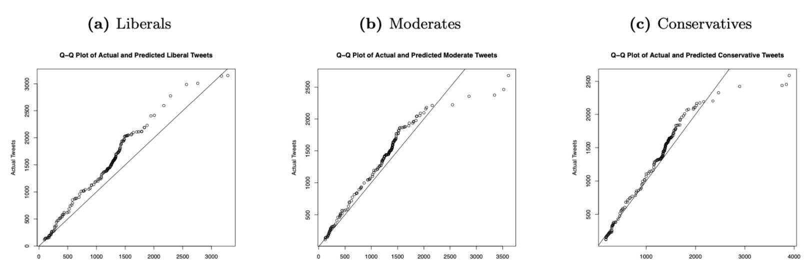

Figure 4. Q-Q plots, by subgroup.

We extend this simple look in Figure 5, where we offer the visualization of traditional two-sample Kolmogorov-Smirnov ("K-S") tests for each subgroup. The generic version of this test measures whether two different samples (in this case, the predicted number of tweets for each group, and the actual tweets sent by the same group) come from the same distribution. The test statistic here is a function of the maximum difference between the cumulative distribution functions, shown in red dotted lines in our figures.

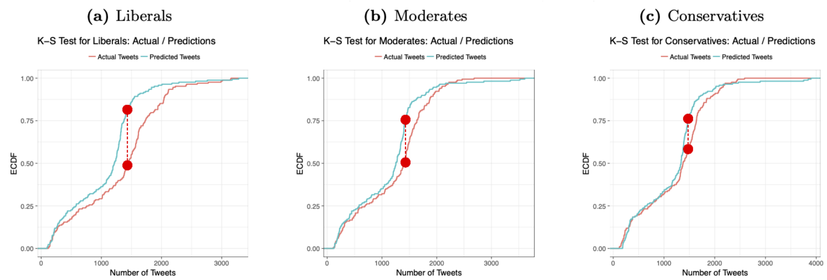

Figure 5. Kolmogorov-Smirnov plots, by subgroup.

We can test more specifically whether the true distribution of our sample of actual tweets is higher than that of the distribution of our predicted tweets with a one-sided test. We do so, and find that in each of our three cases, we can reject the null that our distribution of actual tweets is not greater than our distribution of predicted tweets. Thus, we can confidently say that each of our three groups saw their participation in social media increase in a statistically and substantively significant way.[25] We can also compare the ratios (that is, of actual volume to predicted volume of tweets) for each subgroup. Figures 6 and 7 display the Q-Q and K-S plots for each of the comparisons.

[25] As a robustness check, we repeated the same process of training a model on the 10 weeks of activity prior to two different "random" dates in the lead-up to the American election, and in neither case did we find significant volume differences in the K-S tests for any of our groups.

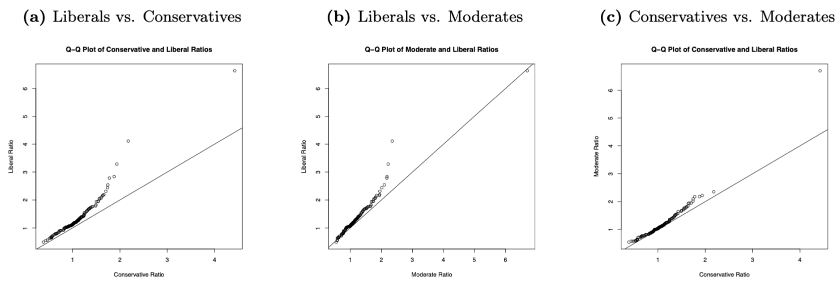

Figure 6. Q-Q plots of comparison between subgroups.

As can be seen in the first two panels within Figure 6, liberal ratios are consistently higher than both conservative and moderate ratios when ordered via their quintiles. This suggests that liberals deviated from their predicted level of tweets (with increased volume) far more than did conservatives and moderates. Both of these groups surpassed their predicted number of tweets, but not nearly to the proportional extent that liberals did. The comparison between moderates and conservative ratios is much closer and difficult to adjudicate without more detailed analyses.

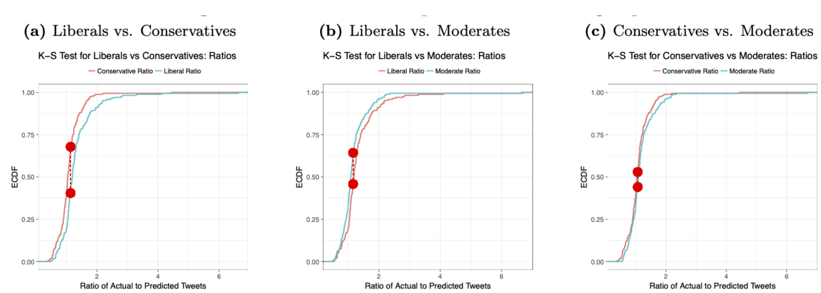

Figure 7. K-S plots of comparison between subgroups.

The K-S graphs display visually what we find in the tests: liberal ratios come from a higher distribution than both conservatives and moderates, but we cannot reject the possibility that the conservative and moderate ratios come from the same distribution.[26] While all groups

[26] All results in this paper that stem from K-S tests find identical support if we choose instead to use Welch two-sample t-tests or Wilcoxon rank sum tests with continuity corrections.

increased their social media use in the post-election period, it was the “losers” who did so the most.

Our second major question pertains to how social media use may shift due to variations in real-life behavior, and specifically how shocks to a group’s natural day-night dynamics might be reflected in social media behavior. We previously identified two daily local maxima, or “peaks,” in social media usage. The first occurs mid-morning, while the second arrives after a traditional workday but before “prime-time.” In order to identify any phase shifts in when our sample was tweeting, we look at the periods surrounding each of these peaks. Consider Figures 8 and 9. Each represents a “heatmap” showing the distribution of tweets in our sample in 15 minute intervals in the 10 weeks before, and 2 weeks after the election (the white bar in each graph represents the day of the election and the day after the election). Each bin is normalized to the day in which it occurs, so any individual rectangle inside the graph represents the ratio of the number of tweets that occurred in a 15 minute time-span (y-axis) on the day (x-axis), divided by the mean number of tweets sent in a 15 minute time-span on that same day. Ratios are represented as reported in the legend, with lighter rectangles reflecting lower ratios, and thus, lower volume of tweets.

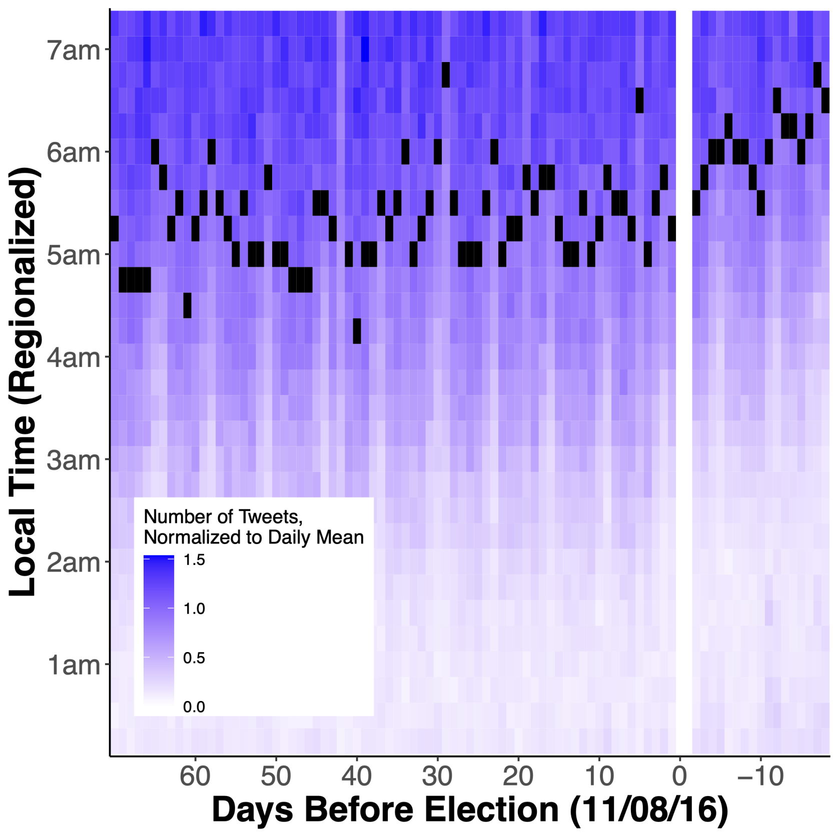

Figure 8. Heatmap of tweet volume in the morning, before and after the election.

Figure 8 displays the heatmap for the time period between midnight and 7:30 am (accounting for time-zones), leading up to the first of the two daily "peaks" in usage. Note the clear upward shift in lighter (and then, darker) colored rectangles in the post-election period. This reflects a fairly substantial phase shift as to when social media use would ramp up in the post-election period. The black rectangles represent the first 15 minute period on each day that the normalized number of tweets surpasses a ratio of 1. These help to draw the reader's eyes to what may not be as starkly clear from the heatmap as a whole: the distribution has shifted to later in the morning, and it remains that way throughout the two week period. By any conservative measure, this phase shift appears to be somewhere between 15 and 45 minutes.[27]

[27]Note that this analysis is somewhat complicated by the fact that Daylight Savings for 2016 ended

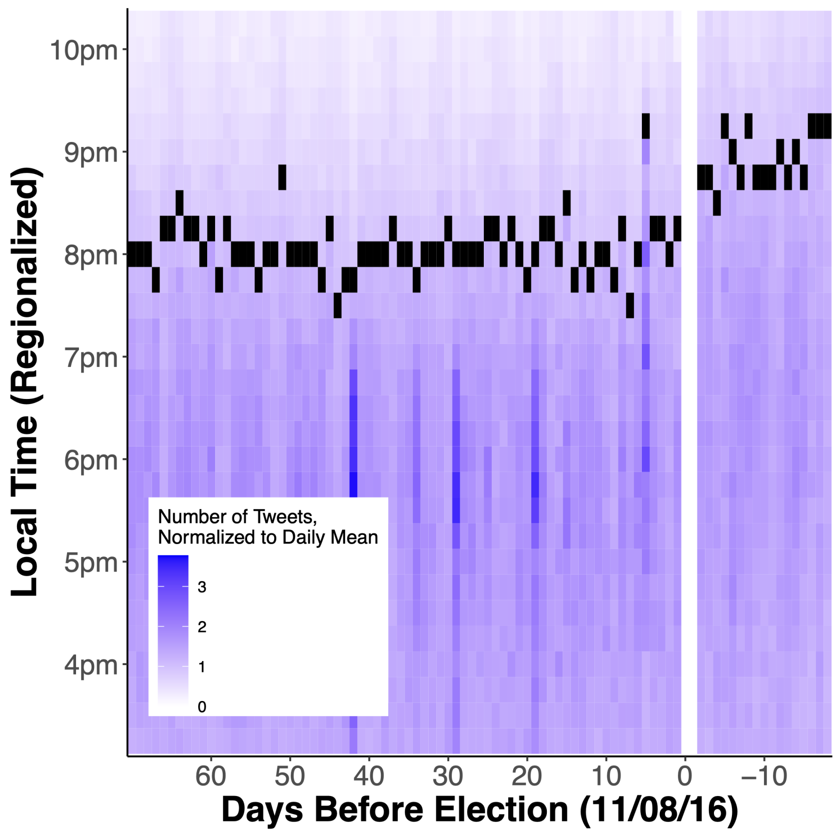

Figure 9. Heatmap of tweet volume in the afternoon, before and after the election.

Figure 9 repeats this visualization for the time period between 3pm and 10:30pm.[28] Here, we not only see the upward shift in darker rectangles (again in the 15-45 minute range), but a widening of the peak time period, cutting into traditional sleep hours. The white range that represented a disengagement with social media in the left-hand side of the graph virtually disappears after the white stripe designating the day of the election.

the Sunday before the election (or Day -2). However, that *should* have moved the distribution "down," as individuals received an extra hour overnight, meaning that they were likely to have bodies adjusted and able to tweet earlier than the listed time. This suggests that the effect of the election is greater than even what is shown here.

[28] Again, adjusted for time zone, with black rectangles now corresponding to the *last* 15 minute period where the normalized value was 1.

Taken together, our findings suggest that individuals in our post election sample seemed more willing to tweet later at night, but also started later in the morning. This had the effect of moving the whole 24-hour distribution (which had the same shape as in the pre-election period, but with greater volume across the board) to the right. We identify this same pattern in each of our ideological subgroups with little difference, suggesting that the entire conversation moved in this fashion, and is not driven by any one particular subgroup. This phase shift in distribution is unlike any other shock in our American or European sample. In all other cases, the immediate effects of a shock wear off in a short period of time (generally less than a day). This case is a good reminder that patterns can be resilient only up to a point, but can be changed by particularly important events. We discuss further implications of all our results below.

Conclusion

Baseline expectations are an important part of any proper analysis, and as researchers continue to shift their focus to the time individuals spend online, this does not change. In this paper, we use a variety of samples and sub-samples to show that social media use on Twitter maintains a steady (largely Circadian) pattern, even across different cultures and geographies. We argue that understanding and properly accounting for this baseline activity is important for inference, both when explaining shifts in use, and using those shifts to explain other outcomes. Finally, we show that while these patterns are resilient over time, large scale events can impact and even change them to some degree, and that these changes may be specific to subgroups within the population, depending on the type of shock you are investigating.

Our findings are mostly methodological and descriptive, but there are some important implications across a variety of substantive topics. First, our example adds to the growing literature on how political outcomes that divide groups into winners and losers can have longer-term knock-on effects to other forms of political participation. Here, we find that, contra other forms of political participation, losers were *more likely* to participate heavily in social media than winners in the aftermath (a finding that will be particularly interesting to revisit after the 2020 US presidential election). This may be due to social media's ability to serve not only as a medium for political participation, but as a gathering place for like-minded individuals to use as an outlet for their concerns, worries, and disappointments. In this way, it plays quite a different role than other political organizations and activities, which are more likely to reinforce the position of the loser in the current political hierarchy. While most of the work in this piece has been descriptive, we encourage further investigation into this dynamic.

Additionally, there are clear implications of our observed phase shift in social media use for public health, among others. Large shifts in daily patterns can be reflective of shocks to mental health, and here the shock was collective and quite large. Generally, shocks to traditional patterns of behavior can result in deteriorating physical health, and reduced productivity. For example, even small disruptions to circadian patterns are thought to increase risk for a number of common cancers ((Davis and Mirick, 2006)). In addition, our analysis is conducted at the aggregate level, but it is a reminder that we can use social media data at the individual level to observe behavior shifts that reflect real-life shifts and shocks we cannot otherwise directly observe.[29]

Having established the expectation of baseline usage, the analyses in this paper also raise additional questions. First, we only observe volume and phase shift here, but just as important is the substance of social media participation. Did the issues talked about change dramatically in the post-election period, or the manner, sentiment, and tone in which they were discussed?

Second, in our paper we find different patterns of social media usage in each of our ideological groups. We adjusted for these different patterns while conducting our analyses, but the existence of different patterns itself is interesting, and worthy of greater study. We know from previous psychology research that there exist elective affinities across different ideological groups, but delving more deeply into how these interact with patterns of social media use would be extremely insightful in a world where polarization seems most acute on these platforms.

Finally, there are platform-specific concerns. In this paper we focused on Twitter, but one might expect that the style and level of interaction on other social media sites may lead to different patterns emerging in their use. This is particularly true on platforms that skew heavily by different demographic groups. Platforms dominated by younger users may operate

[29] Consider, for instance the work in (Jones et al., 2019), using late night social media participation to predict next-day performance among basketball players.

on a different timeline than those dominated by the old, and the introduction of ideology-specific social media sites may see patterns that reflect the differences we discovered in this paper. Thus while accounting for underlying usage dynamics is clearly of methodological value to social media studies, we also hope our analyses here will serve as a launching ground for fruitful substantive empirical studies in the future as well.

Acknowledgments

The Center for Social Media and Politics at New York University is generously supported by funding from the National Science Foundation, the John S. and James L. Knight Foundation, the Charles Koch Foundation, the Hewlett Foundation, Craig Newmark Philanthropies, the Siegel Family Endowment, and NYU?s Office of the Provost and Global Institute for Advanced Study. The authors also wish to thank Yvan Scher and Joanna Sterling for contributions to earlier versions of this project, as well as members of the Center for Social Media and Politics lab who provided helpful comments during the multiple project presentations held there.

References

Anderson, C. J., Blais, A., Bowler, S., Donovan, T., and Listhaug, O. (2005). *Losers' Consent: Elections and Democratic Legitimacy*. OUP Oxford.

Aral, S. and Walker, D. (2012). Identifying influential and susceptible members of social networks. *Science*, page 1215842.

Barberà, P. (2015). Birds of the Same Feather Tweet Together: Bayesian Ideal Point Estimation Using Twitter Data. *Political Analysis*, 23(1):76–91.

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

Barberà, P., Wang, N., Bonneau, R., Jost, J. T., Nagler, J., Tucker, J., and Gonzàlez-Bailòn, S. (2015b). The Critical Periphery in the Growth of Social Protests. *PLOS ONE*, 10(11):e0143611.

Bond, R. M., Fariss, C. J., Jones, J. J., Kramer, A. D. I., Marlow, C., Settle, J. E., and Fowler,

J. H. (2012). A 61-million-person experiment in social influence and political mobilization. *Nature*, 489(7415).

Brown, E. N. and Luithardt, H. (1999). Statistical Model Building and Model Criticism for Human Circadian Data. *Journal of Biological Rhythms*, 14(6):609–616.

Daniel, W. T. and Obholzer, L. (2020). Reaching out to the voter? campaigning on twitter during the 2019 european elections. *Research & Politics*, 7(2):2053168020917256.

Daniller, A. (2016). Can citizens care too much? Investment in election outcomes and perceptions of electoral legitimacy. *Electoral Studies*, 44:151–161.

Davis, S. and Mirick, D. K. (2006). Circadian disruption, shift work and the risk of cancer: a summary of the evidence and studies in seattle. *Cancer causes & control*, 17(4):539–545.

Dzogang, F., Lightman, S., and Cristianini, N. (2017). Circadian mood variations in twitter content. *Brain and neuroscience advances*, 1:2398212817744501.

Feezell, J. T. (2018). Agenda setting through social media: The importance of incidental news exposure and social filtering in the digital era. *Political Research Quarterly*, 71(2):482–494.

Flickinger, R. S. and Studlar, D. T. (2007). One Europe, Many Electorates?: Models of Turnout in European Parliament Elections After 2004. *Comparative Political Studies*, 40(4):383–404.

Gil de Zúñiga, H., Garcia-Perdomo, V., and McGregor, S. C. (2015). What is second screening? exploring motivations of second screen use and its effect on online political participation. *Journal of communication*, 65(5):793–815.

Gonzàlez-Bailòn, S., Borge-Holthoefer, J., Rivero, A., and Moreno, Y. (2011). The dynamics of protest recruitment through an online network. *Scientific reports*, 1:197.

Gonzàlez-Bailòn, S. and Wang, N. (2016). Networked discontent: The anatomy of protest campaigns in social media. *Social networks*, 44:95–104.

Hargittai, E. (2020). Potential biases in big data: Omitted voices on social media. *Social Science Computer Review*, 38(1):10–24.

Hargittai, E. and Litt, E. (2012). Becoming a tweep: How prior online experiences influence twitter use. *Information, Communication & Society*, 15(5):680–702.

Hyndman, R., O'Hara-Wild, M., Bergmeir, C., Razbash, S., and Wang, E. (2017). forecast: Forecasting Functions for Time Series and Linear Models.

Jones, J. J., Kirschen, G. W., Kancharla, S., and Hale, L. (2019). Association between late-night tweeting and next-day game performance among professional basketball players. *Sleep Health*, 5(1):68–71.

Jost, J. T. (2017). Ideological Asymmetries and the Essence of Political Psychology. *Political Psychology*, 38(2):167–208.

Jost, J. T., Stern, C., Rule, N. O., and Sterling, J. (2017). The politics of fear: Is there an ideological asymmetry in existential motivation? *Social cognition*, 35(4):324.

Jungherr, A. (2016). Twitter use in election campaigns: A systematic literature review. *Journal of information technology & politics*, 13(1):72–91.

Karp, J. A. and Milazzo, C. (2015). Democratic Scepticism and Political Participation in Europe. *Journal of Elections, Public Opinion and Parties*, 25(1):97–110.

Klerman, E. B. and Hilaire, M. S. (2007). Review: On Mathematical Modeling of Circadian Rhythms, Performance, and Alertness. *Journal of Biological Rhythms*, 22(2):91–102.

Kostelka, F. and Blais, A. (2018). The Chicken and Egg Question: Satisfaction with Democracy and Voter Turnout. *PS: Political Science & Politics*, pages 1–6.

Larson, J. M., Nagler, J., Ronen, J., and Tucker, J. A. (2019). Social networks and protest participation: Evidence from 130 million twitter users. *American Journal of Political Science*, 63(3):690–705.

Leypunskiy, E., Kıcıman, E., Shah, M., Walch, O. J., Rzhetsky, A., Dinner, A. R., and Rust, M. J. (2018). Geographically resolved rhythms in twitter use reveal social pressures on daily activity patterns. *Current Biology*, 28(23):3763–3775.

Malik, M., Lamba, H., Nakos, C., and Pfeffer, J. (2015). Population bias in geotagged tweets. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 9.

Mocanu, D., Baronchelli, A., Perra, N., Gonçalves, B., Zhang, Q., and Vespignani, A. (2013). The twitter of babel: Mapping world languages through microblogging platforms. *PloS one*, 8(4):e61981.

Obholzer, L. and Daniel, W. T. (2016). An online electoral connection? how electoral systems condition representatives' social media use. *European Union Politics*, 17(3):387–407.

Pierce, L., Rogers, T., and Snyder, J. A. (2016). Losing Hurts: The Happiness Impact of Partisan Electoral Loss. *Journal of Experimental Political Science*, 3(1):44–59.

Rios, M. and Lin, J. (2013). Visualizing the" pulse" of world cities on twitter. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 7.

Tucker, J. A., Theocharis, Y., Roberts, M. E., and Barberà, P. (2017). From Liberation to Turmoil: Social Media And Democracy. *Journal of Democracy*, 28(4):46–59.

Valenzuela, S., Correa, T., and Gil de Zuniga, H. (2018). Ties, likes, and tweets: Using strong and weak ties to explain differences in protest participation across facebook and twitter use. *Political communication*, 35(1):117–134.

Appendix A: High-volume Twitter Users

One might be concerned that the sample for our analysis is structured in such a way to limit generalizability to the users we ultimately exclude, namely users who tweet so frequently that their histories do not extend far enough back to enter the sample. In this appendix, we attempt to allay those fears, by demonstrating that high-volume tweeters follow the same basic patterns as the individuals in our sample, just at much higher amplitudes (akin to the comparison between the US samples and the French and UK ones).

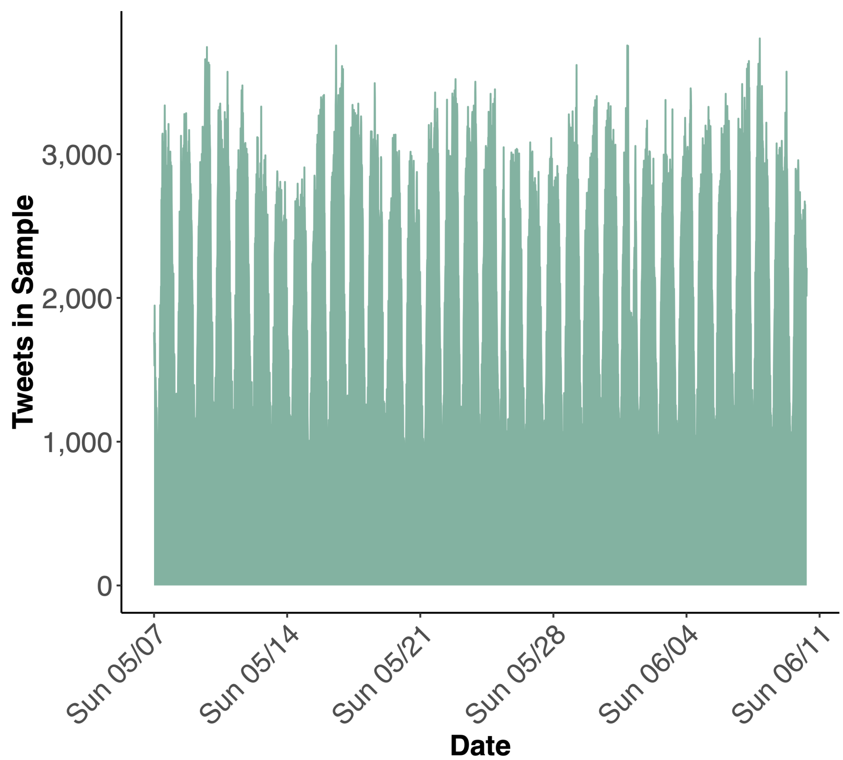

Figure 10. Timeline of tweets by high-volume users in 15-minute intervals.

To do this, we constructed a separate sub-sample of the American data set meant to focus only on high-volume users. After our initial data collection of all histories in May 2017, we also "forward followed" the users for a few months afterward. During this period, we capture all of the tweets users sent, even if that number exceeds 3200. Thus, while this data is not in our time period of interest, it does serve as a period where all types of users (low- and high-volume users) are fully represented. We took this data and selected out the top 1200 users in terms of tweet volume. All 1200 users posted at least 6800 tweets during the period from May 5, 2017 to July 31, 2017, qualifying them as high-volume users of the type that would have been excluded from our analysis sample.

Using this sub-sample, we simply graph the distribution of tweets by high-volume users over a 5 week period, similar to Figure 1. As one can see in Figure 10, there remains a steady pattern in the behavior of high volume users.

In Figure 11, we overlap the data to a single representative week, as we do in Figure 2 of the main paper for the French sample. The pattern is even clearer.

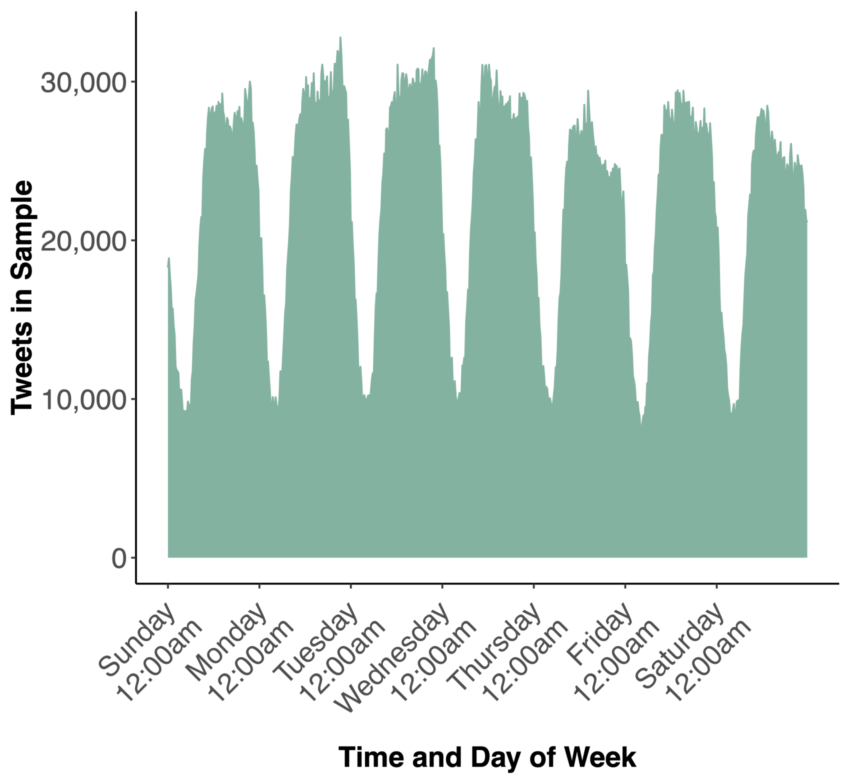

Figure 11. Superimposed weekly timeline (15m intervals ten weeks from 05/07/17 - 07/15/17).

We believe that these two visualizations suggest that high volume tweeters operate in similar patterns to "normal" users, even though some are very likely to be bots or controlled by a process other than an individual interacting on social media in a traditional fashion.

Appendix B: Patterns in Non-Geoboxed Tweets

In order to establish social media usage patterns that are at least in part controlled by time of day, we needed to accurately place individuals in the proper time context. In order to do this, we focused on sampling from geoboxed tweets that allowed us to place a Twitter user in a specific time zone. However, we know that Twitter users who activate geotagging on their accounts are significantly different from those that do not. (Malik et al., 2015) find that individuals that activate geotagging are likely to be from areas with a variety of demographic differences from areas where geotagging is less likely. These demographic differences include higher median incomes, more urbanity, and generally younger populations.

One might worry that our analyses only speak to this specific Twitter sub-population, and may not be repeated in the broader Twitter community. In an attempt to address these concerns, we apply the same type of exploratory analysis to a different group of Twitter users who do not have geotagging activated, but whose presence we have some information about.

For a separate project, we routinely collect the Twitter histories of a random sample of Twitter users. These users were identified by first generating a random number of the size likely to represent a Twitter user's ID number - the number assigned to each user account by Twitter upon account activation. We test whether this account number exists, and when it does, we begin to follow the account and collect its data at regular intervals. This allows us access to a random sample of Twitter users over a long period of time.

We use this random sample in our analysis. Many of the individuals list a location for their account in their profile, information we can use to plausibly locate them within a time zone. While this is an imperfect means of location (individuals sometimes do not list a location, and others list locations that are aspirational, or otherwise inaccurate about their current location), there remains signal in the noise of these imperfections.[30] We scan each user for their profile location, search for this location in the geonames database to verify it is a real place, and then pinpoint that location (where possible) into the appropriate time zone. We finish with 1255 users with profile locations that are real places whose time zones

[30] This general concern with the accuracy of an individual's location is the reason we relied on geoboxed tweets in the first instance.

can be accounted for. It is this data set of users, located in a time zone, that allows us to demonstrate that a broader Twitter audience tends to follow the same rhythmic pattern as the geobox-activating users in our sample.

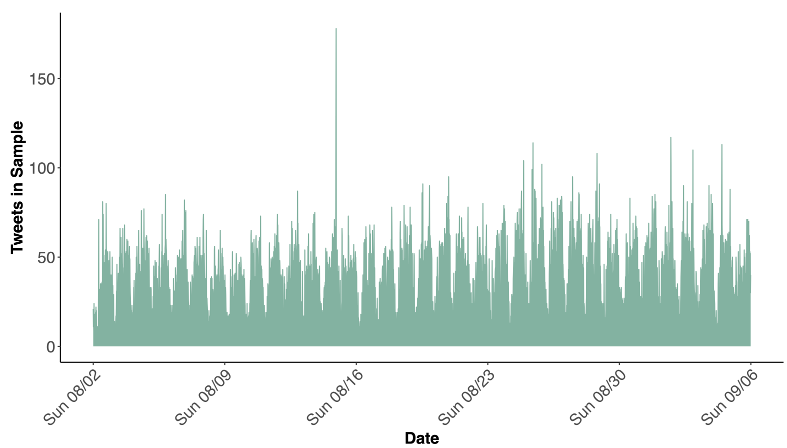

Figure 12. Timeline of tweets by non-geo-boxed users in 15-minute intervals.

Figures 12 and 13 present the two typical basic graphs we use to demonstrate rhythmic patterns in social media use. We note that this sample mimics very closely the general pattern we note in a variety of other sources, suggesting an additional level of robustness to these results, given that we are nearly certain that some users are being "located" in a different time zone than that in which they actually live.

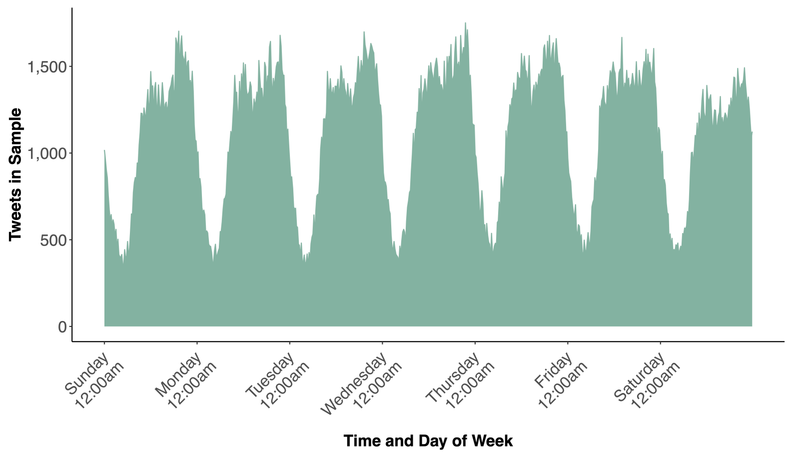

Figure 13. Superimposed weekly timeline (15m intervals ten weeks from 08/02/20 - 10/10/20).