

# #Bigbirds Never Die: Understanding Social Dynamics of Emergent Hashtags

Yu-Ru Lin<sup>1</sup> Drew Margolin<sup>2</sup> Brian Keegan<sup>1</sup> Andrea Baronchelli<sup>3</sup> David Lazer<sup>1,2</sup>

<sup>1</sup>College of Social Sciences and Humanities, Northeastern University, Boston, MA 02115, USA

<sup>2</sup>College of Computer and Information Science, Northeastern University, Boston, MA 02115, USA

<sup>3</sup>Lab for the Modeling of Biological and Socio-technical Systems, Northeastern University, Boston, MA 02115, USA

## Abstract

We examine the growth, survival, and context of 256 novel hashtags during the 2012 U.S. presidential debates. Our analysis reveals the trajectories of hashtag use fall into two distinct classes: “winners” that emerge more quickly and are sustained for longer periods of time than other “also-rans” hashtags. We propose a “conversational vibrancy” framework to capture dynamics of hashtags based on their topicality, interactivity, diversity, and prominence. Statistical analyses of the growth and persistence of hashtags reveal novel relationships between features of this framework and the relative success of hashtags. Specifically, retweets always contribute to faster hashtag adoption, replies extend the life of “winners” while having no effect on “also-rans.” This is the first study on the lifecycle of hashtag adoption and use in response to purely exogenous shocks. We draw on theories of uses and gratification, organizational ecology, and language evolution to discuss these findings and their implications for understanding social influence and collective action in social media more generally.

## Introduction

The hashtag is a ubiquitous and flexible annotation for Twitter users. Hashtags allow users to track ongoing conversations, signal membership in a community, or communicate non-verbal cues like irony. Hashtags vary considerably in their activity and audience, but little is known about how the use of novel hashtags co-evolve with the needs of their users and the presence of other hashtags. Why do some hashtags persist while others are just momentary blips? How do similar hashtags compete for attention?

Hashtags often reflect eccentric topics and their emergence is happenstance. These idiosyncrasies typically limit researchers’ ability to systematically compare features of their emergence and evolution. However, the shared attention to and improvisational nature of the U.S. presidential debates provides a unique opportunity to understand the social dynamics of novel hashtag adoption following repeated exogenous events. Before the October 2012 debates, the importance of concepts such as “big bird”, “malarky”, “binders”, or “bayonets” to a political campaign would be minimal. Yet these unexpected terms now denote salient

moments from these debates and became rallying calls for partisans and punchlines for comedians. More broadly, Twitter users created hundreds of hashtags related to these topics which had no *raison d’être* before the debates. The candidates’ unscripted statements can thus be understood as creating an exogenous shock to the system of political discourse. Combined with fine-grained data about large scale and real time user behavior create, these exogenous shocks create a set of natural experiments to systematically analyze the features that contribute to the growth and stabilization of hashtags.

Previous research has examined the roles of “relevance” and “exposure” for hashtag adoption (Huang, Thornton, and Efthimiadis 2010; Yang et al. 2012; Lehmann et al. 2012; Yang and Counts 2010; Romero, Tan, and Ugander 2011; Romero, Meeder, and Kleinberg 2011). However, these features are inadequate for distinguishing hashtag trajectories in the context of an event as prominent as the U.S. presidential debates where relevance and exposure are effectively “maxed out.” Furthermore, hashtags are more than labels for contextualizing statements, objects for bookmarking, or channels for sharing information, but they are active virtual sites for constructing communities. We propose an alternative framework called *conversational vibrancy* to understand the dynamics of novel hashtag use based on how features such as topicality, interactivity, diversity, and prominence interact with the communities producing and consuming tweets to these hashtags.

We operationalize our conversational vibrancy framework with behavioral and structural variables such as retweets, replies, unique tweets, and audience size. Analysis of the relationship between these features and the rate and sustainability of hashtag adoption reveal the existence of two distinct classes of hashtag trajectories. “Winners” are hashtags characterized by more rapid growth and longer sustained interest than “also-rans.” Statistical models of the growth and persistence of each class of hashtags reveals novel and complex interactions with features from our conversational vibrancy framework as well as the concentration of user attention. Our analysis suggests conversational vibrancy contributes to the growth and sustenance of hashtags in different phases of their life-cycles. For example, while retweets and followers support a hashtag’s growth, they also paradoxically undermine its persistence. We draw on theories of or-

ganizational ecology and language evolution to discuss these findings and their implications for understanding social influence and collective action in social media more generally.

## Background

Twitter is a popular “micro-blogging” platform founded in 2006 that allows users to receive and broadcast (up to) 140-character messages called “tweets”. The brevity and immediacy of the medium contributed to its rapid adoption as a way of sharing information from mobile devices, propagating information by “retweeting”, as well as tracking relevant content by labeling tweets with “hashtags” (Boyd, Golder, and Lotan 2010).

Hashtags first emerged on Twitter during the 2007 San Diego wildfires as a way to track relevant information about a large scale natural disaster by labeling content so that it could be filtered and shared via backchannels (Sutton, Palen, and Shklovski 2008). Unlike traditional tagging systems used for information archival, Twitter hashtags can serve either as a label for identifying topically relevant streams of message or a prompt for commenting and sharing (Huang, Thornton, and Efthimiadis 2010). Hashtags often fill a dual role as both a topical identifier (e.g., #iPhone) and a symbol of a community membership (e.g., #VoteForObama) (Yang et al. 2012). During the “Arab Spring” and other protests, hashtags were used by activists to coordinate their actions and garner support (Starbird and Palen 2012).

Studies of hashtag use have primarily emphasized two dimensions: relevance and exposure. The relevance perspective emphasizes the dependence hashtags have on the topics being discussed: content relevant to the Twitter community draw more tweets and retweets and thus grow sustainably over time (Huang, Thornton, and Efthimiadis 2010; Yang et al. 2012; Lehmann et al. 2012; Yang and Counts 2010). From the view of exposure, however, hashtag’s publicity and legitimacy is driven by social exposure as part of a process of complex contagion: hashtags grow because they are adopted by users with social connections to one another, and by users’ repeated use and spreading the word to others of interests (Romero, Tan, and Ugander 2011; Romero, Meeder, and Kleinberg 2011).

Analyses of hashtag adoption have characterized differences in the persistence of hashtags, modeled the spread of hashtags as a type of complex contagion requiring multiple exposures, and the behavior of hashtags that emerge in response to exogenous shocks (Romero, Meeder, and Kleinberg 2011; Romero, Tan, and Ugander 2011). Other studies of hashtag have focused on the temporal occurrences of the hashtags that are characterized by either “peaky” but ephemeral topics versus “persistent conversations” using less salient terms over longer periods of time (Shamma, Kennedy, and Churchill 2011). In (Lehmann et al. 2012), peaks in popular hashtags are the result of four discrete classes of activity and are primarily driven by exogenous rather than endogenous factors. While these analyses have analyzed the use of hashtags in the context of specific events, topics, or communities, they have not examined the ecology of multiple hashtags competing for attention following an exogenous event.

## Conversational vibrancy

While prior work has emphasized relevance and exposure as mediators of hashtag activity, these features are inadequate for explaining activity during events such as elections, sports, and awards shows that attract unusual levels of shared attention and intense activity. The large audiences who view these popular events predispose them to generating hashtags that become widely adopted as users “dual screen” by tweeting about what they watch (Shamma, Kennedy, and Churchill 2009; 2011). During these events, exogenous shocks can simultaneously expose users to a large number of candidate hashtags that are of similar relevance. These constructs are thus inadequate to distinguish between hashtags that emerge under these conditions. A more interesting question is why are some hashtags related to these events adopted while others are not?

Uses and gratification (U&G) is classic theoretical paradigm from mass communication research that examines why people become involved in some types of mediated communication versus others and what gratifications they receive from it. U&G emphasizes “what people do with media rather than what media do to people” which makes it especially salient for studying social media like Twitter where users both produce and consume content. U&G argues that different forms of media compete to provide distinct uses such as news or entertainment and audiences actively seek out those media that are meaningful to them and can best satisfy their individual needs for information, social connection, or emotional support (Ruggiero 2000; Stafford, Stafford, and Schkade 2004).

We draw from U&G theory to argue that users choose to employ some hashtags while ignoring others because patterns of use around these hashtags provide a better match for their needs. However, we also recognize that patterns of communication (Boyd, Golder, and Lotan 2010) and perceptions of users’ audiences (Marwick and boyd 2011) interact in complex ways with other contextual and content-level features (Yang and Counts 2010; Suh et al. 2010) to influence the spread and adoption of content in social media (Lerman and Ghosh 2010). This suggests the need for an integrative framework that synthesizes these multilevel processes to understand why users adopt some hashtags while rejecting others. We propose *conversational vibrancy* as an explanatory framework to understand the emergence, growth, and persistence of a hashtag and its community of users. Conversational vibrancy consists of four general elements, each of which we operationalize with behavioral features specific to Twitter:

- *Topicality* is a kind of relevance that reflects the extent to which tweets labeled by a hashtag are timely and have contextual relevance for on-going conversations. “#big-bird” was highly topical following the first debate but was not topical during the inauguration because its humor and relevance are lost in this different context. The number of times a hashtagged tweet created by one user is retweeted reflects the value this information has for other users. Hashtags are more likely to grow and persist when users re-share hashtagged content created by others with their own networks.

- *Interactivity* is the extent to which individuals attend to others' statements by responding to them in turn. Hashtags can denote spaces for on-going interpersonal conversation, supportive interest-based communities, or heated debates that bring others into the community and get them to adopt the hashtags as well. A hashtag may encourage people to launch discussions about the topic or people may simply reply to a user with a hashtag to acknowledge his or her wit. The number of replies co-occurring with a hashtag captures the extent to which users are not only paying attention to others' content but actively responding to one another. Hashtags are more likely to grow and persist when there are many users speaking directly to each other while invoking the hashtag.
- *Diversity* captures whether activity is concentrated in the same information or people are attending to many different sources of information. Hashtags that emerge because of a single "one hit wonder" tweet being retweeted over and over are qualitatively different from many unique tweets using the same hashtag. Because counting the absolute frequency of tweets observed using a hashtag can be biased towards hashtags being hustled by spammers, the number of unique tweets that have been retweeted at least once provides a reasonable threshold for capturing the breadth of tweets attracting attention. Hashtags are more likely to grow and persist when many unique tweets with the hashtag are being shared.
- *Prominence* measures the audience exposed to a hashtag by a user mentioning it. The audience of Twitter users plays a crucial and complex role in information sharing behaviors (Marwick and boyd 2011). Features of users themselves sharing information can have greater influence than the content itself being shared (Yang and Counts 2010). Hashtags mentioned by users with large audiences of many followers are more likely to result in rapid growth and sustained activity than hashtags mentioned by users with small audiences.

In summary, hashtags exhibiting higher levels of conversational vibrancy should gratify more users' needs and thus contribute to the success of the hashtag. The user-generated and socially-mediated nature of Twitter will create a positive feedback loop through which hashtags with high conversational vibrancy will preferentially recruit more attention from users who in turn increase the vibrancy of the hashtag while also decreasing the communicative vibrancy of competing hashtags. We discuss these evolutionary dynamics and ecological interactions in more detail in the discussion.

## Our Approach

Contemporary presidential campaigns in the United States are highly scripted and candidates offer repetitive "stump speeches" that rarely provide new information about their position or strategies (Skewes 2007). However, debates are highly ceremonial and enthralling events that not only attract large audiences of pundits, partisans, and undecided voters, but also require candidates to respond directly to critiques and improvise responses outside of well-worn campaign talking points (Schroeder 2008). As a result, these debates

are ripe for unexpected and embarrassing statements. These "gaffes" can be impromptu exchanges that make a candidate look foolish or outright factual errors that undermine a candidate's credibility and often later become fodder for opponents, sometimes substantially changing the dynamics of a race (Clayman 1995). Examples of "gaffes" from the first presidential debate include Mitt Romney's unprompted threat to cut funding for the popular PBS children's television show "Sesame Street" with its iconic lead character Big Bird as well as Barack Obama's widely-derided aloof performance.

The combination of fine-grained behavioral data about a large population of people and high-stakes circumstances of Presidential debates provide a unique opportunity to analyze how a large and engaged audience reacts to statements to which it had not previously been exposed. Because substantial portions of the audience are monitoring or participating in social media while simultaneously watching the live debate (PEW 2012), these gaffes immediately prompt behavior on social media as users try to confirm a statement or alternatively improvise humorous responses to mock it. Unlike prior studies of hashtag adoption, threats to validity from confounded endogenous processes are minimized since gaffes and other unexpected statements made by candidates are inherently unpredictable and exogenous from our analysis.

Examples of tweets containing the hashtags from the first debate are provided in Table 1. The table shows tweets that were posted immediately after or close to the onset time of different hashtags, focusing on the most popular, politically relevant hashtags or their variants. In the first debate, after Mitt Romney mentioned "Big Bird," there were several similar hashtags, including #bigbird2012, #bigbird and #savebigbird that appeared simultaneously in users' tweets (around 21:28 EDT). The #supportbigbird hashtag was popularized 15 minutes after it, largely due to intensive retweeting of a newly-created account under the handle of "@BIG-BIRD".

We examine the dynamics of hashtag adoption during and immediately following each of the four debates during the 2012 presidential campaign. To understand why some novel hashtags are adopted while others are not, we measure two specific outcomes: the initial *growth* of the hashtag as users adopt it and the *persistence* of hashtag use after this initial growth phase. Because the behavior we model in this project is highly domain-specific to politics, we identify and preferentially sample from a population of politically active users to extract their tweeting history. For this sample of users, we examine the emergence of hashtags which are "novel" because they had not appeared in users' feeds beforehand and identify 256 of the most popular and domain-relevant hashtags. Using a variety of inferential statistical modeling approaches, we examine the relationship between the features of the conversational vibrancy framework and the growth and persistence of these hashtags.

## Dataset

While randomly sampling content from Twitter's "garden hose" and aggregating it to understand collective behavior is

Table 1: Examples of debate-related tweets from the first debate on 4 October 2012

hashtag	time (EDT)	tweet text
bigbird2012	21:28:12	#bigbird2012
bigbird	21:28:12	Romney stopping subsidy to PBS telling the host that?! woou #BigBird
savebigbird	21:28:15	#savebigbird #debate
bigbird	21:28:18	How you cut #BigBird though #Romney....Wtf??? Smh
bigbird	21:28:20	@MittRomney says he will cut funding to PBS even thou he likes Lehrer and #bigbird #debates
savebigbird	21:28:50	Cut PBS? Noooooooooooooi #SaveBigBird
supportbigbird	21:46:37	RT @BIGBIRD: Yo Mitt Romney, Sesame Street is brought to you today by the letters F U! #debates #SupportBigBird
supportbigbird	21:51:54	This entire election is now about who will save Big Bird. #supportbigbird #debates
supportbigbird	21:52:45	#OccupySesameStreet #SupportBigBird we are the 47%

Table 2: Summary of datasets. All times are Eastern Daylight Time (EDT).

Debate number	1	2 (Vice Presidential)	3	4
Debate starting time	3 Oct. 21:00	11 Oct. 21:00	16 Oct. 21:00	22 Oct. 21:00
Tweet volume at peak	3,268,918	2,388,963	3,608,291	2,415,703
Unique users at peak	174,297	155,739	181,329	152,538
“Novel” hashtags	92,432	58,165	91,705	77,526
“Pop” hashtags	75	57	82	42
Tracking conclusion time	7 Oct. 02:00	15 Oct. 02:00	20 Oct. 02:00	26 Oct. 02:00

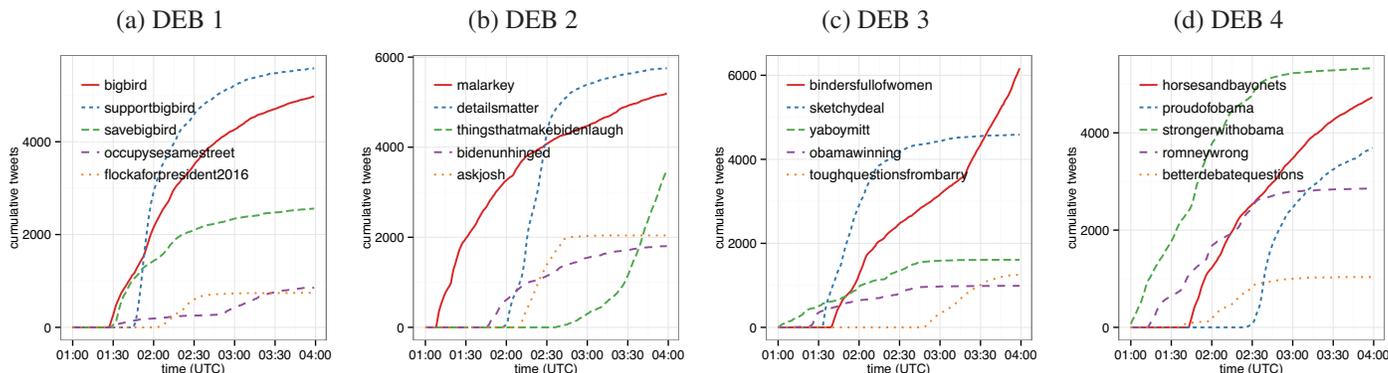


Figure 1: Cumulative tweet volume of top hashtags over time, starting from each debate.

methodologically defensible under some research designs, we argue it is inappropriate for our purposes. We identify a relevant sub-population of users as follows. Users with relevant characteristics are identified first, and tweets for these users are extracted for analysis. This focus on the tweets from specified users reduces the threat of selection bias where inclusion in the observed set is correlated with the dependent variable (Lin et al. 2013). The dependent variables in this study refer to changes in tweet volumes over time. We thus focus on a fixed set of users—those that tweeted heavily during the debates—across the entire observation period. Changes in tweets to the hashtag can thus be attributed to choices made by these users.

**User sampling.** First, we identified politically-active users who tweeted using a hashtag such as “#debate” or mentioned either candidate’s Twitter account<sup>1</sup> during any of the four presidential debates. Using Twitter’s “garden hose” streaming API<sup>2</sup>, users whose tweets appeared in the feed were selected into our population. Second, we extracted

<sup>1</sup>@barackobama, @mittromney

<sup>2</sup><https://dev.twitter.com/docs/streaming-apis>

these users’ tweeting histories beginning in mid-August 2012 through late October 2012 using Twitter’s REST API<sup>3</sup>. Because these queries are expensive owing to rate limits, we prioritized users who tweeted during more of the debates. Thus users who tweeted during all four debates are more likely to be represented in the sample than users who tweeted during only one of the debates. While biased, this population captures users such as journalists, pundits, politicians, and activists who are more politically active and thus relevant to our research setting. The subset of the resulting corpus used in this analysis contains 123,560,785 tweets from 2,516,125 unique users posted between September 29 and October 27.

Descriptive statistics for each of the four debates are summarized in Table 2. Each debate occurred between 21:00 and 22:30 EDT and we tracked the behavior of our population of users for 77 hours following each debate. The tweet volume peaked at the first hour of each debate and remained high at the second hour. We refer the two-hour window as “peak window” and identify “novel” hashtags that were born in

<sup>3</sup><https://dev.twitter.com/docs/api/1.1>

each of the peak windows. Figure 1 shows the growth of top five hashtags in terms of the cumulative number of tweets for each debate.

**Pre-processing.** We extract “novel” hashtags as all hashtags that were contained in tweets posted during one of the peak windows but did not appear in any tweets posted within 96 hours prior to the peak windows. For each novel hashtag, we extract the users who have mentioned the hashtag within the 77 hour post-debate window starting from the peak window. We then select hashtags that have been mentioned by at least 100 unique users to be “pop” hashtags. Despite the features of our politically active user population, many of these hashtags were not related to the debate or politics generally and were removed unless the content of the tweet contained terms such as “debate,” “president,” or any of the four candidates’ names. The resulting corpus contained 256 “pop” hashtags that were novel, widely used, and relevant to the debates. The numbers of pop hashtags from each debate are reported in Table 2.

The conversational vibrancy framework outlined above is operationalized using the following variables.

1. *Number of retweets (rt)* under the hashtag is an indicator of its **topicality**. A retweet is an example of an individual identifying a tweet that is so interesting or relevant that they wish to re-state it, verbatim, to their own followers. Hashtags attracting more retweets have higher topicality.
2. *Number of replies (rp)* to a hashtag is an indicator of its **interactivity**. A heated debate may be more interactive than a collection of statements that agree but do not address one another.
3. *Number of unique retweet sources (src)* measures the **diversity** of the conversation. A hashtag with more unique retweet sources provides more diverse information.
4. *Expected follower size (follow)* is proxy for the **prominence** of the users who tweet to the hashtag. Users with many followers tweet to a hashtag will increase the growth and persistence of the hashtag.

### Characterizing Hashtag Growth and Persistence

Figure 1(a) reveals an interesting hashtag dynamics during the first presidential debate: while the hashtag “#bigbird” was born and rapidly grew at around 21:30 (1:30 UTC), the hashtag “#supportbigbird” which was born 15 minutes after took over in about 10 minutes. To see if this is simply the end of the story, we inspect their growth for a longer period. As shown in Figure 2(a), after six hours from the start of the debate, “#supportbigbird” was still on top of “#bigbird.” However, in a larger time scale as in Figure 2(b), we discover that “bigbird” won back in the 12<sup>th</sup> hour after the debate. The differences between the short-term and long-term dynamics of hashtags lead to an interesting question: How can we characterize the complex dynamics of these emergent hashtags?

From the temporal curves it is clear that the dynamics can not be easily captured by a single process, and hence the parametric modeling approach suggested by prior work (e.g., (Crane and Sornette 2008)) is not appropriate. We focus on analyzing two specific features of the “pop” hashtags’

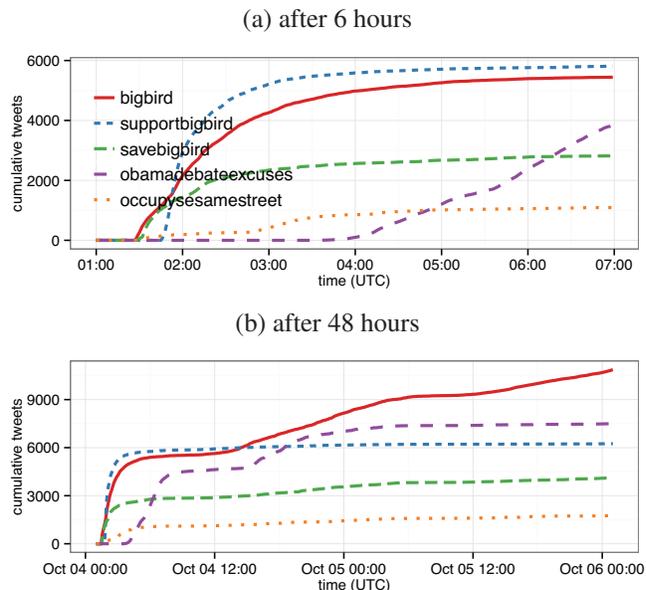


Figure 2: Cumulative tweet volume of top hashtags in the first debate, over the first 6 and 48 hours. Four out of the five hashtags appeared due to Romney’s statement regarding PBS, and one was related to the debate more generally.

adoption dynamics. *Growth* measures the rate of change in hashtag use over observation window. A hashtag that is mentioned 1000 times in 1 hour has higher growth than a hashtag that is mentioned 1000 times in 10 hours. *Persistence* measures the sustained activity of hashtag use in new tweets over time. A hashtag that continues to be mentioned 70 hours after the debate has greater persistence than a hashtag that stopped being mentioned after 12 hours.

Instead of fitting the curves by a parametric model, we capture the shape of curves by fitting a spline function. Figure 3 show three different patterns for the emergent hashtags. A hashtag may grow extremely fast and saturate quickly, as shown in Figure 3(a), grow slower but also slowly saturate, as shown in Figure 3(b), or grow fast and sustain for a longer period, as in Figure 3(c). We then quantify the growth of a hashtag as the largest slope over its fitted spline function. The slope is measured in number of tweets per minute (tpm), so a hashtag with a growth 60 indicates the hashtag gathered 60 tweets per minute in its fast growth phase.

We identify three critical time points along a hashtag’s growth curves:

- *onset time ( $t_0$ )*: the time where the hashtag was first mentioned in our dataset.
- *saturated time ( $t_e$ )*: the time where the hashtag’s size (in terms of total number of tweets) reach 99% of its final size in our dataset. We use 99% instead of final size to avoid the influence from miscellaneous outliers.
- *turning point ( $t^*$ )*: the time point where the hashtag growth curve starts deviating from its tangent line of the largest slope. This turning point can be found by a line search procedure along the tangent line.

The persistence is measured as the duration between its onset time  $t_0$  and the saturated time  $t_e$ . In addition, the turning

point divides a hashtag’s life-cycle into an initial fast growth phase followed by a relatively slow growing phase before reaching its saturated point.

## Results

We use this population of 2.5 million politically-active Twitter users and 256 “pop” hashtags to first describe general patterns of hashtag growth and persistence. Based on behavioral features from the conversational vibrancy framework, we use cluster analysis methods to identify two distinct “winner” and “also-ran” hashtag classes. Integrating these approaches, we develop statistical models to inductively analyze how features of conversational vibrancy covary for both hashtag classes in the growth and persistence phases of their adoption.

### Categorizing popular hashtags

Figure 4 shows the 256 hashtags along the three dimensions: growth, persistence, and final size (the total number of tweets). We use k-means clustering to identify the two distinct classes of hashtags: “winners” (in red) grow rapidly and have high levels of persistence while “also-rans” (in blue) hashtags have either slower growth or less persistence. The “winner” class correspond to the cluster with relatively large final size. The descriptive statistics of the 12 winner hashtags are summarized in Table 3.

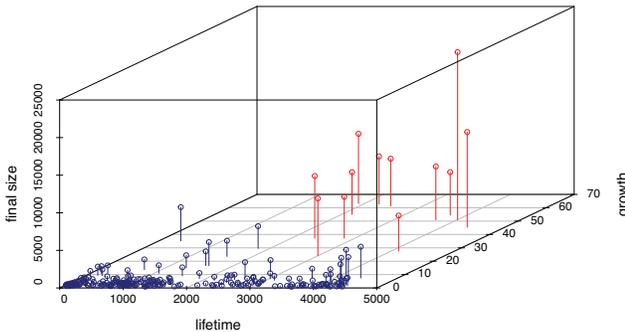


Figure 4: Two classes of popular hashtags. “Winners” (in red) have high growth and persistence (lifetime) with large final size (i.e., total number of tweets). “Also-rans” (in blue) are the rest of “pop” hashtags.

### Explaining Hashtag Growth

We use time-series regression to discover the relationships between the hashtag growth and different aspects of conversational vibrancy. Our goal is to model the minute-by-minute hashtag growth from the onset up to the turning point. A time-series regression allows us to fit a model of dependent variable (growth) on independent variables (conversational vibrancy) where the serially correlated errors are captured by a linear autoregressive moving-average specification. Concretely, we consider the dependent variable  $y_t$  as the number of new tweets for a hashtag at time  $t$  where  $t$  is between the hashtag onset time  $t_0$  and its turning point  $t^*$ .

For each hashtag, we include the following predictors that reflect its associated conversational vibrancy and a control:

- $\mathbf{rt}_{t-1}$ : the number of new retweets at previous time  $t - 1$ .

Table 3: Hashtags in the “winner” class

DEB	Hashtag	total tweets	growth (tpm)	persistence (mins)
1	bigbird	12667	45.45	4405
1	obamadebateexcuses	7617	24.12	2999
1	supportbigbird	6289	61.35	2493
1	savebigbird	4721	27.60	4118
2	thingsthatmakebidenlaugh	9290	63.10	1907
2	malarkey	7108	50.90	3669
2	detailsmatter	6358	62.68	2249
3	bindersfullofwomen	22287	51.10	4003
3	sketchydeal	5704	54.55	3733
4	horsesandbayonets	8266	37.27	2365
4	strongerwithobama	5610	55.07	2162
4	proudofoabama	5502	37.28	2830

- $\mathbf{rp}_{t-1}$ : the number of replies at time  $t - 1$ .
- $\mathbf{src}_{t-1}^\alpha$ : the number of new retweet sources up to time  $t - 1$ .
- $\mathbf{follow}_{t-1}^\alpha$ : the expected largest audience size up to time  $t - 1$ . To compute this, we identify the set of users  $U$  who mentioned the hashtag up to time  $t - 1$ , and then compute the expected follower size from the users whose follower size is on the top  $10^{th}$  percentile in set  $U$ .

We use the notation  $(\cdot)^\alpha$  to indicate whether the predictor is an aggregate measure aggregating data from time  $t_0$  to time  $t - 1$ . To fit the linear model, the variables  $y_t$ ,  $\mathbf{rt}_{t-1}$ ,  $\mathbf{rp}_{t-1}$  and  $\mathbf{follow}_{t-1}^\alpha$  are log-transformed. We use a Box-Jenkins autoregressive integrative moving average (ARIMA) modeling framework (Durbin and Koopman 2001) to evaluate the autocorrelation function and partial autocorrelation function of the residuals, and determine a model with second-order autoregressive process and first-order moving average. The model can be specified as:

$$y_t = \beta^T \mathbf{x}_{t-1} + \epsilon_t$$

$$\epsilon_t = \phi_1 \epsilon_{t-1} + \phi_2 \epsilon_{t-2} + \nu + \psi \nu_{t-1}$$

where  $\mathbf{x}_{t-1}$  is a vector containing the time-dependent predictors,  $\beta$  is a vector of parameters to estimate,  $\epsilon$  is the error,  $\phi_1$  and  $\phi_2$  are the first- and second-order autocorrelation parameters,  $\psi$  is the first-order moving-average parameter. The estimated parameters are summarized in Table 4. The table shows parameter estimates (standard errors in parentheses) for the regression models.

In the “also-ran” class, hashtag growth has significant and positive associations with  $\mathbf{rt}_{t-1}$ ,  $\mathbf{rp}_{t-1}$  and  $\mathbf{follow}_{t-1}^\alpha$ . In the “winner” class, the growth is positively associated with  $\mathbf{rt}_{t-1}$  and  $\mathbf{follow}_{t-1}^\alpha$ , but has a weak and negative associations with  $\mathbf{src}_{t-1}^\alpha$ . In both classes, hashtags whose tweeters have more followers ( $\mathbf{follow}_{t-1}^\alpha$ ) tend to grow faster. This is consistent with the exposure explanation (Romero, Meeder, and Kleinberg 2011). Retweets are significant predictors of growth for both winners and also-rans. This finding is consistent with arguments from organizational ecology (Carroll and Hannan 2000), which suggest that organizations and communities with narrow identities that fit closely to the environment tend to thrive in the short term. Hashtags popu-

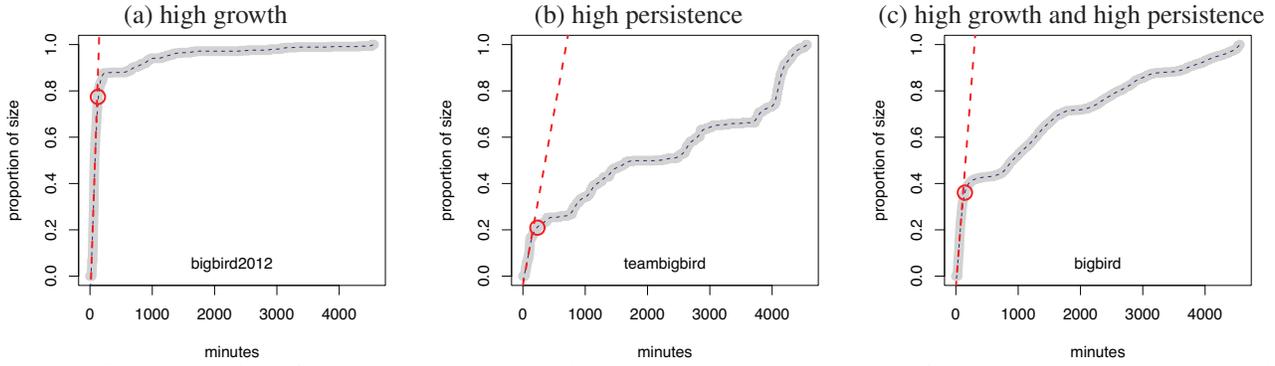


Figure 3: Characterization of hashtag growth and persistence. We capture the growth of a hashtag as the largest slope (in red line) over its fitted spline function (in blue dashed line).

Table 4: Growth models

Variables	Winner	Also-ran
$\mathbf{rt}_{t-1}$	0.0626** (0.0239)	0.2651*** (0.0073)
$\mathbf{rp}_{t-1}$	-0.0128 (0.0161)	0.1642*** (0.0108)
$\mathbf{src}_{t-1}^\alpha$	-0.0016*** (0.0002)	0.0000 (0.0000)
$\mathbf{follow}_{t-1}^\alpha$	0.1048** (0.0375)	0.0891*** (0.0042)
Loglik	-476.54	-26760.69
AIC	971.07	53539.38

lated by many retweets relevant to the Twitter community at a particular moment appear to grow more quickly, as users are drawn in to the “hot” conversation.

The results for replies and growth are not as simple. Replies appear to help also-ran hashtags to grow quickly, but do not appear to help “winner” hashtags. For also-ran hashtags, high interactivity would be associated with high growth, as users may be seeking to participate in conversations where it appears others are paying attention to what they say. The fact that replies do not appear to supply resources for the growth of winning hashtags is surprising, but begins to make sense in light of the findings for uniqueness (see below).

The results for uniqueness ( $\mathbf{src}_{t-1}^\alpha$ ) also show a distinction between the two classes. “Winner” hashtags appear to be constrained by diversity, whereas “also-ran” tags are not. Together with the results for replies, this suggests that these classes are distinguished by different conversational vibrancy. “Winner” hashtags appear to be those that gain strictly from their relevance to the environment, *i.e.*, they have interesting tweets. When fitness is high, diversity becomes a drag (Kauffman 1993). By contrast, “also-ran” tags appear to be less relevant. At the same time, what is important for growth is the ability to bring something beyond relevance—interaction with others.

## Explaining Hashtag Persistence

To examine the effect of earlier conversational vibrancy on a hashtag’s persistence, we use survival analyses based on the Cox proportional-hazards model (Cox 1972). We examine the survival time of a hashtag starting from the turning point  $t^*$  to the saturated time  $t_e$ . Let  $T$  be a continuous random variable, the survival function is defined as:  $S(t) = Pr(T > t)$ , the probability that the hashtag will

survive (*i.e.*, not saturated) beyond  $t$ .

A hazard function assesses the instantaneous risk of demise at time  $t$ , conditional on survival to that time.

$$h(t) = \lim_{\Delta t \rightarrow 0} \frac{Pr[(t \leq T < t + \Delta t)]}{\Delta t}$$

In the Cox regression model, the log hazard can be specified as:

$$h(t) = h_0(t) \exp(\beta^T \mathbf{x})$$

where  $\mathbf{x}$  is a vector of predictors,  $\beta$  is a vector of parameters to estimate,  $h_0(t)$  is the baseline hazard.

To capture a hashtag’s earlier conversational vibrancy, we focus on activities occurring in its fast growth phase. Hence the predictors we choose to examine are aggregate measures for each hashtag at the turning point  $t^*$ , including:

- $\mathbf{rt}_{t^*}^\alpha$ : the total number of retweets received up to the turning time point  $t^*$ , *i.e.*, between  $t_0$  and  $t^*$ .
- $\mathbf{rp}_{t^*}^\alpha$ : the total number of replies posted up to time  $t^*$ .
- $\mathbf{src}_{t^*}^\alpha$ : the total number of retweet sources up to time  $t^*$ .
- $\mathbf{follow}_{t^*}^\alpha$ : the expected largest audience size up to time  $t^*$ . The computation is similar to  $\mathbf{follow}_{t-1}^\alpha$  except that the time is fixed at  $t^*$ .

The notation  $(\cdot)^\alpha$  indicates that the predictor is an aggregate measure. Table 5 summarizes the results for the survival analyses. All coefficient are exponentiated and thus interpretable as multiplicative effects on the hazard (standard error in parenthesis are not exponentiated). For the winner class, the predictors  $\mathbf{rt}_{t^*}^\alpha$  and  $\mathbf{rp}_{t^*}^\alpha$  have significant coefficients. For example, an additional reply reduces the per minute hazard of hashtag saturation by 0.65%. A positive coefficient increases the value of the hazard function and therefore indicates a negative effect on survival time. In other words,  $\mathbf{rp}_{t^*}^\alpha$  has a positive association with a hashtag’s persistence and  $\mathbf{rt}_{t^*}^\alpha$  has a negative association with the persistence. The other predictors,  $\mathbf{src}_{t^*}^\alpha$  and  $\mathbf{follow}_{t^*}^\alpha$ , have no significant effect. For the also-ran class,  $\mathbf{rt}_{t^*}^\alpha$  has negative effect while  $\mathbf{src}_{t^*}^\alpha$  has positive effect on the hashtag persistence.

As expected, and consistent with exposure explanations, the number of followers is not a significant predictor of a tag’s longevity. Also consistent with the explanation for growth is the finding for retweets. Organizational ecology predict that organizations and communities with specific identities strongly coupled with the environment are risk of

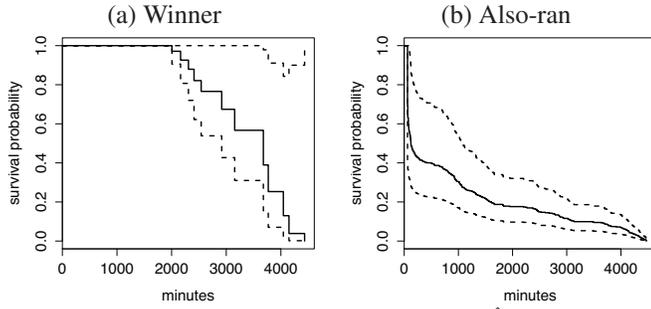


Figure 5: Estimated survival function  $\hat{S}(t)$  for the two classes. Half of the hashtags in the winner class lived beyond 2.5 days, and half of the hashtags in the also-ran class died out (became saturated) in two hours. The dash lines show a point-wise 95% confidence envelope around the survival function.

Table 5: Persistence models

Variables	Winner	Also-ran
$\mathbf{rt}_{t^*}^\alpha$	1.001* (0.0003)	1.001* (0.0003)
$\mathbf{rp}_{t^*}^\alpha$	0.9935* (0.0003)	1.000 (0.0002)
$\mathbf{src}_{t^*}^\alpha$	1.001 (0.0014)	0.9933*** (0.0019)
$\mathbf{follow}_{t^*}^\alpha$	1.000 (0.0000)	1.000 (0.0000)
Loglik	-15.92968	-1028.731
AIC	39.85936	2065.461

demise when the environment changes (Carroll and Hannan 2000). The very thing that made them appealing when the environment was suited to what they had to offer now makes them less useful. For example, a hashtag that mainly contains jokes about the Big Bird reference may be popular in the initial frenzy after the comment. As the debate moves on or ends and serious issues come to the fore, the humorous hashtags lose their relevance to the new environment.

Interactivity and diversity present a more complex set of findings. More replies increase the longevity of “winner” hashtags while having no effect for “also-rans.” Simultaneously, diversity extends the life of also-rans but has no effect for winners. This is additional evidence of a qualitative difference between the two classes of hashtags. Replies appear to buffer winner tags from the inevitable decline in their relevance over time (Baum and Oliver 1991). That is, “winner” hashtags with many replies may be able to fall back into the niches that were originally occupied by the “also-rans”—a place to exchange ideas about a topic, rather than just to broadcast one’s opinions or humorous insights. The fact that “also-rans” benefit from diversity, but “winners” do not, suggests that some “also-rans” may build up a sufficient community that wishes to continue the conversation even as new topics emerge and relevance declines. Fig. 5 shows the estimated survival function  $\hat{S}(t)$  for the two classes. In the winner class, half of the hashtags lived beyond 2.5 days. In the also-ran class, half of the hashtags died out (became saturated) in two hours.

### Examining Environmental Context

In our analysis, all the “pop” hashtags emerged from a unusual condition where users’ attention was concentrated on

Table 6: Growth models with co-occurring hashtags

Variables	Winner	Also-ran
$\mathbf{rt}_{t-1}$	0.0596* (0.0240)	0.2817*** (0.0073)
$\mathbf{rp}_{t-1}$	-0.0141 (0.0161)	0.1560*** (0.0109)
$\mathbf{src}_{t-1}^\alpha$	-0.0016*** (0.0002)	0.0013*** (0.0003)
$\mathbf{follow}_{t-1}^\alpha$	0.0989** (0.0380)	0.0811*** (0.0042)
$\mathbf{rtEnv}_{t-1}$	-0.0028 (0.0271)	0.0530*** (0.0043)
$\mathbf{rpEnv}_{t-1}$	0.0113 (0.0174)	0.0173*** (0.0038)
$\mathbf{srcEnv}_{t-1}^\alpha$	0.0014 (0.0016)	0.0000 (0.0000)
Loglik	-475.89	-26793.36
AIC	975.79	53610.72

Table 7: Persistence models with co-occurring hashtags

Variables	Winner	Also-ran
$\mathbf{rt}_{t^*}^\alpha$	1.001 (0.0004)	1.001** (0.0003)
$\mathbf{rp}_{t^*}^\alpha$	0.9904* (0.0045)	1.000 (0.0002)
$\mathbf{src}_{t^*}^\alpha$	1.002 (0.0021)	0.9928*** (0.0019)
$\mathbf{follow}_{t^*}^\alpha$	1.000 (0.0000)	1.000 (0.0000)
$\mathbf{rtEnv}_t$	1.428 (2.688)	0.7173* (0.1502)
$\mathbf{rpEnv}_t$	1.548 (2.288)	0.6976* (0.1709)
$\mathbf{srcEnv}_t^\alpha$	2.023 (1.560)	1.020 (0.01799)
Loglik	-11.65331	-1023.476
AIC	37.30663	2060.952

debate events. It is possible that a hashtag’s growth and death is largely determined by this particular environmental condition rather than the conversational vibrancy of the hashtag itself. To test the robustness of our findings against these concerns, we re-construct the prior growth and persistence models and include variables for users’ tweeting activity for other hashtags at the same moment.

We extend the growth models to include three time-dependent environmental covariates. For each hashtag  $i$ , we have:

- $\mathbf{rtEnv}_{t-1}$ : the total number of new retweets that do not contain hashtag  $i$ , at previous time  $t - 1$ .
- $\mathbf{rpEnv}_{t-1}$ : the number of replies that do not contain hashtag  $i$ , at time  $t - 1$ .
- $\mathbf{srcEnv}_{t-1}^\alpha$ : the number of new retweet sources that do not contain hashtag  $i$ , up to time  $t - 1$ .

Based on the time-series regression described above, we include the three new predictors into new models of growth. The estimated results are summarized in Table 6. Interestingly, none of the environmental covariates has significant effect on the growth of winner hashtags. However, the predictors  $\mathbf{rtEnv}_{t-1}$  and  $\mathbf{rpEnv}_{t-1}$  are positively correlated to the growth of “also-ran” hashtags. While these effect are not as strong as  $\mathbf{rt}_{t-1}$  and  $\mathbf{rp}_{t-1}$ , the results suggest that “also-ran” hashtags benefit from other tweeting activities.

To study the relationship between hashtag’s persistence and the environment, we include three new predictors in the survival analyses. Unlike the original four predictors in the persistence models which were measured at a fixed point in time ( $t^*$ ), the new predictors are measured at each time  $t$ , where  $t$  is between the turning point  $t^*$  and the saturated time  $t_e$ . The time-dependent covariates measures “what was

going on in the environment” before the hashtag’s saturation. For each hashtag  $i$ , we have:

- $\mathbf{rtEnv}_t$ : the total number of new retweets that do not contain hashtag  $i$ , at time  $t$ .
- $\mathbf{rpEnv}_t$ : the number of replies that do not contain hashtag  $i$ , at time  $t$ .
- $\mathbf{srcEnv}_t^\alpha$ : the number of new retweet sources that do not contain hashtag  $i$  up to time  $t$ .

These three new time-dependent covariates, together with the four fixed point covariates, are included in the new persistence models. The estimated results are summarized in Table 7. Like the growth model, we see no impact from the environment on this model for predicting winner hashtags’ persistence. For also-ran hashtags, the  $\mathbf{rtEnv}_t$  and  $\mathbf{rpEnv}_t$  increase the hazard of becoming saturated and they are negatively associated with persistence. This reflects the fact that many of the also-ran hashtags died out in a couple of hours where there were still high tweeting activities going on in the environment.

## Discussion

User-generated content with high levels of conversational vibrancy should be a focal point for attracting the attention and participation of users. These elements should satisfy a confluence of users’ needs by providing topical information, interactive sociality, diverse and novel content, and involvement from well-connected users and lead to a feedback loop that reinforces these tendencies. Operationalizing this model for Twitter in particular, we investigated how these features of conversational vitality influenced both the growth and persistence of adoption for different classes of hashtags.

**Summary of findings.** The number of times a hashtag is retweeted (topicality) as well as the popularity of the users mentioning the hashtag (prominence) lead to more rapid growth of hashtags for “winner” and “also-ran” hashtags alike. Additional replies (interactivity) and the number unique retweet sources (diversity) support the persistence of “winner” and “also-ran” hashtags, respectively. In addition, our findings unexpectedly suggest that the number of retweets inhibit the growth of hashtags. Because these findings remain robust after controlling for other contextual behavior of users, these findings provide mixed evidence for the simple cumulative vibrancy model we proposed. In particular, the fact that additional incremental activity inhibits the growth and persistence of hashtags suggests there are higher-order processes that lead to limits or tipping points.

One direction for extending this conversational vibrancy model would be to draw upon theories from organizational ecology that describe the birth, growth, and death of organizations and communities. Ecologists treat organizations as entities which, like hashtags, coordinate behavior through identities that suggest the kinds of behaviors and messages that are appropriate within organizations of a particular form (Carroll and Hannan 2000). In particular, ecological research focuses on the importance of balancing consistency and coherence with flexibility and access to resources. The environment provides a limited supply of resources for organizations (laborers, customers, *etc.*) analogous to how the

Twitter environment provides a limited supply of resources for hashtags (attention, users, ideas to express, *etc.*). Organizations survive and thrive when their identities are specific enough such that individuals know what to expect from them but broad enough that they can address a range of needs and access a variety of resources. It has been recognized that organizations with similar identities tend to thrive as their population grows, with each gaining more attention and legitimacy. At some point a limit is reached and the environment can no longer support most of these organizations, and only a few remain (Carroll and Hannan 2000).

A substantial portion of the research on growth and sustainability involves the comparison of cultural forms, such as words, names, memes, or networks. Both exposure and “fitness” play important roles in these models, suggesting rationales for both self-reinforcing growth and saturation (Steyvers and Tenenbaum 2005; Barabasi 2003). Studies of baby naming conventions have found that burstiness is associated with short life-time because collective aversion to “faddishness” limits the growth of a name (Berger and Le Mens 2009). These findings are corroborated by other work on memes that suggest social media have carrying capacities (Leskovec, Backstrom, and Kleinberg 2009). Our model introduces an additional factor to consider: the endogenous and developing qualities of the communication and interactions between those that use the form, serving to weaken the determinism suggested by fitness and first-mover based explanations.

The application of this model to linguistic evolution may be particularly appealing. The ways in which linguistic conventions are shared by a group have been investigated in the field of language dynamics (Loreto, Baronchelli, and others 2011). Computational models have shown that simple linguistic interactions can lead the group to reach a consensus on a given linguistic convention by the progressive elimination of competing synonyms (Steels 1995; Baronchelli et al. 2006). Interestingly, these models assume that the competing synonyms are perfectly equivalent. While the situation is clearly more complex with hashtags, our research demonstrates that the observed winner-takes-all dynamics are not necessarily driven only by intrinsic properties of hashtags but also by other contextual and structural features which suggest new avenues for research across fields like communication, social computing, linguistics, and political science.

Our approach in this analysis also has limitations that merit future research. We do not examine more platform-specific contextual features such as the effect of a hashtag’s length (number of characters) on its life-cycle, which has significance due to Twitter’s 140-character limit for a message (Yang et al. 2012). Despite the highly situational and socially-embedded interactions in social media, our modeling approach assumed the observation of hashtags on users are the result of an independent process. Future work could specify alternative models that account for the local clustering of ties or the modular structure of these communities which likely strongly influence the growth and persistence of hashtag use.

## Conclusion

We introduced a theoretical framework called conversational vibrancy that synthesizes work from social computing and mass communication to understand why some forms of user-generated content are adopted while others are abandoned. The U.S. presidential debates allowed us to understand the lifecycle of hashtag adoption and use in response to exogenous shocks that allow us to separate out the confounding influence of topical relevance and social exposure. Using simple computational means to characterize the key aspects of hashtags's complex dynamics, statistical models reveal qualitative differences in how conversational vibrancy influences the growth and persistence of distinct classes of hashtags. The study also demonstrates the power of computational social science (Lazer et al. 2009) approaches to meld data-driven computational and statistical approaches to investigate fundamental questions about human sociality.

## Acknowledgments

The authors gratefully acknowledge the support of the Lazer Lab at Northeastern University, supported in part by MURI grant #504026, DTRA grant #509475, and ARO #504033.

## References

- Barabasi, A.-L. 2003. Linked: How everything is connected to everything else and what it means.
- Baronchelli, A.; Felici, M.; Loreto, V.; Caglioti, E.; and Steels, L. 2006. Sharp transition towards shared vocabularies in multi-agent systems. *JSTAT* 2006(06):P06014.
- Baum, J. A., and Oliver, C. 1991. Institutional linkages and organizational mortality. *Administrative Science Quarterly* 187–218.
- Berger, J., and Le Mens, G. 2009. How adoption speed affects the abandonment of cultural tastes. *PNAS* 106(20):8146–8150.
- Boyd, D.; Golder, S.; and Lotan, G. 2010. Tweet, tweet, retweet: Conversational aspects of retweeting on twitter. In *HICSS'10*, 1–10. IEEE.
- Carroll, G. R., and Hannan, M. T. 2000. *The demography of corporations and industries*. Princeton University Press.
- Clayman, S. E. 1995. Defining moments, presidential debates, and the dynamics of quotability. *J. of Communication* 45(3):118–147.
- Cox, D. R. 1972. Regression models and life-tables. *Journal of the Royal Statistical Society. Series B (Methodological)* 187–220.
- Crane, R., and Sornette, D. 2008. Robust dynamic classes revealed by measuring the response function of a social system. *PNAS* 105(41):15649–15653.
- Durbin, J., and Koopman, S. 2001. Time series analysis by state space methods.
- Huang, J.; Thornton, K. M.; and Efthimiadis, E. N. 2010. Conversational tagging in twitter. In *Proc. of HT'10*, 173–178. ACM.
- Kauffman, S. A. 1993. *The origins of order: Self-organization and selection in evolution*. Oxford University Press, USA.
- Lazer, D.; Pentland, A.; Adamic, L.; et al. 2009. Social science. computational social science. *Science* 323(5915):721.
- Lehmann, J.; Gonçalves, B.; Ramasco, J. J.; and Cattuto, C. 2012. Dynamical classes of collective attention in twitter. In *Proc. of WWW'12*, 251–260. ACM.
- Lerman, K., and Ghosh, R. 2010. Information contagion: An empirical study of the spread of news on digg and twitter social networks. In *Proc. ICWSM'10*. AAAI.
- Leskovec, J.; Backstrom, L.; and Kleinberg, J. 2009. Meme-tracking and the dynamics of the news cycle. In *Proc. of SIGKDD'09*, 497–506. ACM.
- Lin, Y.-R.; Margolin, D.; Keegan, B.; and Lazer, D. 2013. Voices of victory: A computational focus group framework for tracking opinion shift in real time. In *Proc. of WWW'13*. ACM.
- Loreto, V.; Baronchelli, A.; et al. 2011. Statistical physics of language dynamics. *JSTAT* 2011(04):P04006.
- Marwick, A. E., and boyd, d. 2011. I tweet honestly, i tweet passionately: Twitter users, context collapse, and the imagined audience. *New Media & Society* 13(1):114–133.
- PEW. 2012. One-in-ten "dual-screened" the presidential debate. Technical report.
- Romero, D. M.; Meeder, B.; and Kleinberg, J. 2011. Differences in the mechanics of information diffusion across topics: idioms, political hashtags, and complex contagion on twitter. In *Proc. of WWW'11*, 695–704. ACM.
- Romero, D. M.; Tan, C.; and Ugander, J. 2011. Social-topical affiliations: The interplay between structure and popularity. *arXiv preprint arXiv:1112.1115*.
- Ruggiero, T. E. 2000. Uses and gratifications theory in the 21st century. *Mass Communication & Society* 3(1):3–37.
- Schroeder, A. 2008. *The Presidential Debates: Fifty Years of High-Risk TV*. Columbia University Press, second edition.
- Shamma, D.; Kennedy, L.; and Churchill, E. 2009. Tweet the debates: understanding community annotation of uncollected sources. In *Proc. SIGMM'09*, 3–10. ACM.
- Shamma, D. A.; Kennedy, L.; and Churchill, E. F. 2011. Peaks and persistence: modeling the shape of microblog conversations. In *Proc. of CSCW'11*, 355–358. ACM.
- Skewes, E. A. 2007. *Message Control: How News is Made on the Campaign Trail*. Rowman & Little.
- Stafford, T. F.; Stafford, M. R.; and Schkade, L. L. 2004. Determining uses and gratifications for the internet. *Decision Sciences* 35(2):259–288.
- Starbird, K., and Palen, L. 2012. (how) will the revolution be retweeted?: information diffusion and the 2011 egyptian uprising. In *Proc. of CSCW'12*, 7–16. ACM.
- Steels, L. 1995. A self-organizing spatial vocabulary. *Artificial life* 2(3):319–332.
- Steyvers, M., and Tenenbaum, J. B. 2005. The large-scale structure of semantic networks: Statistical analyses and a model of semantic growth. *Cognitive science* 29(1):41–78.
- Suh, B.; Hong, L.; Pirolli, P.; and Chi, E. 2010. Want to be retweeted? large scale analytics on factors impacting retweet in twitter network. In *Proc. SocialCom'10*, 177–184. IEEE.
- Sutton, J.; Palen, L.; and Shklovski, I. 2008. Backchannels on the front lines: Emergent uses of social media in the 2007 southern california wildfires. In *Proc. of ISCRAM'08*, 624–632. Washington, DC.
- Yang, J., and Counts, S. 2010. Predicting the speed, scale, and range of information diffusion in twitter. In *Proc. ICWSM'10*. AAAI.
- Yang, L.; Sun, T.; Zhang, M.; and Mei, Q. 2012. We know what @you #tag: does the dual role affect hashtag adoption? In *Proc. of WWW'12*, 261–270. ACM.