

Hurricanes and hashtags: Characterizing online collective attention for natural disasters

Michael V. Arnold,^{1,*} David Rushing Dewhurst,^{1,2} Thayer Alshaabi,¹ Joshua R. Minot,¹ Jane L. Adams,¹ Christopher M. Danforth,¹ and Peter Sheridan Dodds¹

¹*MassMutual Center of Excellence in Complex Systems and Data Science,
Computational Story Lab, Vermont Complex Systems Center,
and the Department of Mathematics and Statistics,
University of Vermont, Burlington, Vermont 05405, USA*

²*MassMutual Data Science, Boston, Massachusetts 02110, USA*

(Dated: April 1, 2020)

We study collective attention paid towards hurricanes through the lens of n -grams on Twitter, a social media platform with global reach. Using hurricane name mentions as a proxy for awareness, we find that the exogenous temporal dynamics are remarkably similar across storms, but that overall collective attention varies widely even among storms causing comparable deaths and damage. We construct ‘hurricane attention maps’ and observe that hurricanes causing deaths on (or economic damage to) the continental United States generate substantially more attention in English language tweets than those that do not. We find that a hurricane’s Saffir-Simpson wind scale category assignment is strongly associated with the amount of attention it receives. Higher category storms receive higher proportional increases of attention per proportional increases in number of deaths or dollars of damage, than lower category storms. The most damaging and deadly storms of the 2010s, Hurricanes Harvey and Maria, generated the most attention and were remembered the longest, respectively. On average, a category 5 storm receives 4.6 times more attention than a category 1 storm causing the same number of deaths and economic damage.

I. INTRODUCTION

The collective understanding and memory of historic events shapes the common world views of societies. In a narrative economy, attention is a finite resource generating intense competition [1–9]. As commerce and communication shift to online platforms, so too has the narrative economy moved to the digital realm. In 2018, over \$100 billion dollars were spent on internet advertising in the United States, nearly overtaking the \$110 billion spent on traditional media advertising—about 1% of the US GDP [10]. Today, social media both facilitates and records an extraordinary percentage of the world’s public communication [11, 12]. For computational social scientists, the migration of parts of the narrative economy to the web continues to present an immense opportunity, as the discipline becomes data-rich [13, 14].

Academics have become interested in narrative spreading around newsworthy events on social media platforms such as Twitter, as increasingly political fights for influence or narrative control are fought by actors as wide ranging from activists and police departments [15], to state censors suppressing discourse internally and state supported troll factories spreading divisive narratives internationally [16–21]. In 2019, the social media platform Twitter boasted over 145 million daily active users [22].

Quantifying the spread of narratives and the total attention commanded by them is a daunting task. Recent

work has made progress in tracking the spread of quoted and modified phrases through the news cycle, and others have worked to identify actant-relationships and compile contextual story graphs from social media posts [3, 23]. In comparison, quantifying attention directed towards a topic, person or event is a somewhat easier task. Rather than identifying actors and identifying what they act on, as is the case for narrative attention, we can simply count mentions of an entity. Since increasing raw attention or number of mentions is often the zeroth order activity in public relations campaigns, quantifying the volume of attention, irrespective of the sentiment or narrative within which the attention is embedded, seems a natural first step [24].

An understanding of attention has typically focused on time dynamics as measured by the number of mentions in a given corpus, explaining either temporal decay of interest or heavy-tailed allocation of attention given to a spectrum of topics through some preferential attachment mechanism. [25–32]. Another group of studies have worked to classify attention time series from social media as either exogenous or endogenous to the system, modeling the functional form of collective attention decay, or determining if spreading crosses a critical threshold [33–36]. While these studies have typically focused on scientific works, patents, or cultural products such as movies, the rise of large social media datasets have enabled the investigation of a wider range of topics in online public discourse [37].

In this study we examine the collective attention focused on hurricanes, using Twitter, which allows us to capture more natural speech intended for human readers as opposed to search terms. Twitter data has been

* michael.arnold@uvm.edu

used to measure shifts in collective attention surrounding exogenous events like earthquakes by looking for jumps in the Jensen-Shannon divergence between tweet rate distributions between days, or creating real-time earthquake detection using keyword based methods [38, 39].

Here, we use collective attention in a more narrow sense. Instead of looking for anomalous tweet rates, we study n -gram usage rates for hashtags and 2-grams associated with individual events. Specifically, we examine the usage rates of hashtags and 2-grams matching the case-insensitive pattern “#hurricane*” and “hurricane*”, respectively. Natural disasters provide an ideal case study, since they are generally unexpected, producing the signature of an exogenous event. However, the volume of attention given to any particular hurricane varies widely across several orders of magnitude, as does the severity of the storm in terms of the lives lost and damages caused.

Prior efforts have examined the attention received by disasters by type and location, as measured by time devoted on American television news network coverage, and striking discrepancies: for example, to have the same estimated probability of news coverage as a disaster in Europe, a disaster in Africa would need to cause 45 times as many deaths [40]. The same study found that in order to receive equivalent coverage to a deadly volcano, a flood would need to cause 674 times as many deaths, a drought 2,395 times as many, and a famine 38,920 times as many casualties.

Strong hurricanes are more likely to capture attention than weak hurricanes, and hurricanes impacting the continental United States capture much more attention than those failing to make landfall. To what degree does attention shrink when hurricanes make landfall outside of the continental US? The 2017 hurricane season is a particularly stark example, showing that for comparably powerful storms above category 4, those projected to make landfall over the continental United States were talked about nearly an order of magnitude more than Hurricane Maria, which impacted Puerto Rico, and two orders of magnitude more than Hurricane Jose, which never made landfall.

Given the attention received by some hurricanes so unbalanced, we must ask the question: Do government or humanitarian relief resources get dispersed with greater generosity for storms that capture public attention, or are these organizations insulated from popular attention? For the 2017 hurricane season, more money was spent more quickly to aid the victims of hurricanes Harvey and Irma than victims of Hurricane Maria, contributing to the significantly higher death toll and adverse public health outcomes in Puerto Rico [41]. While the attention and policies of government agencies are not usually dictated from Twitter, public attention certainly has some effect on the focus of agencies and allocation of government resources, and recently more attention has been focused on understanding the discourse on social media before, during, and after natural disasters [42–45]

We structure our paper as follows. In Section II, we

outline our methods and data sources, covering the collection of n -gram usage rate data in English tweets as well as data sources for hurricane locations and impacts. In Section III, we examine the spatial associations between hurricanes and the attention they receive, we compute and compare measures of total attention, maximum daily attention, and non-parametric measures of the rate of attention decay for the most damaging hurricanes in the past decade. We present conclusions in Section IV.

II. METHODS

A. n -gram usage rates

We query the daily usage rate of hashtags referencing hurricanes are queried from a corpus of 1-gram—words or other single word-like constructs—usage rate time series, computed from approximately 10% of all posts (“tweets”) from 2009 to 2019 collected from Twitter’s “decahose” [46]. We define usage rate, f , as

$$f(t) = c_\tau(t) / \sum_{\tau' \in \mathcal{D}_t} c_{\tau'}(t),$$

with count, c_τ , of a particular 1-gram divided is by the count of all 1-grams occurring on a given day, \mathcal{D}_t . The usage rates are based only on the usage rate of 1-grams observed in tweets classified as English by FastText, a language classification tool [47, 48]. Our usage rate data set includes separate usage rates for 1-grams in “organic” tweets, tweets that are originally authored, as well as usage rates of 1-grams in all tweets (including retweets and quote tweets). More details about the parsing of the Twitter n -gram data set are available in [49].

For the purpose of studying attention, our usage rates are derived from the corpus with all tweets, including retweeted text, to better reflect not only the number of people tagging a storm, but also the number of people who decide the information contained therein was worth sharing.

We studied the usage rate of 1-grams exactly matching the form “#hurricane*”, where * represents a storm’s name. We also measured the usage rate of 2-grams matching the pattern “hurricane *” for each storm name. All string matching is case-insensitive.

For the ten years covered by the HURDAT2 dataset overlapping with our Twitter dataset, there have been 75 storms reaching at least category 1 in the North Atlantic Basin. Within our 10% sample of tweets, we count over all storms a total of 1,824,842 hashtag usages within a year of each storm, and 3,643,411 instances of the matching 2-gram.

B. Deaths, damages, and locations

To augment our usage rate data set, we downloaded data associated with all hurricanes in the North Atlantic basin from 2008 to 2019 from Wikipedia [50]. Included in the Wikipedia data are the damage estimates (US\$) and deaths caused by each storm, as well as the dates of activity and areas effected. We also used the HURDAT2 data set containing the positions and various meteorological attributes of all North Atlantic hurricanes from 1900 to 2018 for the spatial component of this work [51]. For the time range overlapping with the Twitter derived data set, HURDAT2 has 3 hour resolution.

III. RESULTS AND DISCUSSION

A. Hurricane Attention Maps

In Fig. 1, we show hurricane positions as well as their hashtag usage rate timeseries with a time series indicating the usage rate of the hashtag of the form `#hurricane*`.

We plot the same hashtag usage rate time series below on both linear and logarithmic axes, as well as 2-gram usage rates. For clarity, we only include hurricanes reaching at least category 4.

The hurricane map tracks are meant to show the spatial dependence of attention given to hurricanes, while giving enough visual cues to connect locations along the path to the time the attention was observed. We generated the map shown in Fig. 1 by filling in the polygon defined by the set of points lying at the end of a line segment of length proportional to the smoothed usage rate of the related hashtag, along the vector normal to the current velocity of the hurricane, and centered at the hurricane position at the given time.

Our hashtag usage rate is at the day scale, while HURDAT has 3 hour resolution, so the wrapped attention volume is smoothed with a moving average with a window size of one day to avoid discontinuous jumps. This method obscures any sub-day scale resolution on the map, which could be related to the daily fluctuation of tweet volume as well as varying interest in the hurricanes. While we lose some granularity using daily usage rates, the decays in attention are spread out over days and weeks for smaller storms, and months for larger storms. Daily resolution is sufficient to capture the longer decays in attention, which are our primary interest.

Examining the map, we can see the minimal attention paid to Hurricane Harvey as it traveled across the Caribbean sea and made landfall in Mexico. It is only after crossing the Gulf of Mexico that the hashtag registered on our instrument, and only when it was about to make landfall over Texas did the hashtag usage rate approach its maximum rate, approximately 3 of every 10,000 1-grams in English tweets. It appears that the devastation wrought by Harvey primed hurricane-related

conversation, as the next hurricane, Irma was talked about long before it made landfall. While Irma was talked about with a similar usage rate as Harvey as it impacted Puerto Rico, Hispaniola, and Cuba, it spiked while making landfall in the Florida keys.

Comparing the attention generated by the previous two storms, Hurricane Maria generated substantially less hashtag usage. The peak of its attention gathered as it made landfall over Puerto Rico as a category 4 storm, with less than a fifth of the attention as the hurricanes making landfall on the US. Part of the reason may be due the affected area being Spanish speaking, while our hashtag usage measurement only counts occurrences in English tweets. We find that usage rates of the 2-gram “Huracn Maria” in Spanish tweets were also lower than the usage rates for “Huracn Irma”, but comparable to those for “Huracn Harvey.” See Fig. S1 to compare top hurricane related 2-gram time series for the 2017 hurricane season in English and Spanish.

Another potential contributing factor for the low volume of Hurricane Maria tweets could be that Puerto Rico’s electric grid was destroyed and 95% of cell towers were down in the aftermath of the storm, making it impossible for those directly affected to communicate about the storm [52]. Unfortunately, due to Twitter’s usage norms in this time period, we do not have locations for the vast majority of tweets. The number of people affected by the storms could also help explain the different levels of attention, as both Hurricane Harvey and Irma affected 19 million people, while Maria affected about 4 million [53].

B. Hurricane Attention Comparison

To compare the variation in attention received by different storms, we combined measurements of the hashtag usage rate with deaths and damages caused by each storm from 2009 to 2019. The supplementary materials, Section S1, shows these raw measured values for the most damaging hurricanes in this period.

In Fig. 2, we show radar plots (radial, categorical charts) comparing six measurements of impact and attention for each of the eight most damaging hurricanes in the time period of study [54]. Included measurements are:

- Max Usage Rate—peak attention on any single day
- Integrated Usage Rate—total attention over the entire hurricane season
- Quantile 0.9: $Q_{0.9}$ —days to 90% attention
- Quantile 0.99: $Q_{0.99}$ —days to 99% attention
- Damage—total damage caused by the storm in US dollars
- Deaths—total deaths associated with the storm (both direct and indirect)

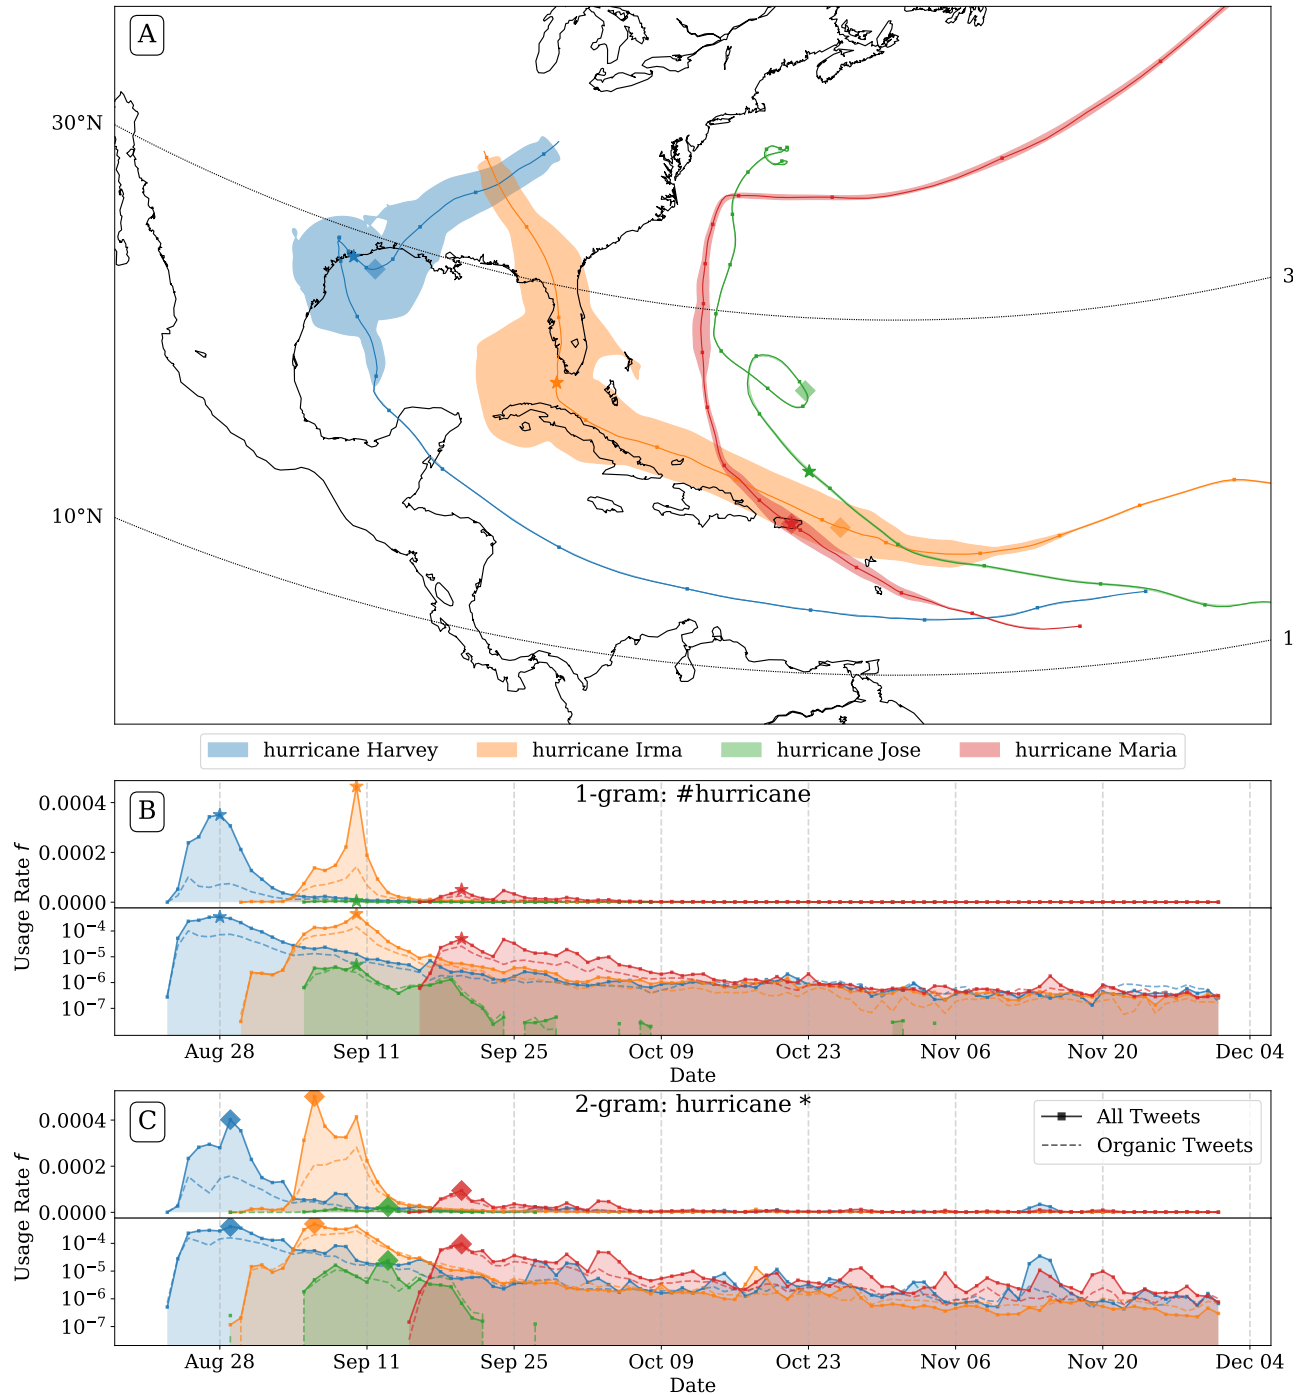


FIG. 1. **Hashtag attention map and usage rate time series** for 1-grams matching the case-insensitive pattern “#hurricane*” for all four hurricanes reaching at least category 4 in the 2017 hurricane season. Markers along the hurricane trajectory indicate the National Oceanic and Atmospheric Administration (NOAA) reported position for every day at noon UTC. On the map, the smoothed rate of hashtag usage is wrapped in an envelope around the hurricane trajectory in panel A, showing the spatial dependence of attention on Twitter. In the lower two plots, panels B and C, we show the usage rates for hashtags and 2-grams matching hurricane* in English language tweets on linear and logarithmic scales. Usage rates within all tweets are indicated with a solid line, while usage rates in ‘organic’ tweets (tweets that are not retweets), are represented by a dashed line. The day of maximum attention on Twitter is marked with a star or a diamond for hashtags or 2-grams, respectively. Generally, hurricanes making landfall on the continental United States received greater attention than those not making landfall. The hashtag usage rate for hurricanes Harvey and Irma at their maximum were approximately an order of magnitude larger than the maximum hashtag usage corresponding to hurricane Maria, and two orders of magnitude larger than Hurricane Jose.

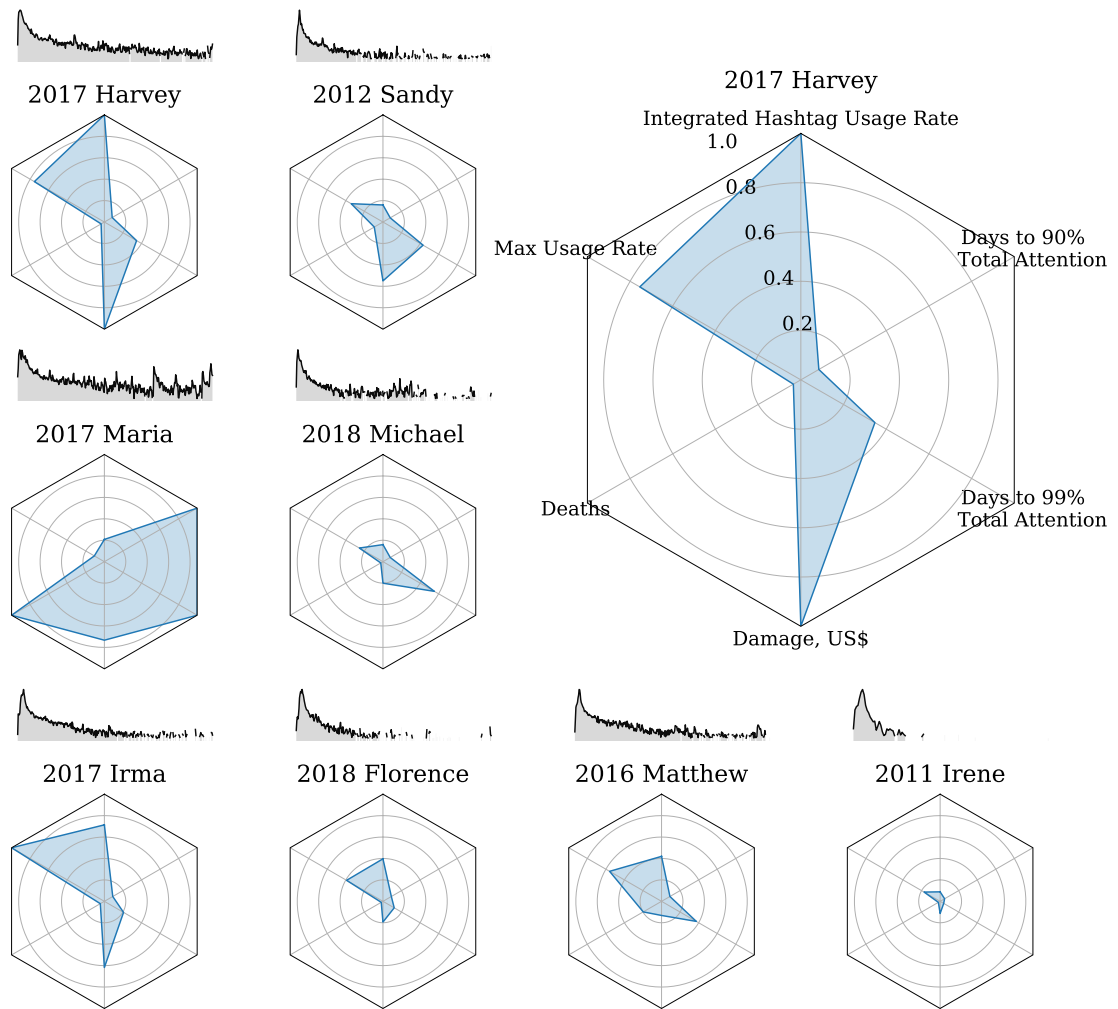


FIG. 2. Radar plots comparing the eight most monetarily damaging hurricanes in the North Atlantic basin from 2009 to 2018. For each plot, starting at the top position and rotating clockwise the measures are: the sum of usage rate of the hashtag, the number of days to reach 90% and 50% of the total attention received during that season, the total cost in dollars attributed to damage caused by the hurricane (in its year), the number of deaths attributed to the hurricane, and maximum usage rate of the hashtag during the year of interest. All measurements are normalized to the maximum value achieved by any hurricane. Hurricane Harvey was the most talked about hurricane, as well as the most damaging. Hurricane Irma was the most talked about on any single day. Hurricane Maria caused the most deaths, and had the longest attention half-life of all measured hurricanes. Raw values for this figure are shown in Section S1. Hashtag usage rate spark lines above each radar plot are normalized to show the common decay shape, and can not be compared to evaluate *relative* volume, and are shown on a log scale.

The relative magnitude of each quantity is shown as a fraction of the maximum value for any storm in the study. The quantile values are non-parametric measurements of the attention time scale—comparable to half-lives but without the assumption of an exponential decay. Some storms receive significant interest months after they pass, usually related to the recovery efforts. Spark lines above

each plot show the attention time series for the year after each storm, as measured by the log usage rate, but do not convey relative scale.

The three most damaging storms, Hurricanes Harvey, Maria, and Irma, all destroyed tens of billions of dollars of property. Storms in Fig. 2 are ordered by damage, with the least damaging being Hurricane Irene in 2011,

which still destroyed an estimated \$14 billion in property.

The most deadly North Atlantic hurricane in the past decade was Hurricane Maria, killing over 3000 people over the course of the extended disaster. The next most deadly storms were Hurricanes Matthew, Sandy, Irma, and Harvey, all killing at least 100 people. Among the storms shown in the Fig. 2, Hurricanes Florence and Irene were the least deadly, causing 58 and 57 deaths, respectively.

The highest hashtag usage rate on a single day was associated with Hurricane Irma, reaching $\max f_{\tau} = 4.6 \times 10^{-4}$, or 4.6 of every 10,000 1-grams, as the storm made landfall over the Florida Keys. Other storms reached comparable single day usage rates, such as Hurricanes Harvey and Matthew, reaching $\max f = 3.5 \times 10^{-4}$ and $\max f = 2.6 \times 10^{-4}$, respectively. Within the top eight most damaging storms, the hashtag associated with Hurricane Maria had the lowest maximum usage rate. The hashtag “#hurricanemaria” appeared only five times for every 100,000 1-grams as Maria made landfall in Puerto Rico.

The highest integrated hashtag usage rate was associated with Hurricane Harvey, followed by Hurricanes Irma, Matthew, and Florence. The integrated hashtag usage rate for “#hurricaneharvey”, $I = 2.3 \times 10^{-3}$. Hashtags associated with Hurricanes Sandy and Irene had the total attention, with $I = 3.7 \times 10^{-4}$ and $I = 2.0 \times 10^{-4}$, respectively.

Due to the extended crisis in the aftermath of Hurricane Maria, the hashtag continued to be used at relatively high volumes even a year after the storm had passed, leading to much larger value for $Q_{0.9}$ of 175 days [55, 56]. Typical values for $Q_{0.9}$ were around 1–4 days, with more prolonged and damaging storms like Harvey in 2017 taking 15 days to reach 90% total attention. In comparison no other storm took longer than 100 days to reach this benchmark. We chose the longer term attention timescale benchmark, $Q_{0.99}$, to describe how long until nearly all storm focused attention has passed. We observe the hashtag associated with Hurricane Maria is the largest for this measurement as well, with $Q_{0.99}$ of 363 days, which should be interpreted as attention not dying away within a year, since we truncate the timeseries after one year. Hurricane Michael, Sandy, and Harvey also have triple digit values for $Q_{0.99}$, as they continued to be talked about, albeit at much lower levels than their peak. Other storms quickly lose attention, such as Hurricane Irene, which took only 12 days to reach 99% total attention.

We observed variation in the overall radar plot shape. More recent storms have been more damaging and deadly, and we find higher measures of total attention and attention decay. A number of storms like Sandy, Michael, and Matthew have relatively higher values for both maximum usage rate and number of days to reach 99% total attention. While there is significant variation in the magnitude of these measurements, the essential exogenous shape of the hashtag usage rate timeseries, f , is consis-

tent.

C. Attention and Impact Regressions by Category

We next explore the associations between damage, deaths, and attention given to hurricanes. In Fig. 3, we show the scaling relationship between attention and impacts for each category storm on the Saffir-Simpson wind scale [57]. Each sub-panel plots the integrated usage rate, $I = \sum_t f(t)$ for hashtag or 2-gram τ , against a measure of storm impact, where t runs over an index of the 365 days after each storm began. I is chosen as a measure of total attention given to the storm during its respective hurricane season, which can be compared across years since it is already normalized to the total volume of conversation on Twitter. Color represents the maximum category storm reached, and the smaller sub-plots are breakout panels for each category. We include Spearman’s ρ , a non-parametric measure of rank correlation, in each panel.

We perform linear regressions on storms in each category separately, a choice that models the attention received by different category storms as separate processes. With models in Section III D, we separately consider attention as a singular process where we account for the hurricane’s maximum category rating using an explicit indicator variable.

1. Model Choice and Fitting Procedure

For each category and each impact, we model total attention as

$$\log_{10} I = a_0 + a_{\text{impact}} X_{\text{impact}} + \varepsilon_{\tau}, \quad (1)$$

where X_{impact} is either \log_{10} deaths or \log_{10} damages caused by each storm. We use a logarithmic model both to capture the scaling relationships between impacts and attention and to inform on the relative changes in attention associated with storm impacts. We offset I by 10^{-8} and the log impacts, X_{impact} by \$10,000 and 0.1 deaths, respectively to avoid divergent log data where observed values are equal to zero.

We set a zero-centered normal prior on the slope of the regression model as $a_1 \sim \mathbf{normal}(0, 1)$. We set a normal prior on the intercept of the model with mean equal to $\log_{10} I = -8$, the minimum value of the offset added to I . We did not have strong beliefs about the likely precision of a_0 since it was not *a priori* clear how much attention would be paid to hurricanes with very little associated monetary damage or few deaths. We thus set a weak hyper-prior on the precision of a_0 , $\tau \sim \mathbf{gamma}(3, 1)$; the intercept of the regression is distributed as $a_0 \sim \mathbf{normal}(-8, \tau^{-1})$.

We found regression coefficients by sampling with the No-U-Turn-Sampler (NUTS), using 8 chains with 2000

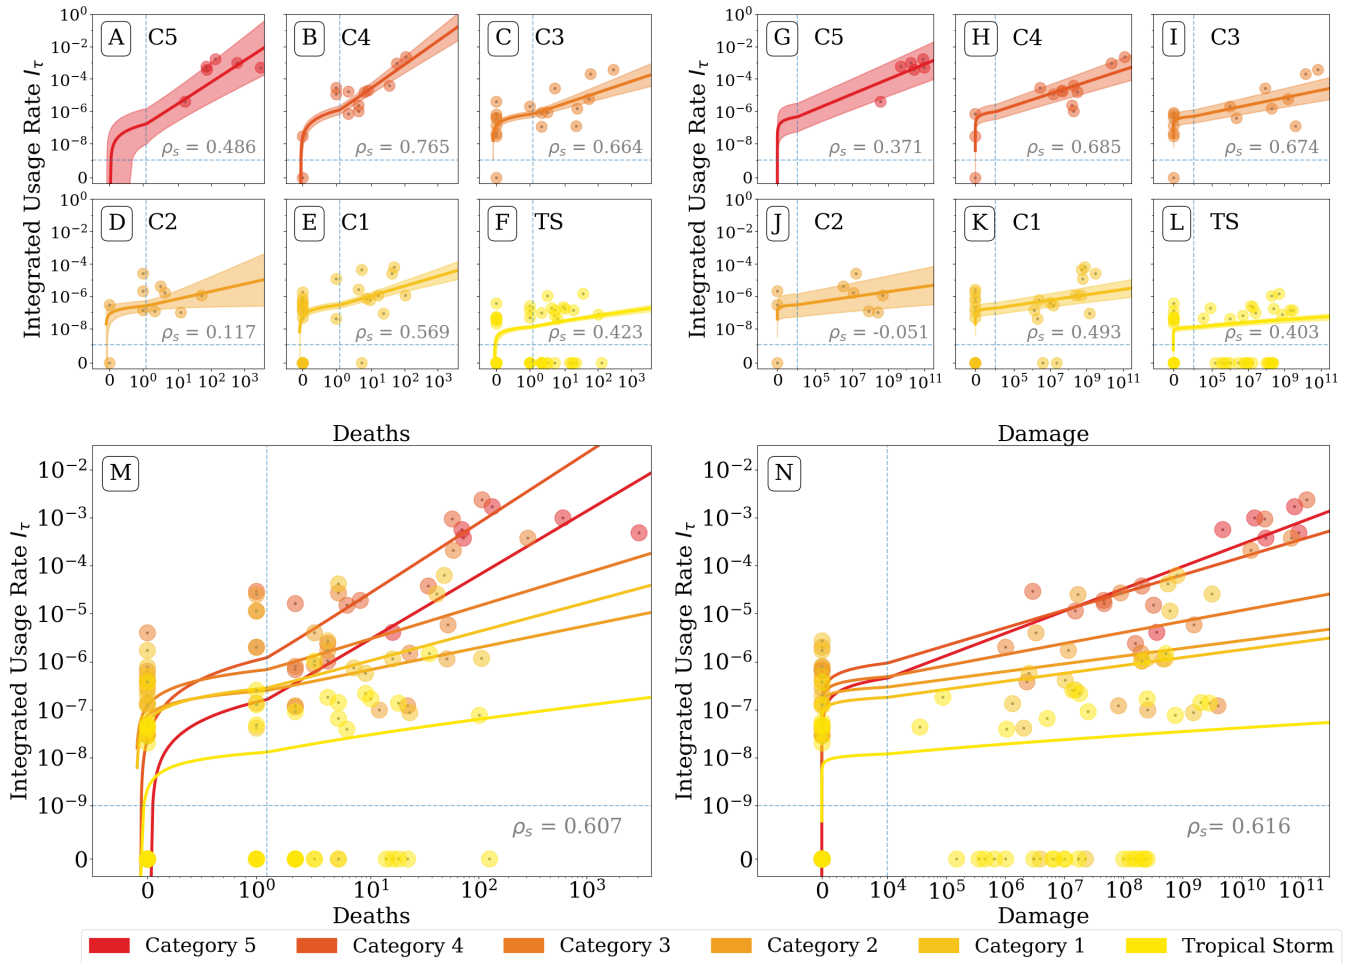


FIG. 3. **Scatter plots for integrated hashtag usage rate versus the deaths and damages caused by each storm.** There is a clear positive association between the total attention represented by hashtags and the impacts of these storms. We reported Spearman's rho, ρ_s , in the top left corner of each plot. While for some categories, there is little evidence for a positive association, for the entire dataset $\rho_s \sim 0.54$. We perform a Bayesian linear regression for each category storm between the log I and log impacts. We show the mean model, along with the credible interval within a standard deviation of the mean model. We use hybrid axis with logarithmic scaling for most horizontal and vertical values and linear scaling near zero, in order to show storms that caused zero deaths or damages, as well as storms for which we measured a hashtag usage rate of zero. Changes in axis scaling occur at the blue dashed lines. Generally, more powerful storms received more attention, higher category storms received more attention even when causing minimal damage, and high category storms had a higher regression slope. These results suggest that for powerful storms, a given increase in impact was associated with a larger increase in attention. While for category one storms a 10-fold increase in deaths is associated with a two-fold increase in attention, for category five hurricanes, this same 10-fold increase in attention is associated with a 27-fold increase in attention.

draws each after 1000 steps of burn-in [58]. Our models converged, with the Gelman-Rubin statistic, \hat{R} , never exceeding 1.004 for any parameter in the 12 models fit.

2. Model Posteriors and Discussion

In Fig. 3, we show the fitted regressions for each category. The size of the impact and attention variables vary over many orders of magnitude, but also include zero values, corresponding to storms that cause no deaths or damage, or had zero usage of the hashtag associated with their name during the year the storm was active.

Note that it should not be surprising that tropical storms appear to receive less attention via our hashtag usage rate measurement, since they never officially become hurricanes, and thus many of the tropical storm hashtags have an integrated usage rate, $I = 0$.

To display all data, we use symmetric log axes: logarithmic for large values and linear for small values. We indicate the switch point from linear to log space axis as blue dotted lines. This choice of axes causes the linear regressions on the log transformed data to appear curved for small values.

In each of the small subplots of Fig. 3, we show the 1σ credible interval for the model as a band around the

mean regression model. The credible interval is noticeably wider for category five storms, which is reasonable given there are only seven storms reaching this category. Generally the mean regression lines are ordered such that higher category storms are receiving more attention than lower category storms. The slopes of the regressions are also higher for higher category storms. However, to better understand the models, we need to compare the model parameters individually.

In Fig. 4 we provide posterior distributions for model parameters, which show that, as expected, more intense storms receive more attention per unit of log impact than weaker storms. For category five storms, we find a mean regression co-efficient of $a_{\text{deaths}} = 1.35 \pm 0.39$, using the format $\mu \pm \sigma$ where μ is the mean and σ is the standard deviation, while for category one storms we find a mean regression co-efficient of $a_{\text{deaths}} = 0.61 \pm 0.18$.

Looking at associations between log damages and log attention we find $a_{\text{damage}} = 0.46 \pm 0.07$ for category 5 storms, while for category one storms we find $a_{\text{damage}} = 0.17 \pm 0.05$.

To interpret the regression coefficients, a_{impact} , as representing proportional increases in attention per proportional increase in impact, we exponentiate the coefficient. Thus, our model shows a 10-fold increase in deaths for a category 5 storm is associated with a 22-fold increase in attention, while for a category 1 storm the same 10-fold increase in deaths is associated with a 4-fold increase in attention.

The intercepts, a_0 , for higher category storms tend to be larger, meaning that for a theoretical minimally disruptive storm causing exactly \$1 of damages or one death, a powerful storm would be talked about more, as shown in Fig. 4. We believe this trend could continue for category 5 storms, but we have observed only $n = 6$ such storms for the duration of our attention dataset. We interpret the intercepts as indications of how much attention low-impact storms receive on average.

In Fig. 4, we fit another regression model on all hurricanes examining log deaths and log attention. We find a 10-fold increase in deaths is associated with a 14-fold increase in attention, since the mean value of $\bar{a}_{\text{deaths}} = 1.16 \pm 0.15$. For damages, coefficients tend to be lower than those for deaths: $\bar{a}_{\text{damage}} = 0.31 \pm 0.05$. We interpret this coefficient as a 10-fold increase in damage being associated with no more than a 2-fold increase in attention.

D. Regression Models for Impacts, Impact Interactions and Hurricane Category

In order to better understand the scaling of attention with hurricane impacts, we fit a number of models on the log transformed data. We applied the same offsets as in the previous section to avoid non-finite log transformed data. We exclude tropical storms, since their attention is not captured in same way as our string matching for

hurricanes.

1. Regression 1

We fit the regression model,

$$\log_{10} I = a_0 + a_{\text{death}} X_{\text{death}} + a_{\text{damage}} X_{\text{damage}} + \varepsilon, \quad (2)$$

where both predictors X are log impacts, which we be referred to as regression 1. The regression coefficients can be interpreted as the increase in log attention received for every unit increase in log impact. Likewise, the intercept can be interpreted as the expected attention for a minimally damaging storm causing one death and \$1 of damage. This model is distinguished from the previous section by including both log impacts in a single model, while not including an interaction term as later models will.

We set priors for the model as shown in Section S1. We chose the intercept, $a_0 \sim \text{normal}(-8, 3)$, to be centered around -8, approximately the lowest usage rate captured in our data, as we guess storms causing 1 death and \$1 worth of damage are talked about relatively little, but wish to allow a wide range of uncertainty spanning a few orders of magnitude. We chose the priors for the regression coefficients $a_{\text{death}} \sim a_{\text{damage}} \sim \text{normal}(0, 1)$, to be weakly informative and centered around zero, as to not bias towards any association. We sampled the coefficients' posterior distributions using NUTS, using 8 chains with 2000 draws each, after 500 steps of burn-in [58]. We found the model converged, with the maximum value of $\hat{R} = 1.000$.

We show the posterior distributions of model parameters for regression one in Panel A of Fig. 5, which have a positive scaling between both deaths and damages, and the amount of attention commanded by the storm, as measured by the log hashtag usage rate. We interpret the mean value of $a_0 = -7.57 \pm 0.5$ for the regression constant as the expected log hashtag usage rate for a minimally destructive storm, i.e., that in English tweets, the hashtag usage rate would integrate to $10^{-7.57}$ over the season. We provide summary statistics in Table S3.

At first glance, this level of attention seems remarkably low: if occurring all in a single day, this is little more than 1 usage for every 100 million 1-grams. The most devastating storms can have integrated usage rates of $I = 2.3 \times 10^{-3}$, five orders of magnitude more attention than our regression constant. However, the least impactful storms affect relatively few people, while the most destructive storms significantly disrupt the lives of tens of millions, so the differences in the scale of total hashtag usage rate are not unreasonable. See Section S1 for measured values corresponding to each storm.

We find $a_{\text{death}} \simeq 0.49$ and $a_{\text{damage}} \simeq 0.24$. Because $10^{0.24} \simeq 1.7$, considering the results in linear space, a 10-fold increase in damages is associated with a 1.7-fold increase in hashtag usage rates, while a 10-fold increase in deaths is associated with a 3-fold increase.

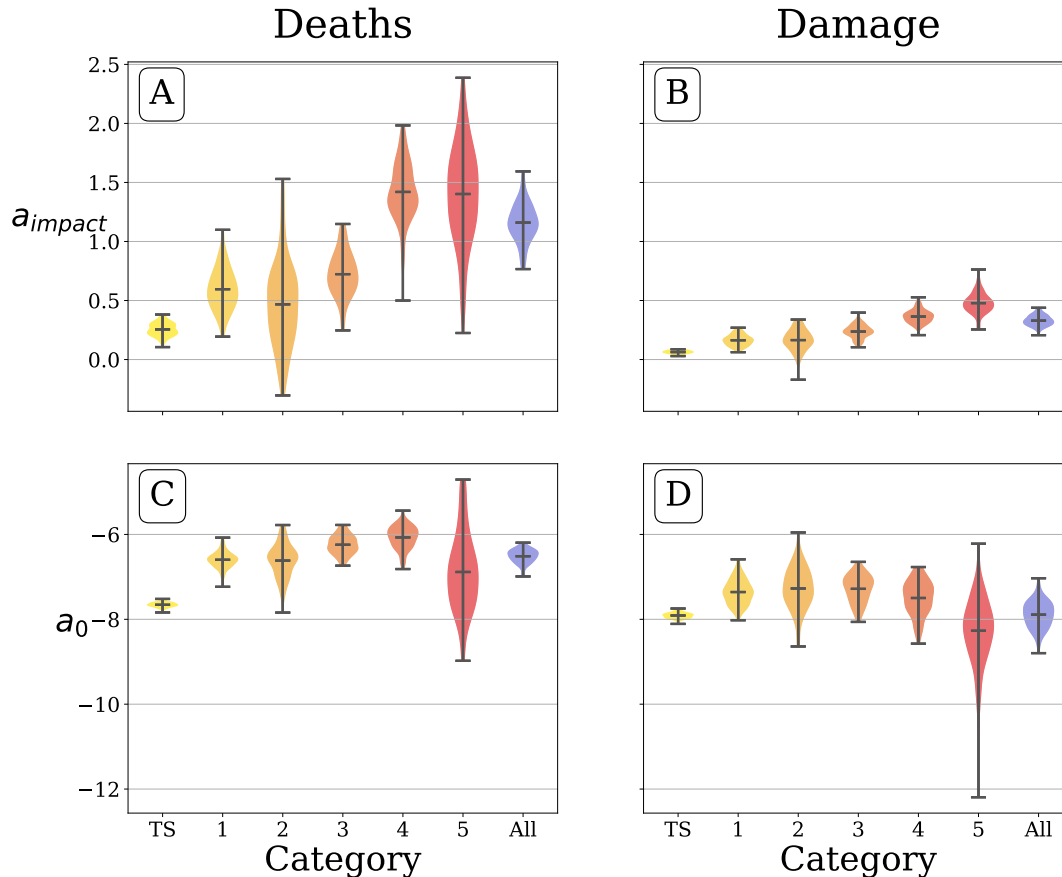


FIG. 4. **Posterior distributions of regression parameters** for the model $\log_{10} I \sim a_0 + a_1 X_i$, where X_i is either the log number of deaths (A and C) or log damages in dollars associated with the storm (B and D), and $\log_{10} I$ is the log integrated hashtag usage rate. The trend in regression coefficients for association between the log attention and log deaths suggests that higher category storms receive more attention per unit impact, while the trend of intercepts shows increasing baseline attention for a hypothetical minimally disruptive storm causing exactly \$1 in damages or one death. For regression coefficients relating log attention to log damages, Category 4 and 5 storms receive more attention per unit increase in log damages than lower category storms. However, the coefficients are smaller in magnitude due to damages varying across 7 orders of magnitude, as compared to deaths varying over 4 orders of magnitude. There is a larger uncertainty for the category 5 intercept values, as only 6 storms of this intensity formed between 2009 and 2019 in the Atlantic basin. At the right of each plot, we show the coefficients for the model fit for all hurricanes (blue violin), excluding tropical storms. Above each category, we show the value of the mean posterior distribution for each parameter. For a table of mean parameter values, see Table S1.

2. Regression 2

For the second regression, an interaction term was introduced between the log number of deaths and the log damages,

$$\log_{10} I = a_0 + a_{\text{death}} X_{\text{death}} + a_{\text{damage}} X_{\text{damage}} + a_{d,D} X_{\text{death}} X_{\text{damage}} + \varepsilon. \quad (3)$$

Prior distributions for the intercept and main effect coefficients are unchanged from regression 1, and we set the prior distribution for the interaction coefficient to be $a_{d,D} \sim \text{normal}(0, 1)$, a standard weakly informative prior for regression coefficients. We used identical fitting

procedures as above, and found the models converged with a maximum value of $\hat{R} = 1.0001$.

Here, the intercept is largely the same as the simplest regression model. Interpreting a_{death} as the conditional relationship between log usage rate and log deaths when total damage is \$1, the $a_{\text{death}} = 0.05$ implies that for a 10-fold increase in deaths is associated with a 1.12-fold increase in hashtag usage rate, though the standard error includes zero. Similarly, $a_{\text{damage}} = 0.22$ implies a 10-fold increase in damage is associated with a 1.6-fold increase in hashtag usage rate. Finally, the interaction coefficient $a_{d,D}$ is small, but positive: a 10-fold increase in $X_{\text{death}} X_{\text{damage}}$ is associated with a 1.14-fold increase in

hashtag usage rate. Notably, the inclusion of the interaction term significantly reduces the regression coefficient associated with deaths, while the coefficient associated with damage is largely unchanged. This provides evidence that storms that cause a large number of deaths and damages are associated with higher volumes of attention, while a storm causing a large number of deaths but relatively less damage will attract much less attention for Twitter users. This leads us to believe that damages could act as a priming factor for human attention, in part explaining why deadly disasters in capital-poor countries often receive less attention than when similarly deadly storms occur in wealthy areas.

3. Regression 3

To better understand the effect of hurricane category on attention, we performed a regression including this categorical variable, modeled as

$$\log_{10} I = a_0 + a_{\text{death}} X_{\text{death}} + a_{\text{damage}} X_{\text{damage}} + a_{d,D} X_{\text{death}} X_{\text{damage}} + \sum_j a_{C_j} X_{C_j} + \varepsilon, \quad (4)$$

where the index j runs from 2 to 5. We did not include a variable for category 1 hurricanes to avoid issues of multicollinearity. Fitting procedures were identical to above, and we found the model converged with the max value of $\hat{R} = 1.0003$.

We did not change priors for the model coefficients from above for existing parameters, and we set the coefficients for category indicator variables to a weakly informative prior, $a_{C_i} \sim \text{normal}(0, 1)$. Since we have included our hurricane categories, the interpretation of the intercept a_0 is now the expected log integrated hashtag usage rate I for a category one hurricane, which causes one death and \$1 of damage. The value is similar to the other regression models. Effect sizes for a_{damage} and $a_{d,D}$ are reduced in magnitude slightly compared to the preceding regression.

As measured by the integrated hashtag usage rate, compared to a category 1 storm causing the same deaths and damages, hurricanes in:

- category 2 receive 1.14 times more attention,
- category 3 receive 1.5 times more attention,
- category 4 receive 5.6 times more attention,
- and category 5 receive 4.6 times more attention.

We show the posterior distributions for regression three in Panel C of Fig. 5.

IV. CONCLUDING REMARKS

We have explored the attention given to hurricanes as measured by the hashtag and 2-gram usage rate. We

quantify the relative volume of attention time series for major storms. We find evidence that not only are more powerful—higher maximum category rating—storms talked about more than weaker storms, but they are talked about more when they inflict the same amount of damage or take the same number of lives. Further, different attention scaling relationships exist for different category storms. For the most destructive storms, we demonstrate that a 10-fold increase in deaths is associated with a 27-fold increase in attention, while for weaker storms the same proportional increase in deaths would lead to only a 3-fold increase in attention on average.

How people outside of the government agencies and non-governmental organizations (NGOs) tasked with responding to natural disasters perceive the importance of disasters have real-world consequences [59, 60]. We hypothesize that monetary donations to NGOs that assist with hurricane disaster relief efforts are strongly associated with the amount of attention attracted by the hurricane. If this is true, it could be advantageous for NGOs to prospect for financial contributions while collective attention is focused most strongly on a storm [61]. It is also possible that the speed and scale of governmental relief programs are influenced by popular attention paid to storms, and previous work has shown that relief has been inequitable in the past [41]. Future work could compare the quantities of non-profit and governmental assistance with attention volume.

While the users of Twitter are certainly not representative of the world, or even English speakers, measuring the text they generate approaches measurement of the population at large, at least more-so than published books or edited newspaper columns [62–66]. The digital signatures left behind by our collective online presence offers rich data for observational studies of everyday language with unprecedented time resolution. Of course, many tweets referencing hurricanes are authored by journalists or news organizations and future efforts could attempt to disentangle the various motivations contributing to the overall usage rate of hashtags and other n -grams.

Another limitation of our work, particularly relevant to any geospatial findings, is that we only consider tweets classified as English. While the density of English speakers closely mirrors the population density for much of the United States, we observe much lower usage rates for the English language hashtags and 2-grams over predominately Spanish speaking areas. While different populations may use different n -grams to reference the same storm, for the purposes of our study we have focused only on the English-speaking population of Twitter.

Future work could consider how to better quantify the total fraction of conversation of Twitter focused on a storm or event of interest. Our current method only includes counts for individual n -grams, which we believe acts as a proxy of total attention, but almost certainly underestimates the total fraction of text devoted to discussing a topic. Hashtag co-occurrence network-based methods could help to identify the most prominent hash-

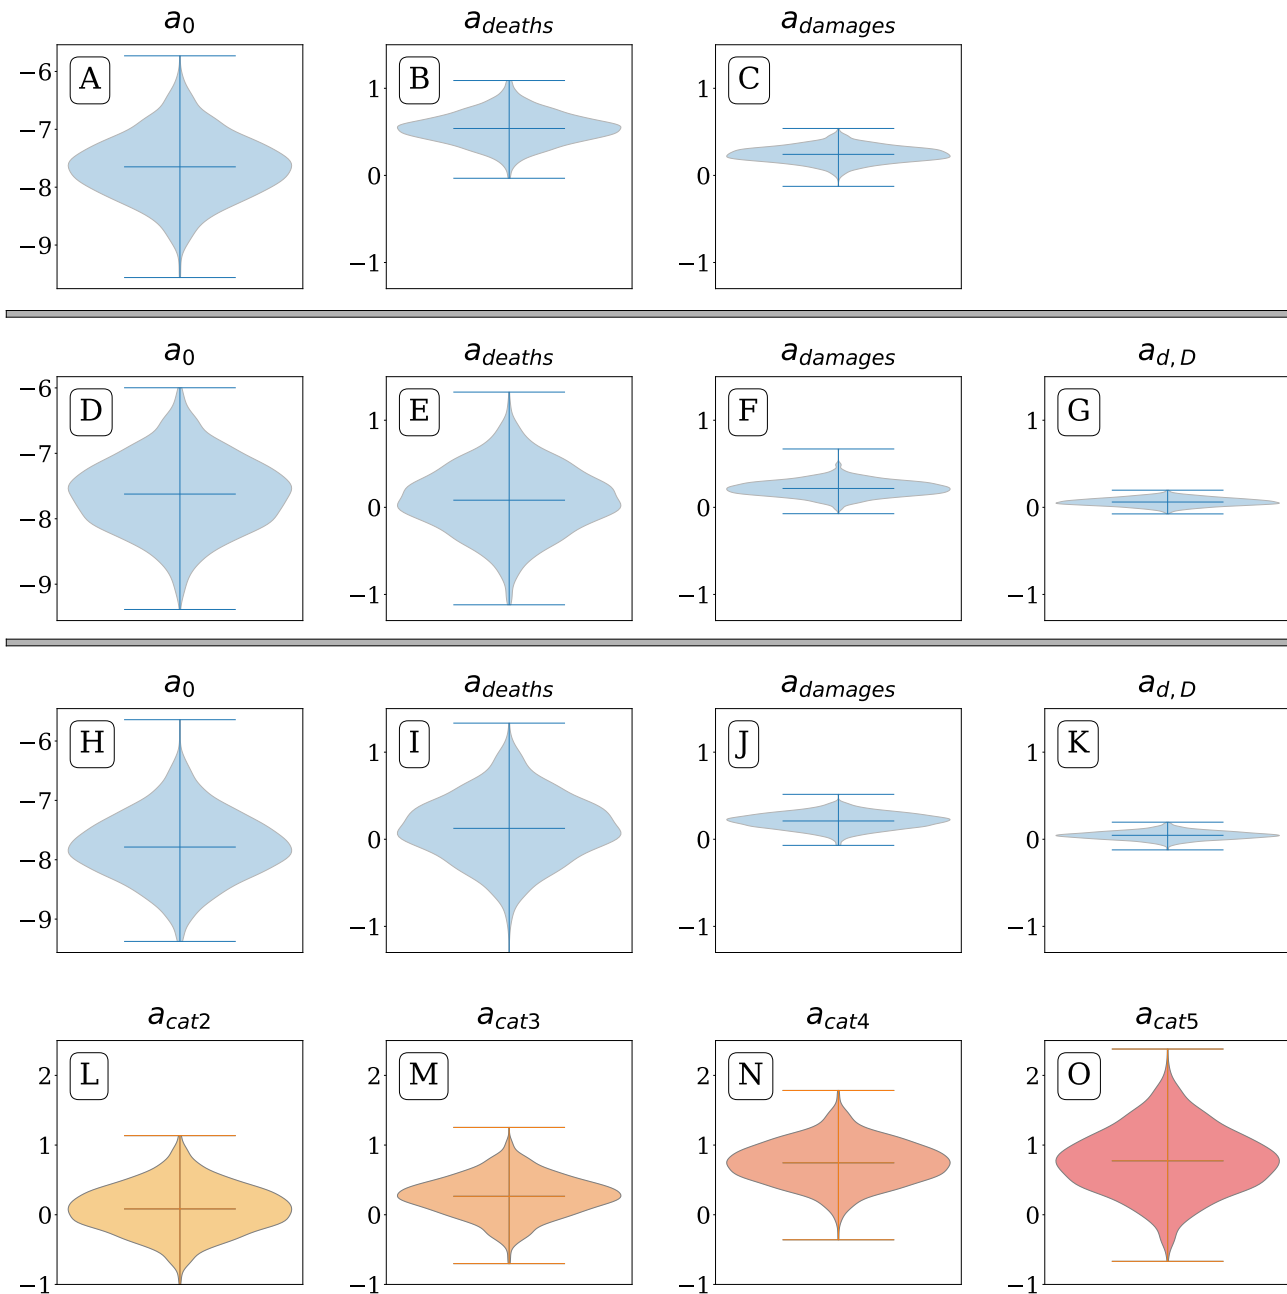


FIG. 5. **Parameter distributions for models 1, 2 and 3** (Sections IIID1 to IIID3). Plots A–C show posterior distributions for regression 1, $\log_{10} I \sim a_0 + a_{\text{deaths}}X_{\text{deaths}} + a_{\text{damage}}X_{\text{damage}}$, plots D–G show distributions for regression 2, $\log_{10} I \sim a_0 + a_{\text{deaths}}X_{\text{deaths}} + a_{\text{damage}}X_{\text{damage}} + a_{d,D}X_{d,D}$, which includes the addition of an interaction term, and plots H–O showing distribution for regression 3, $\log_{10} I \sim a_0 + a_{\text{deaths}}X_{\text{deaths}} + a_{\text{damage}}X_{\text{damage}} + a_{d,D}X_{d,D} + \sum_{j=2}^5 a_{\text{cat } j}X_{\text{cat } j}$, which includes indicator variables for hurricane categories two through five. The addition of the interaction term, $a_{d,D}$ increases posterior variance for a_{deaths} as well as reducing its mean from $a_{\text{deaths}} = 0.49$ in regression 1 to $a_{\text{deaths}} = 0.05$ in regression 2 and $a_{\text{deaths}} = 0.12$ in regression 3, suggesting that while the number of deaths is associated with increased attention, attention response is primed by destruction. Additionally, the hurricane category indicator variables in regression 3 show the progressive increase in attention given to higher category storms compared to category 1 hurricanes.

tags associated with a given storm, or any event of interest, and to classify tweets as relevant. Examining properties of this network changing in time, such as the inte-

grated usage rate of all significant hashtags within one degree could give a more unbiased view of the total attention surrounding the hurricane than our current method.

Other dynamics of hurricanes could be explored in this way, perhaps by encoding Jenson-Shannon Divergence shifts between hashtags as a node attribute [67], or more simply how the most frequently used hashtags in this ego network change in rank over time, as different phases of the storm occur. Authors of previous works studying the effectiveness of NGO hashtag usage following natural

disasters could exploit these network based methods [68].

ACKNOWLEDGMENTS

The authors are grateful for the computing resources provided by the Vermont Advanced Computing Core and financial support from the Massachusetts Mutual Life Insurance Company and Google.

-
- [1] R. J. Shiller. Narrative economics. *American Economic Review*, 107(4):967–1004, Apr. 2017.
- [2] R. Shiller. *Narrative Economics: How Stories Go Viral and Drive Major Economic Events*. Princeton University Press, 2019.
- [3] J. Leskovec, L. Backstrom, and J. Kleinberg. Meme-tracking and the dynamics of the news cycle. In *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 497–506, 2009.
- [4] Z. Tufekci. “Not this one” social movements, the attention economy, and microcelebrity networked activism. *American Behavioral Scientist*, 57(7):848–870, 2013.
- [5] G. Franck. The economy of attention. *Journal of sociology*, 55(1):8–19, 2019.
- [6] A. Humphreys and R. V. Kozinets. The construction of value in attention economies. *ACR North American Advances*, 2009.
- [7] Y. Citton. *The ecology of attention*. John Wiley & Sons, 2017.
- [8] G. Franck. Scientific communication—a vanity fair? *Science*, 286(5437):53–55, 1999.
- [9] A. Nowak, M. Kacprzyk-Murawska, and E. Serwotka. Social psychology and the narrative economy. In *Non-Equilibrium Social Science and Policy*, pages 45–58. Springer, Cham, 2017.
- [10] IAB internet advertising revenue report, May 2018.
- [11] N. Newman. The rise of social media and its impact on mainstream journalism. 2009.
- [12] A. Perrin. Social media usage. *Pew research center*, pages 52–68, 2015.
- [13] Michel, Jean-Baptiste, Shen, Yuan Kui, Aiden, Aviva Presser, Veres, Adrian, Gray, Matthew K, Team, The Google Books, Pickett, Joseph P, Hoiberg, Dale, Clancy, Dan, Norvig, Peter, Orwant, Jon, Pinker, Steven, Nowak, Martin A, and Aiden, Erez Lieberman. Quantitative analysis of culture using millions of digitized books. *Science*, 331(6014):176–182, Jan. 2011.
- [14] E. A. Pechenick, C. M. Danforth, and P. S. Dodds. Characterizing the google books corpus: Strong limits to inferences of socio-cultural and linguistic evolution. *PLoS ONE*, 10:e0137041, 2015.
- [15] R. J. Gallagher, A. J. Reagan, C. M. Danforth, and P. S. Dodds. Divergent discourse between protests and counter-protests: #blacklivesmatter and #allivesmatter. *PLOS ONE*, 13(4):e0195644, Apr. 2018.
- [16] N. N. S. C. of Excellence and 2015. Internet trolling as a hybrid warfare tool: The case of Latvia.
- [17] E. Colleoni, A. Rozza, and A. Arvidsson. Echo chamber or public sphere? predicting political orientation and measuring political homophily in Twitter using big data. *Journal of Communication*, 64(2):317–332, 03 2014.
- [18] A. Gruzdt and J. Roy. Investigating political polarization on Twitter: A Canadian perspective. *Policy & Internet*, 6(1):28–45, 2014.
- [19] P. Barber and G. Rivero. Understanding the political representativeness of Twitter users. *Social Science Computer Review*, 33(6):712–729, 2015.
- [20] D. A. Broniatowski, A. M. Jamison, S. Qi, L. AlKulaib, T. Chen, A. Benton, S. C. Quinn, and M. Dredze. Weaponized health communication: Twitter bots and Russian trolls amplify the vaccine debate. *American Journal of Public Health*, 108(10):1378–1384, 2018. PMID: 30138075.
- [21] V. S. Subrahmanian, A. Azaria, S. Durst, V. Kagan, A. Galstyan, K. Lerman, L. Zhu, E. Ferrara, A. Flammini, and F. Menczer. The DARPA Twitter bot challenge. *Computer*, 49(6):38–46, June 2016.
- [22] S. Salinas. Twitter stock plunges as company blames ad targeting problems for earnings miss. *CNBC*, Oct 2019.
- [23] B. Shahbazi. *StoryMiner: An Automated and Scalable Framework for Story Analysis and Detection from Social Media*. PhD thesis, UCLA, 2019.
- [24] P. S. Dodds, J. R. Minot, M. V. Arnold, T. Alshaabi, J. L. Adams, D. R. Dewhurst, A. J. Reagan, and C. M. Danforth. Fame and ultrafame: Measuring and comparing daily levels of ‘being talked about’ for United States’ presidents, their rivals, God, countries, and k-pop. *arXiv.org*, Sept. 2019.
- [25] S. N. Dorogovtsev and J. F. F. Mendes. Evolution of networks with aging of sites. *Physical Review E*, 62(2):1842–1845, Aug. 2000.
- [26] M. Golosovsky and S. Solomon. Stochastic dynamical model of a growing citation network based on a self-exciting point process. *Physical Review Letters*, 109(9):098701, Aug. 2012.
- [27] S. Valverde, R. V. Solé, M. A. Bedau, and N. Packard. Topology and evolution of technology innovation networks. *Physical Review E*, 76(5):056118, Nov. 2007.
- [28] K. W. Higham, M. Governale, A. B. Jaffe, and U. Zülicke. Unraveling the dynamics of growth, aging and inflation for citations to scientific articles from specific research fields. *Journal of Informetrics*, 11(4):1190–1200, Nov. 2017.
- [29] K. W. Higham, M. Governale, A. B. Jaffe, and U. Zülicke. Fame and obsolescence: Disentangling growth and aging dynamics of patent citations. *Physical Review E*, 95(4):042309, Apr. 2017.
- [30] D. Wang, C. Song, and A.-L. Barabási. Quantifying long-term scientific impact. *Science*, 342(6154):127–132, Oct.

- 2013.
- [31] C. Candia, C. Jara-Figueroa, C. Rodriguez-Sickert, A.-L. Barabási, and C. A. Hidalgo. The universal decay of collective memory and attention. *Nature Human Behaviour*, 3(1):82–91, Jan. 2019.
- [32] P. Lorenz-Spreen, B. M. Mønsted, P. Hövel, and S. Lehmann. Accelerating dynamics of collective attention. *Nature Communications*, 10(1):1–9, Apr. 2019.
- [33] R. Crane and D. Sornette. Robust dynamic classes revealed by measuring the response function of a social system. *Proceedings of the National Academy of Sciences*, 105(41):15649–15653, Oct. 2008.
- [34] J. Lehmann, B. Gonçalves, J. J. Ramasco, and C. Cattuto. *Dynamical classes of collective attention in Twitter*. ACM, New York, New York, USA, Apr. 2012.
- [35] F. Wu and B. A. Huberman. Novelty and collective attention. *Proceedings of the National Academy of Sciences*, 104(45):17599–17601, Nov. 2007.
- [36] D. Kiley. *Characterizing the Shapes of Collective Attention using Social Media*. PhD thesis, The University of Vermont, May 2014.
- [37] R. J. Ladle, R. A. Correia, Y. Do, G. J. Joo, A. C. Malhado, R. Proulx, J. M. Roberge, and P. Jepson. Conservation culturomics. *Frontiers in Ecology and the Environment*, 14(5):269–275, June 2016.
- [38] K. Sasahara, Y. Hirata, M. Toyoda, M. Kitsuregawa, and K. Aihara. Quantifying collective attention from tweet stream. *PLOS ONE*, 8(4):e61823, Apr. 2013.
- [39] T. Sakaki, M. Okazaki, and Y. Matsuo. Earthquake shakes Twitter users: Real-time event detection by social sensors. In *Proceedings of the 19th International Conference on World Wide Web, WWW 10*, page 851860, New York, NY, USA, 2010. Association for Computing Machinery.
- [40] T. Eisensee and D. S. . News droughts, news floods, and US disaster relief. *The Quarterly Journal of Economics*, 2007.
- [41] C. E. Willison, P. M. Singer, M. S. Creary, and S. L. Greer. Quantifying inequities in US federal response to hurricane disaster in Texas and Florida compared with Puerto Rico. *BMJ Global Health*, 4(1):e001191, Jan. 2019.
- [42] D. E. Allen and M. McAleer. President trump tweets supreme leader kim jong-un on nuclear weapons: A comparison with climate change. *Sustainability*, 10(7):2310, 2018.
- [43] M. T. Niles, B. F. Emery, A. J. Reagan, P. S. Dodds, and C. M. Danforth. Social media usage patterns during natural hazards. *PLOS ONE*, 14(2):1–16, 02 2019.
- [44] E. M. Cody, J. C. Stephens, J. P. Bagrow, P. S. Dodds, and C. M. Danforth. Transitions in climate and energy discourse between hurricanes Katrina and Sandy. *Journal of Environmental Studies and Sciences*, 7(1):87–101, Mar 2017.
- [45] M. A. Ahmed, A. M. Sadri, and P. Pradhananga. Social media communication patterns of construction industry in major disasters. 02 2020.
- [46] Q. Li, S. Shah, M. Thomas, K. Anderson, X. Liu, A. Nourbakhsh, and R. Fang. How much data do you need? twitter decahose data analysis. 07 2016.
- [47] A. Joulin, E. Grave, P. Bojanowski, and T. Mikolov. Bag of tricks for efficient text classification. *arXiv.org*, July 2016.
- [48] T. Alshaabi, D. R. Dewhurst, J. R. Minot, M. V. Arnold, J. L. Adams, C. M. Danforth, and P. S. Dodds. The growing echo chamber of social media: Measuring temporal and social contagion dynamics for over 150 languages on twitter for 2009–2020. *arXiv preprint arXiv:2003.03667*, 2020.
- [49] T. Alshaabi, J. Minot, M. Arnold, D. R. Dewhurst, T. Gray, C. Danforth, and P. Dodds. Curating a decade of daily counts of words, phrases, and emojis on Twitter for over 150 languages. 2020, forthcoming.
- [50] Wikipedia contributors. 2010 atlantic hurricane season — Wikipedia, the free encyclopedia, 2019. [Online; accessed 31-October-2019].
- [51] J. Weinkle, C. Landsea, D. Collins, R. Musulin, R. P. Crompton, P. J. Klotzbach, and R. Pielke. Normalized hurricane damage in the continental united states 1900–2017. *Nature Sustainability*, 1(12):808–813, 2018.
- [52] M. Scott. Hurricane Maria’s devastation of Puerto Rico, Jan 2020. [Online; accessed 6. Jan. 2020].
- [53] U. C. Bureau. Hurricanes, Dec 2019. [Online; accessed 4. Dec. 2019].
- [54] Wen-Yuan Liu, Bao-Wen Wang, Jia-Xin Yu, Fang Li, Shui-Xing Wang, and Wen-Xue Hong. Visualization classification method of multi-dimensional data based on radar chart mapping. In *2008 International Conference on Machine Learning and Cybernetics*, volume 2, pages 857–862, July 2008.
- [55] M. O. Román, E. C. Stokes, R. Shrestha, Z. Wang, L. Schultz, E. A. S. Carlo, Q. Sun, J. Bell, A. Molthan, V. Kalb, et al. Satellite-based assessment of electricity restoration efforts in puerto rico after hurricane maria. *PloS one*, 14(6), 2019.
- [56] C. D. Zorrilla. The view from Puerto Rico hurricane Maria and its aftermath. *New England Journal of Medicine*, 377(19):1801–1803, 2017. PMID: 29019710.
- [57] H. T. Taylor, B. Ward, M. Willis, and W. Zaleski. The Saffir-Simpson hurricane wind scale. *Atmospheric Administration: Washington, DC, USA*, 2010.
- [58] M. D. Hoffman and A. Gelman. The no-u-turn sampler: adaptively setting path lengths in hamiltonian monte carlo. *Journal of Machine Learning Research*, 15(1):1593–1623, 2014.
- [59] L. M. Miller. Collective disaster responses to katrina and rita: Exploring therapeutic community, social capital and social control. *Southern Rural Sociology*, 22(2):45–63, 2007.
- [60] R. Burnside, D. S. Miller, and J. D. Rivera. The impact of information and risk perception on the hurricane evacuation decision-making of greater new orleans residents. *Sociological Spectrum*, 27(6):727–740, 2007.
- [61] M. Halloran. Analysis finds disaster relief support swift but short, recurring donors crucial | classy, Sep 2018. [Online; accessed 3. Dec. 2019].
- [62] Mislove, A, Lehmann, S, Ahn, Y Y, and ICWSM, JP Onnela. Understanding the demographics of Twitter users. *aaai.org*, 2011.
- [63] A. Java, X. Song, T. Finin, and B. Tseng. Why we twitter. In *the 9th WebKDD and 1st SNA-KDD 2007 workshop*, pages 56–65, New York, New York, USA, 2007. ACM Press.
- [64] W. Housley, R. Procter, A. Edwards, P. Burnap, M. Williams, L. Sloan, O. Rana, J. Morgan, A. Voss, and A. Greenhill. Big and broad social data and the sociological imagination: A collaborative response. *Big*

- Data & Society*, 1(2):205395171454513, Aug. 2014.
- [65] L. Sloan, J. Morgan, P. Burnap, and M. Williams. Who tweets? deriving the demographic characteristics of age, occupation and social class from Twitter user meta-data. *PLOS ONE*, 10(3):e0115545, Mar. 2015.
- [66] S. Wojcik and A. Hughes. How Twitter users compare to the general public, Apr 2019. [Online; accessed 7. Jan. 2020].
- [67] P. S. Dodds, J. R. Minot, M. V. Arnold, T. Alshaabi, J. L. Adams, D. R. Dewhurst, T. J. Gray, M. R. Frank, A. J. Reagan, and C. M. Danforth. Allotaxonomy and rank-turbulence divergence: A universal instrument for comparing complex systems. *arXiv preprint arXiv:2002.09770*, 2020.
- [68] C. Wukich and A. Steinberg. Nonprofit and public sector participation in self-organizing information networks: Twitter hashtag and trending topic use during disasters. *Risk, Hazards & Crisis in Public Policy*, 4(2):83–109, June 2013.

SUPPLEMENTARY MATERIAL

S1. SUMMARY TABLES FOR REGRESSIONS

Provided for the reader here are tables of summary statistics of the estimated parameters in the regression models in Section III C and Section III D.

S2. 2-GRAM ATTENTION PROPORTION OF “HURRICANE” USAGE RATE

Examining the top 2-grams matching the pattern “hurricane *” in Fig. S1, we can get a sense of what are the top storms during the season, and how much attention is allocated to each at a given time. For English tweets, the first major spike of the 2017 hurricane season is surrounding Hurricane Harvey, though attention also spikes for Hurricane Katrina, in reference to the 2005 storm that affected a nearby region of the gulf coast. As attention begins to decay for Hurricane Harvey, a spike in usage for the 2-gram “hurricane relief” is observed, though it reaches only $f = 3 * 10^{-5}$. Next, attention turns to Hurricane Irma, which reaches the highest 2-gram usage rate of any hurricane in our dataset. Finally, one week after attention for Irma begins to decay, attention spikes for Hurricane maria, though at a level noticeably lower than for Harvey or Irma.

We notice that during storm events the 2-gram usage rates for storms “hurricane *” is often between only half or a fifth the usage rate of the 1-gram “hurricane”, meaning that about one in every 5 times the name of the storm follows the word hurricane in English tweets during active storms.

In Spanish tweets the usage rates of “Huracn Harvey” only reach a maximum of around $f \sim 10^{-4}$, while “Huracn Irma” receives much more relative attention. “Huracn Mara” receives about as much attention as Harvey, and also occupies a space similar to “Hurricane Maria” in English, around $f \sim 10^{-4}$.

S3. BI-EXPONENTIAL DECAYS

To quantify the characteristic time scales of attention given to storms, we examined usage rates by fitting the bi-exponential model introduced by Candia *et al.* [31]. Not all storms receive enough attention, but 50 of 75 in the Atlantic basin recorded at least 6 days of consecutive 2-gram usage within the year of the hurricane, and these storms were had both their hashtag and 2-gram usage rate fit with the bi-exponential model of Candia *et al.* The model here assumes two populations, u and v , which become interested in a given event. Population u , comparable to the general population starts with a peak

interest, and losses attention as $\frac{du}{dt} = -(p+r)u$. During every unit time pu attention is lost from the system and ru is transferred to population v . The dynamics of population v are as follows: $\frac{dv}{dt} = ru - qv$, so attention decays from v with rate q , but increases proportionally to the total attention of population u . The final bi-exponential model is

$$S(t) = \frac{N}{p+r-q} [(p-q)e^{-(p+r)t} + re^{-qt}],$$

and we present the half-lives associated with this model as $\tau_1 = \frac{\ln(2)}{(p+r)}$ and $\tau_2 = \frac{\ln(2)}{q}$, which are the rates of decay from the two populations u and v . The distributions of τ_1 and τ_2 for both hashtag usage rates and 2-gram usage rates are shown in Fig. S3. The mean half-life for population u , the population with faster attention decay, is $\bar{\tau}_1 = 1.3$ days for hashtags, and $\bar{\tau}_1 = 1.1$ days for 2-grams. The decays for population v were not uni-modal, due to some storms regaining attention long after their initial impact, deviating from the model and receiving poor fits, and resulting in very large values of τ_2 , but median values were approximately 24 days. All summary statistics are reported in Section S3. We speculate that for this model the population u is largely people effected by the storm, while population v is largely people writing about the storms or sharing information about the storm response, eg, reporters and non-profit professionals. Further work could look to confirm who is behind the tweets.

The fitting procedure was to first find the maximum value of the usage rate for each storm, before fitting the above model to the decay of log usage rate after this maximum. The resulting fits are shown in Fig. S6 and Fig. S7. The fits generally appear sensible, but there are sometimes issues for noisy time series, where the rate parameter r becomes very small, corresponding to a very long half-life, and misfitting the early decay. This occurs in the time series for Hurricane Florence. The distributions of Mean Squared Error (MSE) are shown in Fig. S5.

Looking at the decay half-lives in Table S3 we notice can see that most hurricane hashtags lose half their volume on the order of 1 or 2 days. The storms with relatively more attention on Twitter, Harvey, Irma, Matthew, and Sandy, all initially decay quickly, with a half-life on the order of a few days, but then have much longer decays associated with τ_2 , on the order of a few weeks. There are some aberrations where the bi-exponential model does a poor job of explaining the data, such as for hurricane Joaquin, where a fight between Governor Bobby Jindal and the Obama administration over the size of a recovery package spurred news stories and attention long after the initial activity associated with the storm itself. This leads to increases in hashtag usage rate, and thus negative half-lives. The longest half-life is associated with hurricane Maria, τ_2 was approximately twice as long as the next largest hurricane. The extended crisis in Puerto Rico caused by Maria may be a reason this exceedingly long lifetime, even though the initial attention received by the

Mean Regression Parameters – Deaths

	Tropical Storms	Cat 1	Cat 2	Cat 3	Cat 4	Cat 5	All Hurricanes
a_{deaths}	0.25	0.61	0.31	0.72	1.39	1.35	1.16
a_0	-7.65	-6.63	-6.58	-6.25	-6.01	-6.91	-6.56

Mean Regression Parameters – Damages

	Tropical Storms	Cat 1	Cat 2	Cat 3	Cat 4	Cat 5	All Hurricanes
a_{damage}	0.06	0.17	0.17	0.24	0.37	0.46	0.31
a_0	-7.91	-7.41	-7.27	-7.21	-7.60	-8.22	-7.92

TABLE S1. Mean Regression Parameters fit for storms of each category. See Fig. 4 for full parameter distributions.

hashtag was less than storms of comparable strength.

We also fit a simple exponential model $S(t) = Ne^{-pt}$. For high attention storms for which we have more than a week of data, this model is unable to capture decays occurring on different time scales, and thus has poor fits. For smaller storms for which attention is lower than the resolution of our data set, the exponential model is perhaps more appropriate. A distribution of half-lives for hashtags and 2-grams is shown in Fig. S4. While for larger storms, the fits did not capture the changing rates of attention decay, it was adequate for smaller storms that decay quickly below our instrument’s resolution. However, for storms for which we have data for an extended

decay, the bi-exponential model is more appropriate.

S4. HURRICANE ATTENTION MAPS

The remaining Hurricane Attention Map and time series from 2009 to 2018 are presented for the reader’s perusal. Only storms reaching at least Category 2 are shown, and Seasons 2013 and 2014 are omitted. Earlier storms in our dataset mostly did not make landfall, and thus appear to receive relatively little attention. The scale of attention on the maps is held constant between years.

a_0	a_{death}	a_{damage}
$\text{normal}(-8, 3)$	$\text{normal}(0, 1)$	$\text{normal}(0, 1)$

TABLE S2. Priors for Regression 1

	mean	sd	mc_error	hpd_2.5	hpd_97.5	n_eff	Rhat
a_0	-7.57	0.52	0.01	-8.60	-6.56	4182	1.0
Deaths	0.49	0.16	0.00	0.16	0.80	4660	1.0
Damage	0.24	0.08	0.00	0.08	0.40	4108	1.0
sd	0.89	0.08	0.00	0.75	1.05	8449	1.0

TABLE S3. Results for Regression 1

a_0	a_{death}	a_{damage}	$a_{\text{d,D}}$
$\text{normal}(-8, 3)$	$\text{normal}(0, 1)$	$\text{normal}(0, 1)$	$\text{normal}(0, 1)$

TABLE S4. Priors for Regression 2

	mean	sd	mc_error	hpd_2.5	hpd_97.5	n_eff	Rhat
a_0	-7.58	0.51	0.01	-8.58	-6.58	8085	1.0
Deaths	0.05	0.34	0.00	-0.65	0.70	8326	1.0
Damage	0.22	0.08	0.00	0.06	0.38	8151	1.0
Interaction	0.06	0.04	0.00	-0.02	0.14	8676	1.0
sd	0.88	0.08	0.00	0.74	1.04	10843	1.0

TABLE S5. Results for Regression 2

a_0	a_{death}	a_{damage}	$a_{\text{d} \times \text{D}}$	a_{C_i}
$\text{normal}(-8, 3)$	$\text{normal}(0, 1)$	$\text{normal}(0, 1)$	$\text{normal}(0, 1)$	$\text{normal}(0, 1)$

TABLE S6. Priors for Regression 3

	mean	sd	mc_error	hpd_2.5	hpd_97.5	n_eff	Rhat
a_0	-7.64	0.51	0.01	-8.60	-6.60	9916	1.0
Deaths	0.09	0.36	0.00	-0.60	0.81	9892	1.0
Damage	0.20	0.08	0.00	0.05	0.35	10580	1.0
Interaction	0.05	0.04	0.00	-0.04	0.13	10424	1.0
Cat2	0.07	0.31	0.00	-0.55	0.66	15415	1.0
Cat3	0.21	0.26	0.00	-0.32	0.72	14877	1.0
Cat4	0.76	0.28	0.00	0.20	1.29	15063	1.0
Cat5	0.66	0.44	0.00	-0.17	1.57	13237	1.0
sd	0.84	0.08	0.00	0.70	1.00	14240	1.0

TABLE S7. Results for Regression 3

	Integrated Frequency	Max Frequency	Deaths	Damage	Quantile 0.99	Quantile 0.9
2017 Harvey	2.3×10^{-3}	3.5×10^{-4}	1071.2×10^{11}		126	14
2017 Maria	4.9×10^{-4}	5.0×10^{-5}	30579.1×10^{10}		363	166
2017 Irma	1.6×10^{-3}	4.6×10^{-4}	1347.7×10^{10}		75	15
2012 Sandy	3.7×10^{-4}	1.5×10^{-4}	2866.8×10^{10}		157	13
2018 Michael	3.7×10^{-4}	1.1×10^{-4}	722.5×10^{10}		201	13
2018 Florence	9.3×10^{-4}	1.8×10^{-4}	572.4×10^{10}		44	15
2016 Matthew	9.9×10^{-4}	2.6×10^{-4}	6031.6×10^{10}		136	15
2011 Irene	2.0×10^{-4}	8.0×10^{-5}	581.4×10^{10}		14	8
2019 Dorian	5.7×10^{-4}	1.2×10^{-4}	$70 \ 4.6 \times 10^9$		36	12
2012 Isaac	2.6×10^{-5}	6.1×10^{-6}	$41 \ 3.1 \times 10^9$		192	97
2010 Alex	5.8×10^{-6}	2.5×10^{-6}	$52 \ 1.5 \times 10^9$		15	7
2017 Nate	6.3×10^{-5}	3.1×10^{-5}	$48 \ 7.8 \times 10^8$		8	5
2019 Barry	1.1×10^{-5}	3.8×10^{-6}	$1 \ 6.0 \times 10^8$		8	4
2016 Hermine	4.1×10^{-5}	1.9×10^{-5}	$5 \ 5.5 \times 10^8$		7	3
2019 Lorenzo	4.1×10^{-6}	1.0×10^{-6}	$16 \ 3.6 \times 10^8$		11	9
2014 Gonzalo	1.5×10^{-5}	6.4×10^{-6}	$6 \ 3.1 \times 10^8$		14	11
2015 Joaquin	3.7×10^{-5}	1.1×10^{-5}	$34 \ 2.0 \times 10^8$		11	5
2017 Ophelia	2.7×10^{-5}	1.2×10^{-5}	$5 \ 8.7 \times 10^7$		15	7
2009 Bill	1.6×10^{-5}	9.4×10^{-6}	$2 \ 4.6 \times 10^7$		11	7
2010 Earl	1.9×10^{-5}	4.9×10^{-6}	$8 \ 4.5 \times 10^7$		8	6
2014 Arthur	2.5×10^{-5}	1.3×10^{-5}	$1 \ 1.6 \times 10^7$		9	5
2016 Nicole	1.1×10^{-5}	5.3×10^{-6}	$1 \ 1.5 \times 10^7$		13	9
2017 Katia	4.0×10^{-6}	1.1×10^{-6}	$3 \ 3.2 \times 10^6$		7	4
2017 Jose	2.9×10^{-5}	4.7×10^{-6}	$1 \ 2.8 \times 10^6$		22	13
2014 Bertha	2.7×10^{-6}	1.1×10^{-6}	4 0.0		11	8
2015 Danny	4.0×10^{-6}	1.8×10^{-6}	0 NaN		6	3

TABLE S8. The unnormalized values associated with radar plots in Section III




FIG. S1. Word usage rate proportions of “hurricane *” in English tweets




FIG. S2. Attention proportions of “Huracán *” in Spanish. We can see that the word usage rate surrounding “Hurricane Maria” captures a similar amount of the total attention for the 1-gram hurricane as “Huracán Mara” captures. Additionally, hurricane Harvey’s 2-gram usage rate is lower in Spanish than in English, while Hurricane Katrina is talked about considerably in English but does not rise about the 50000th most used 2-gram in Spanish. As always, usage rates are case-insensitive.




FIG. S3. **Hurricane decay half-lives:** Distributions of fitted half-lives for the populations u and v . The mean half-lives for $\tau_1 = 1.3$ days and $\tau_2 = 156$ days for hashtags and $\tau_1 = 1.1$ days and $\tau_2 = 241$ days for 2-grams. For τ_2 the median half-lives are also interesting since we suspect the longest half-lives are due to poor fits. For hashtags $\tau_2 = 23$ days, and for 2-grams $\tau_2 = 24$ days.




FIG. S4. **Simple Exponential Hurricane decay half-lives:** Distributions of fitted half-lives for a single population. The median half-lives for $\tau = 5.3$ days a for hashtags and $\tau = 5.2$ days for 2-grams. The simple exponential model fails to explain the break in attention decay for larger storms, receiving more attention. The bi-modal distribution of half-lives for 2-grams suggests that there are two categories of storms, ones with larger half-lives have more data, and thus the longer decay increases the fitted half-life. Meanwhile, smaller storms receive so little attention, that we don't measure any after a week or so, leading to a much smaller half-life, which corresponds to τ_1 in our bi-exponential fit.




FIG. S5. **Decay Model Comparison:** Distributions of Mean Squared Error (MSE). The bi-exponential model has the lowest average MSE, followed by the simple exponential decay. The power law decay fails to capture the dynamics of attention decay, when the fit is compared to the data visually, and is reflected in the higher average MSE.

	Max Usage Rate	τ_1 [Days]	τ_2 [Days]	Season
#hurricanealex	2.5×10^{-6}	0.7	8.6	2010
#hurricanearthur	1.3×10^{-5}	0.9	190.3	2014
#hurricanebarry	3.8×10^{-6}	0.7	16.0	2019
#hurricanebertha	1.1×10^{-6}	0.6	6.9	2014
#hurricanebill	9.4×10^{-6}	0.2	693.1	2009
#hurricanechris	8.9×10^{-7}	0.6	693.1	2018
#hurricane cristobal	2.0×10^{-7}	2.0	6.9	2014
#hurricanedanielle	1.9×10^{-7}	0.7	693.1	2010
#hurricanedanny	1.8×10^{-6}	0.7	6.9	2015
#hurricanedorian	1.2×10^{-4}	1.6	8.8	2019
#hurricaneearl	5.0×10^{-6}	0.4	6.9	2010
#hurricaneeflorence	1.8×10^{-4}	2.8	323.3	2018
#hurricane gert	3.6×10^{-7}	0.4	6.9	2017
#hurricanegonzalo	6.4×10^{-6}	0.9	693.1	2014
#hurricaneharvey	3.5×10^{-4}	2.5	30.6	2017
#hurricanehermine	1.9×10^{-5}	0.8	15.9	2016
#hurricaneida	8.3×10^{-7}	0.8	9.7	2009
#hurricaneigor	2.2×10^{-7}	1.1	693.1	2010
#hurricaneirene	8.0×10^{-5}	0.7	26.5	2011
#hurricaneirma	4.6×10^{-4}	1.0	20.0	2017
#hurricaneisaac	6.1×10^{-6}	0.7	693.1	2012
#hurricanejoaquin	1.1×10^{-5}	1.2	57.7	2015
#hurricanejose	4.7×10^{-6}	2.0	23.1	2017
#hurricanekarl	7.4×10^{-8}	0.6	68.9	2010
#hurricanekatia	8.7×10^{-7}	0.2	6.9	2011
#hurricanelorenzo	1.0×10^{-6}	1.3	64.2	2019
#hurricanemaria	5.0×10^{-5}	4.1	43.4	2017
#hurricanemathew	2.6×10^{-4}	1.4	27.4	2016
#hurricanemichael	1.1×10^{-4}	1.8	20.2	2018
#hurricanenate	3.1×10^{-5}	0.5	10.6	2017
#hurricane nicole	5.3×10^{-6}	0.6	6.9	2016
#hurricaneophelia	1.2×10^{-5}	0.3	6.9	2017
#hurricanesandy	1.5×10^{-4}	1.1	23.0	2012
#hurricanetomas	3.0×10^{-7}	0.9	6.9	2010

TABLE S9. Fitted half-lives τ_1 and τ_2 for all storms with at least 10 days of hashtag usage.

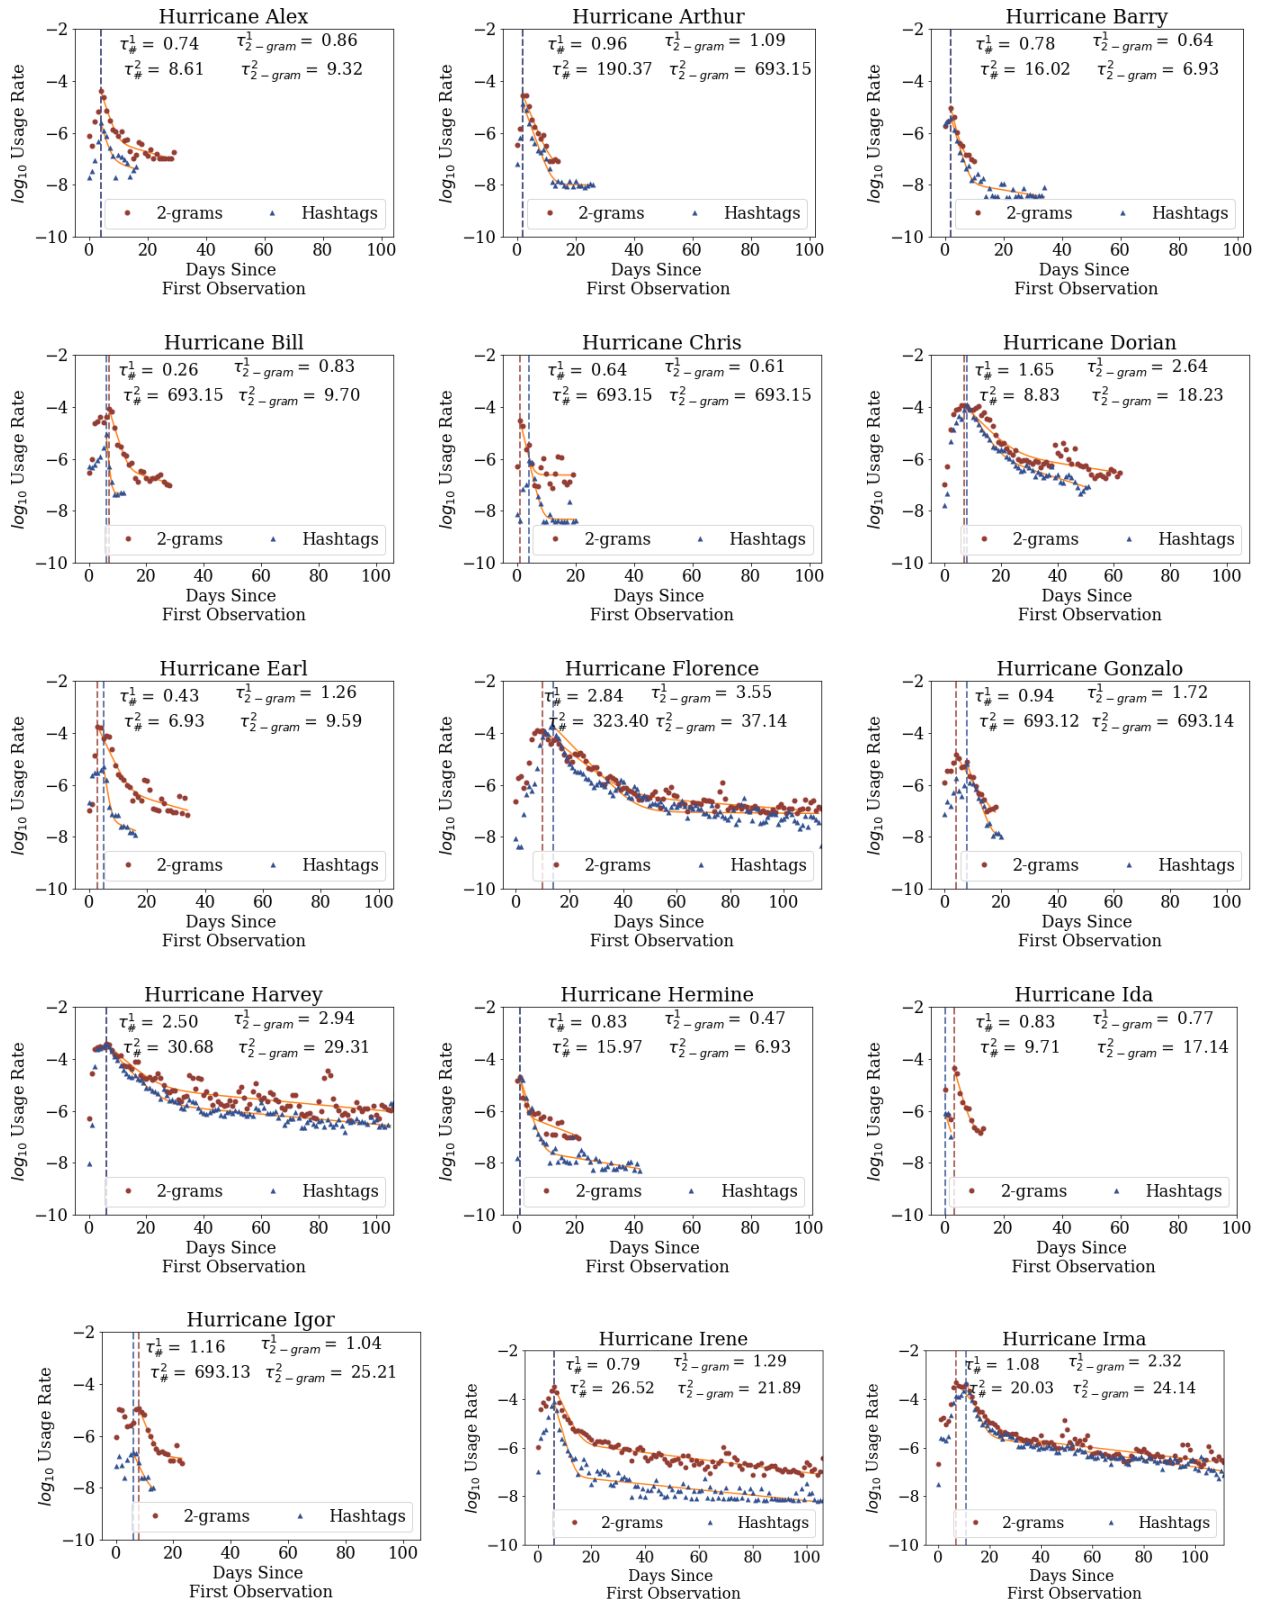


FIG. S6. Hurricane bi-exponential decay fits for hashtag usage rates and 2-gram usage rates for “hurricane *”

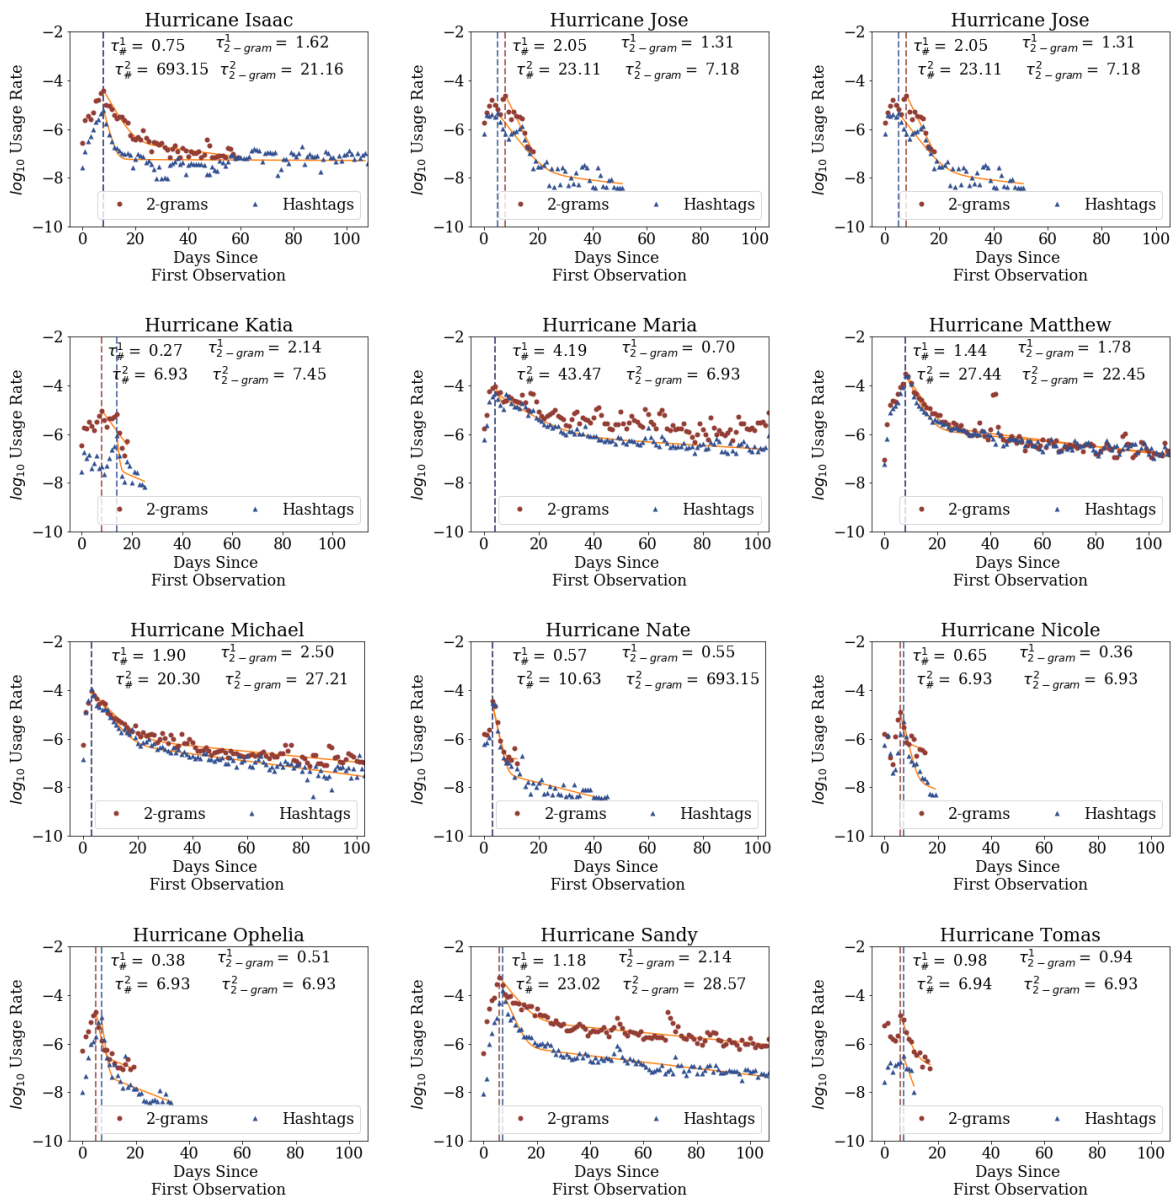


FIG. S7. Hurricane decays fits for all hurricanes for which we have at least 10 days of 2-gram usage rate data. Fits are performed for the function $y = \frac{N}{p+r-q}[(p-q)e^{-(p+r)t} + re^{-qt}]$, a simple two population decay model as proposed by Candia *et al.* [31]. Here p would be interpreted as rate of decay from population 1, r would be the transfer rate from population 1 to population 2, and q would be the rate of decay from population 2. Population 1 might be thought of as bystanders with a shorter attention span, while population two are those living with the ramifications, or working on the recovery who lose attention more slowly. Reported on the graph are the half lives associated with fitting this model for both the hashtag usage rate and 2-gram usage rate, $\tau_1 = \frac{\ln 2}{p+r}$ and $\tau_2 = \frac{\ln 2}{q}$

	Max Usage Rate	τ_1 [Days]	τ_2 [Days]	Season
Hurricane Alex	4.1×10^{-5}	0.8	9.3	2010
Hurricane Arthur	2.8×10^{-5}	1.0	693.1	2014
Hurricane Barry	8.9×10^{-6}	0.6	6.9	2019
Hurricane Bertha	8.2×10^{-6}	0.4	693.1	2014
Hurricane Bill	8.2×10^{-5}	0.8	9.7	2009
Hurricane Chris	3.0×10^{-5}	0.6	693.1	2018
Hurricane Cristobal	1.9×10^{-6}	1.5	693.1	2014
Hurricane Danielle	1.0×10^{-5}	0.9	7.1	2010
Hurricane Danny	7.6×10^{-6}	0.6	693.1	2015
Hurricane Dorian	1.1×10^{-4}	2.6	18.2	2019
Hurricane Earl	1.7×10^{-4}	1.2	9.5	2010
Hurricane Florence	1.3×10^{-4}	3.5	37.1	2018
Hurricane Gert	1.0×10^{-6}	2.1	321.9	2017
Hurricane Gonzalo	1.4×10^{-5}	1.7	693.1	2014
Hurricane Harvey	4.0×10^{-4}	2.9	29.3	2017
Hurricane Hermine	2.0×10^{-5}	0.4	6.9	2016
Hurricane Ida	4.5×10^{-5}	0.7	17.1	2009
Hurricane Igor	1.1×10^{-5}	1.0	25.2	2010
Hurricane Irene	3.3×10^{-4}	1.2	21.8	2011
Hurricane Irma	5.0×10^{-4}	2.3	24.1	2017
Hurricane Isaac	3.8×10^{-5}	1.6	21.1	2012
Hurricane Joaquin	4.4×10^{-5}	1.2	144.5	2015
Hurricane Jose	2.4×10^{-5}	1.3	7.1	2017
Hurricane Karl	1.6×10^{-5}	0.3	6.9	2010
Hurricane Katia	9.3×10^{-6}	2.1	7.4	2011
Hurricane Lorenzo	2.7×10^{-6}	1.7	8.1	2019
Hurricane Maria	1.1×10^{-4}	0.7	6.9	2017
Hurricane Matthew	2.9×10^{-4}	1.7	22.4	2016
Hurricane Michael	9.3×10^{-5}	2.5	27.2	2018
Hurricane Nate	3.5×10^{-5}	0.5	693.1	2017
Hurricane Nicole	1.2×10^{-5}	0.3	6.9	2016
Hurricane Ophelia	1.9×10^{-5}	0.5	6.9	2017
Hurricane Sandy	5.3×10^{-4}	2.1	28.5	2012
Hurricane Tomas	1.4×10^{-5}	0.9	6.9	2010

TABLE S10. Fitted half-lives τ_1 and τ_2 for all storms with at least 10 days of 2-gram usage.

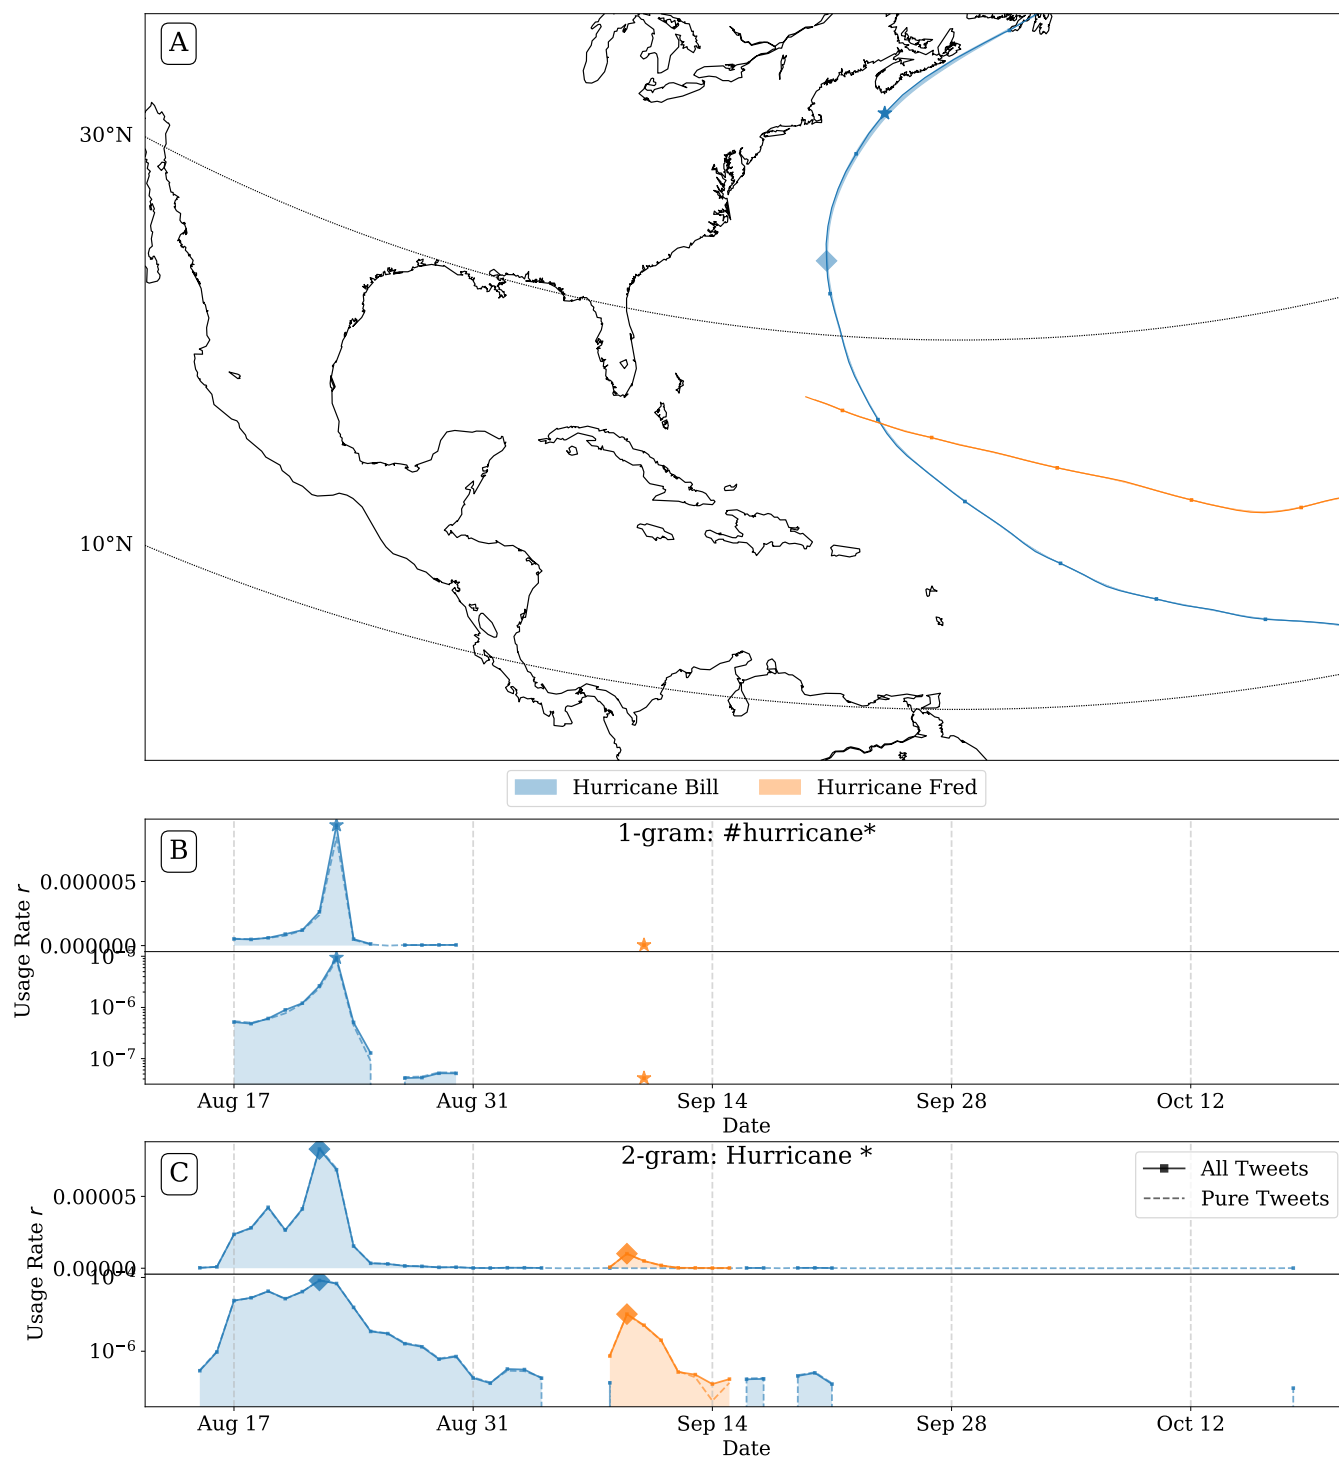


FIG. S8. Hurricane Attention Map and time series for 2009

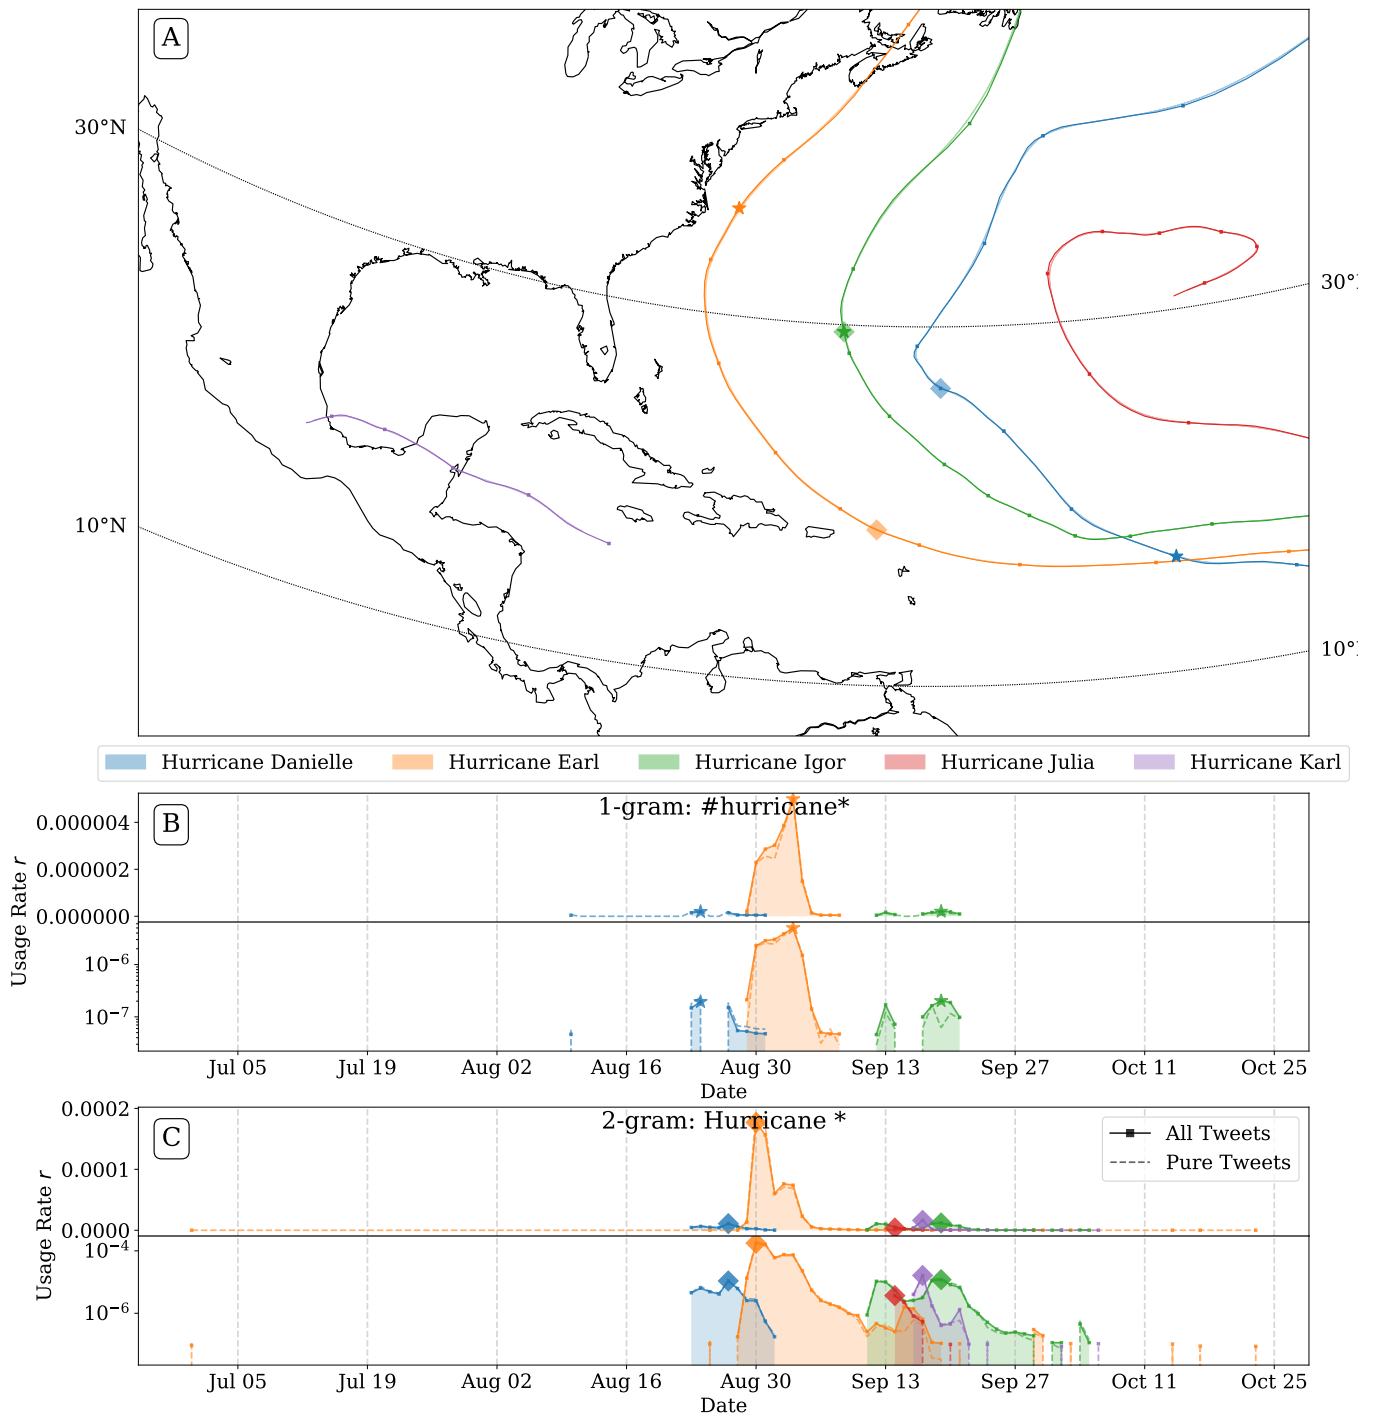


FIG. S9. Hurricane Attention Map and time series for 2010

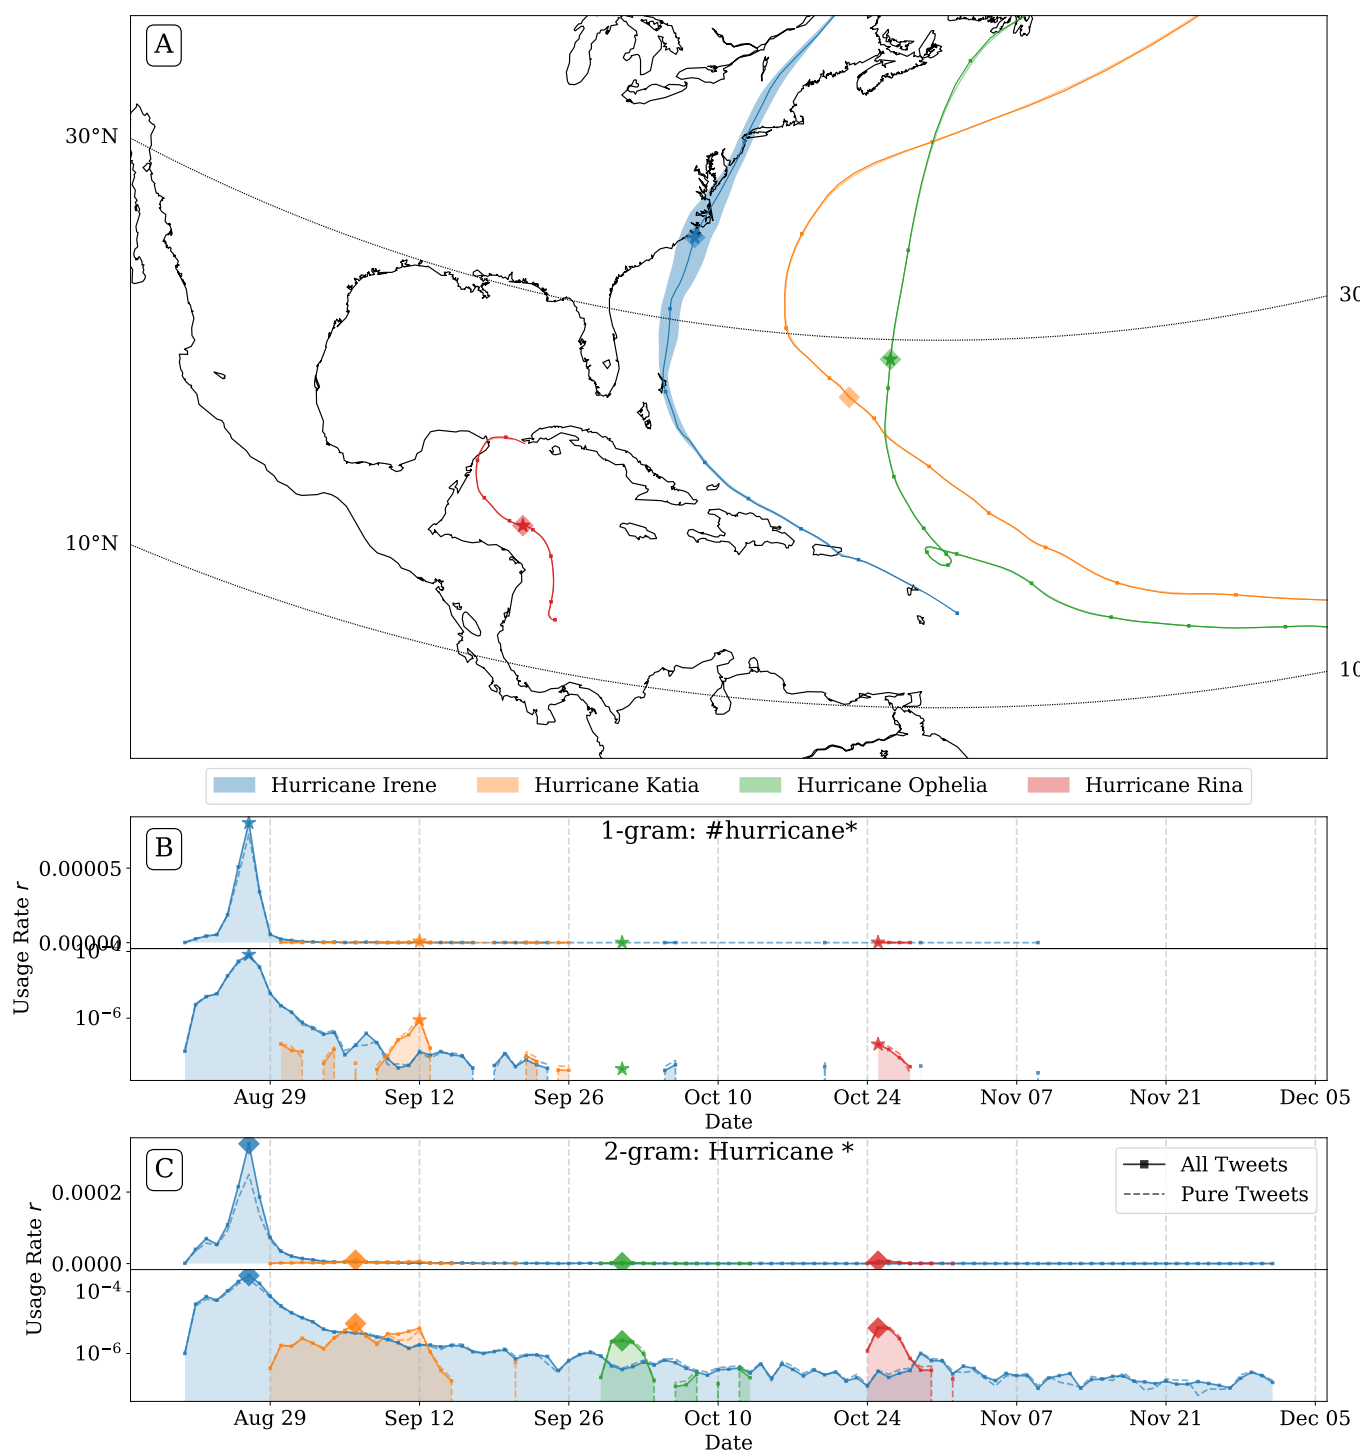


FIG. S10. Hurricane Attention Map and time series for 2011

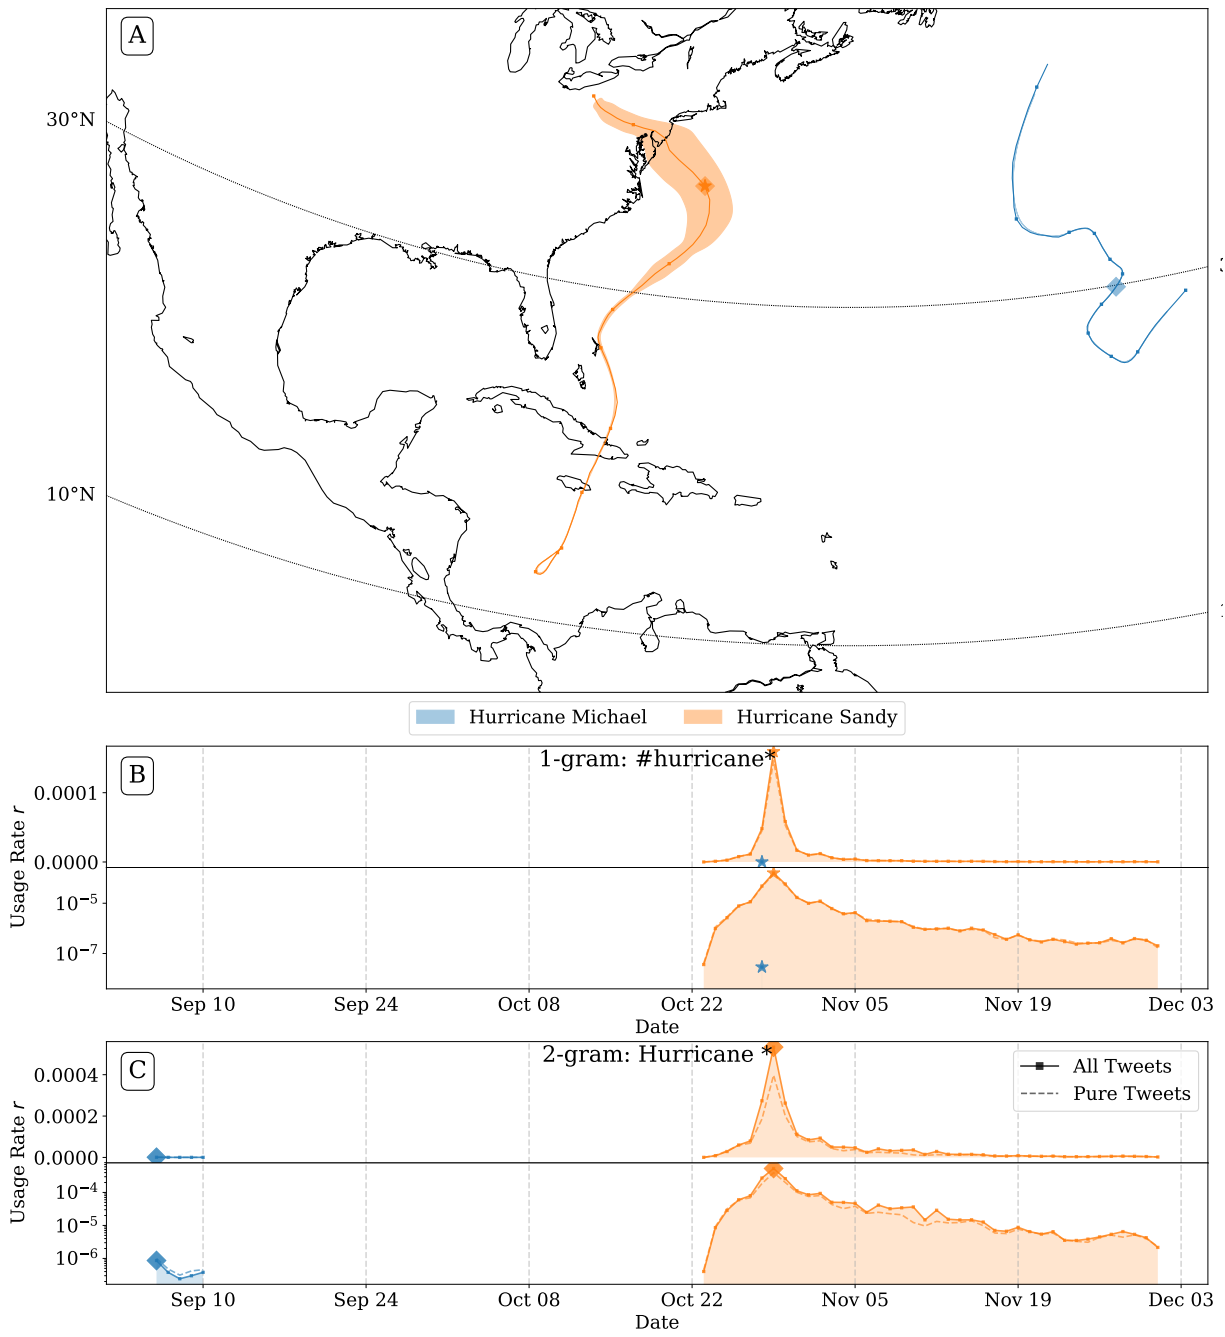


FIG. S11. Hurricane Attention Map and time series for 2012

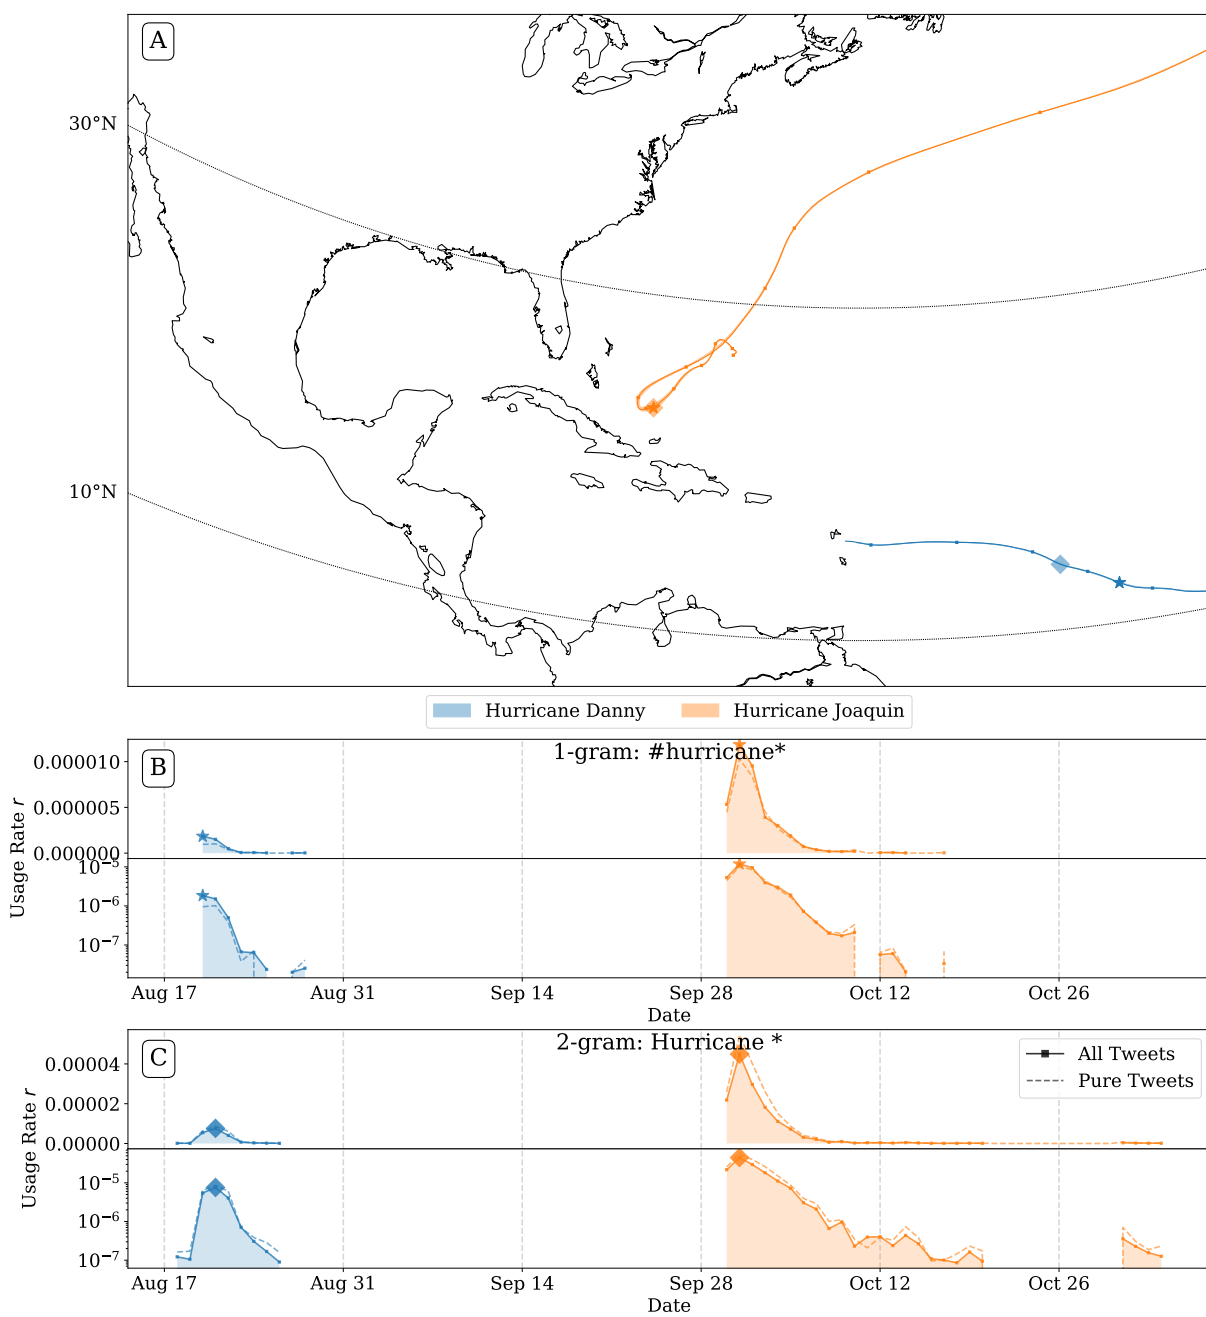


FIG. S12. Hurricane Attention Map and time series for 2015

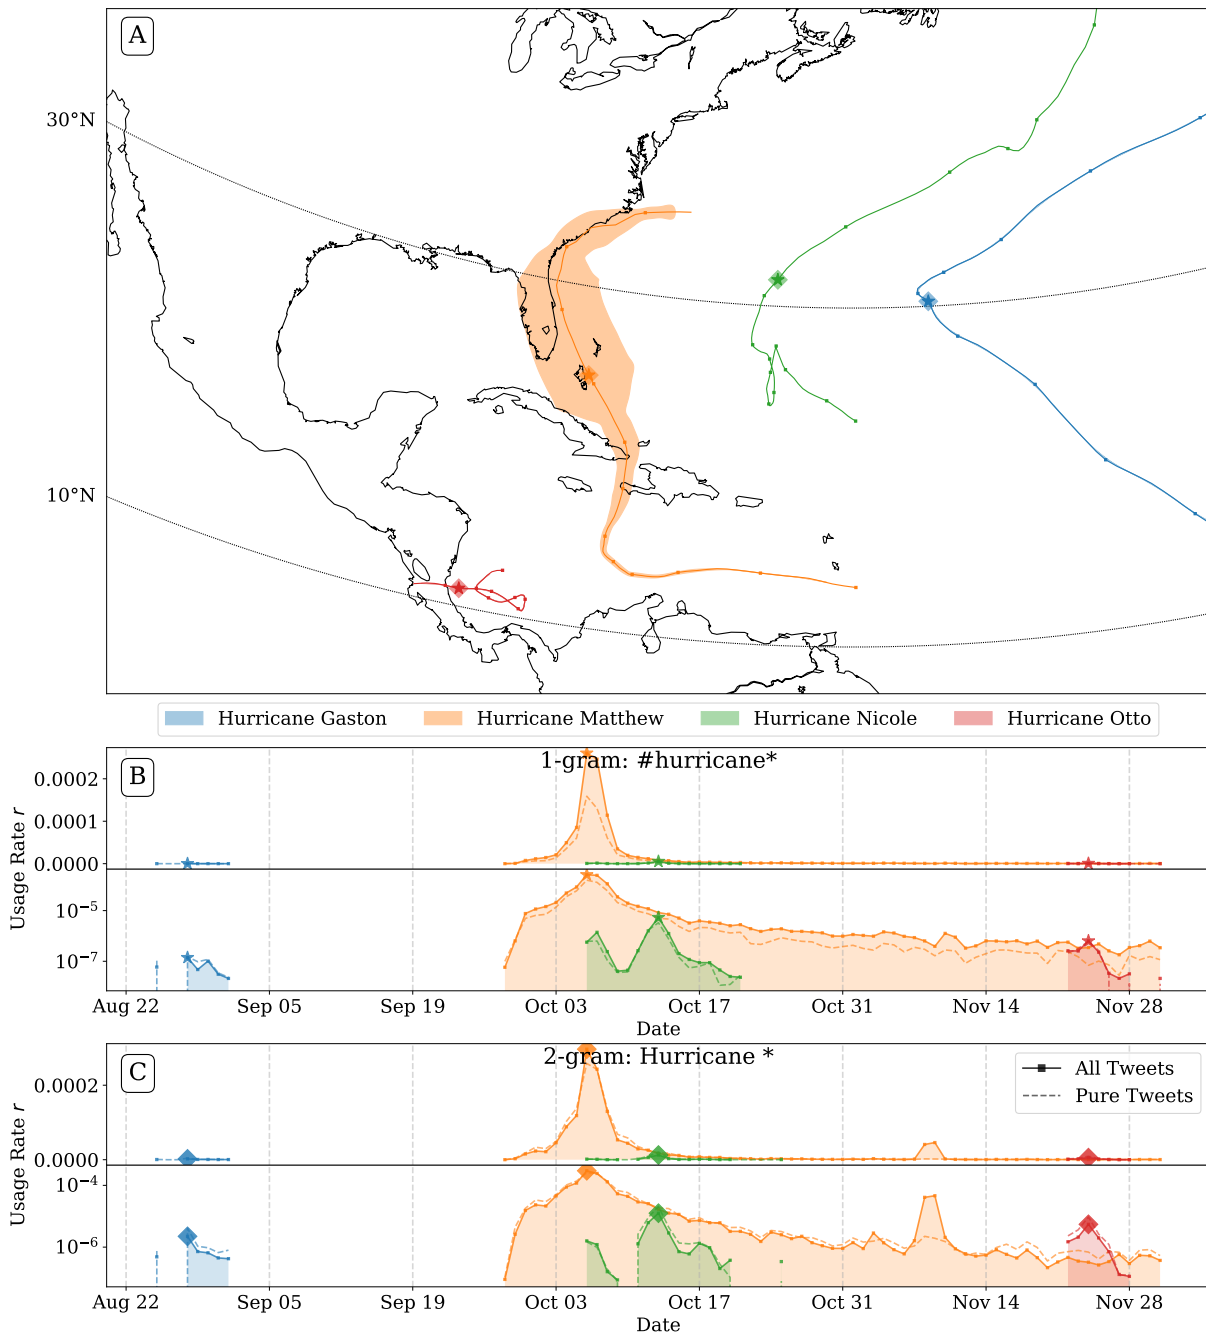


FIG. S13. Hurricane Attention Map and time series for 2016

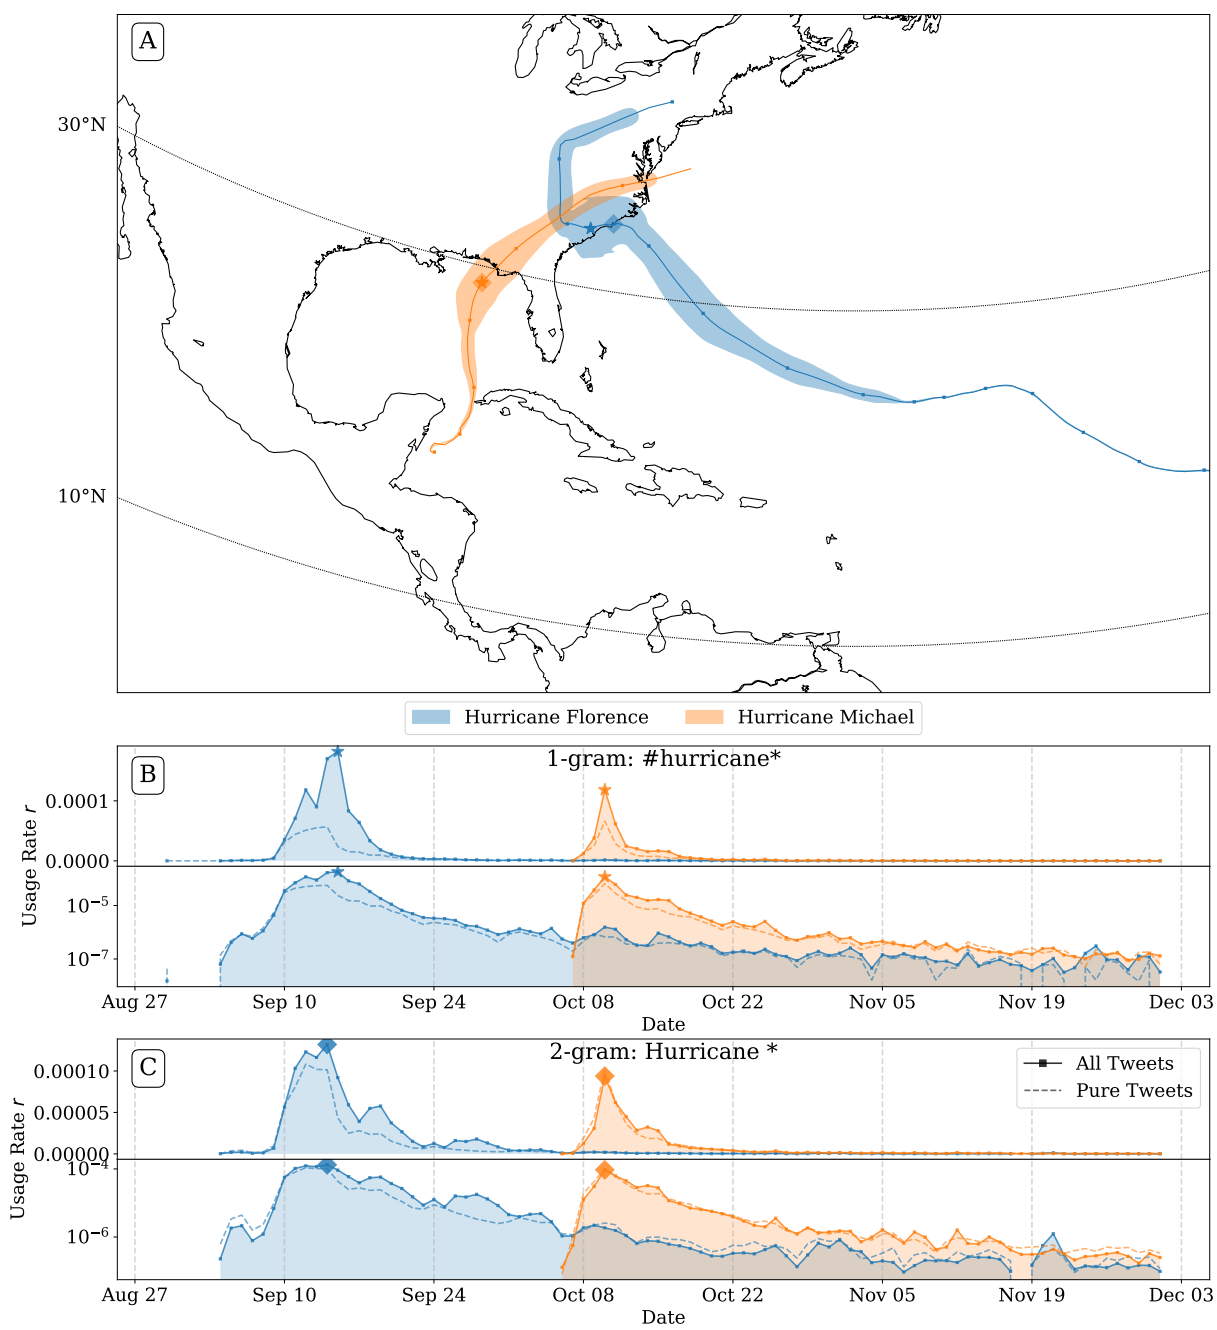


FIG. S14. Hurricane Attention Map and time series Map and time series for 2018