

Investigating the Observability of Complex Contagion in Empirical Social Networks

Clay Fink, Aurora Schmidt
JHU Applied Physics Lab., Laurel, MD
{clayton.fink, aurora.schmidt}@jhuapl.edu

Vladimir Barash, John Kelly
Graphika, Inc., New York, NY
{vlad.barash, john.kelly}@graphika.com

Christopher Cameron, Michael Macy
Cornell University, Ithaca, NY
{cjc73, m.macy}@cornell.edu

Abstract

Social contagion is the mechanism by which ideas and behaviors spread across human social networks. *Simple* contagion models approximate the likelihood of adoption as constant with each exposure to an “infected” network neighbor. However, social theory postulates that when adopting an idea or behavior carries personal or social risk, an individual’s adoption likelihood also depends on the number of distinct neighbors who have adopted. Such *complex* contagions are thought to govern the spread of social movements and other important social phenomena. Online sites, such as Twitter, expose social interactions at a large scale and provide an opportunity to observe the spread of social contagions “in the wild.” Much of the effort in searching for complex phenomena in real world contagions focuses on measuring user adoption thresholds. In this work, we show an alternative method for fitting probabilistic complex contagion models to empirical data that avoids measuring thresholds directly, and our results indicate bias in observed thresholds under both complex and simple models. We also show 1) that probabilistic models of simple and complex contagion are distinguishable when applied to an empirical social network with random user activity; and 2) the predictive power of these probabilistic adoption models against observed adoptions of actual hashtags used on Twitter. We use a set of tweets collected from Nigeria in 2014, focusing on 20 popular hashtags, using the follow graphs of the users adopting the tags during their initial peaks of activity.

Introduction

It is thought that ideas and behaviors spread through human populations in ways similar to the spread of biological infections. By way of analogy, such processes are called social contagions. When there is little or no cost associated with adopting an idea or behavior, only a single exposure to an “infected” person is needed for another person to adopt. Such *simple* contagions contrast with cases where adoption carries a risk of losing social capital, or exposing oneself to official sanction or physical harm. In such cases, a person’s decision to adopt may require social reinforcement; his or her likelihood of adoption may greatly

increase after knowing that some members of their social network have already adopted. These threshold-based processes are termed *complex* contagions and they are thought to play a role in the outbreak of riots, mass protests, and the spread of social and political change (Granovetter 1973; Coleman, Katz, and Menzel 1966; Valente 1996). While collective behaviors are directly observable, theories about the underlying social processes at play have been difficult to confirm outside of laboratory experiments and simulations. Social networking sites, such as Twitter and Facebook, expose social interactions at a larger scale, giving us the ability to test statistical models related to the spread of social contagions. Given the assumptions about the role of complex contagions in the spread of controversial information and the formation of transformative social movements, identifying their onset and modeling their spread is of direct applicability for journalists, scholars, and policy makers.

In this paper, we present an approach for accomplishing this goal. We use Twitter data collected from users geographically localized to Nigeria in 2014, including tweets that used one or more of a set of popular hashtags. We collected follow information on a set of 52,689 of these users, comprising the first 1000 adopters of each of 82 popular hashtags that were initially propagated in 2014. Using the actual follow network for these individuals, we ran a set of experiments where we simulated probabilistic models of simple and complex contagion. We first show that by using the true follow network over the same set of users, simple and complex contagions predict statistically differing adoption curves, suggesting these phenomena are distinguishable in real world network contagions. We then investigate the degree to which the probabilistic simple and complex contagion models can predict the actual spread of a set of popular hashtags, optimizing infection probability parameters to best fit the empirical adoption curve¹ for each model. We show that for a sample hashtag that is likely to carry social risk, the adoptions observed are better explained by the complex contagion model than by the simple contagion model, lending evidence that social reinforcement models of contagion spread may be at work.

¹By “empirical adoption curve” we mean the cumulative number of adoptions versus time.

Our work differs in a number of ways from previous efforts. First, most other work looks at the spread of information online without regard to geography. In our work, by focusing on particular region we can observe online diffusion events that are more likely to be relevant to important offline events in that region. Secondly, our work differs from prior investigations that use simulated rather than observed network topologies: either randomly rewired lattice graphs (Centola and Macy 2007; Barash, Cameron, and Macy 2012) or preferential attachment graphs (Goel et al. 2015). Thirdly, our simulations use an asynchronous model of user activity rather than a global synchronous clock for all agents. This more accurately models the way Twitter is actually used, and enables us to model the timing of information spread more realistically. Finally, some of the work comparing simple and complex contagion uses deterministic rules for adoption rather than modeling probabilistic adoption as a function of exposures and sources of influence (Centola and Macy 2007; Barash, Cameron, and Macy 2012). While we do not claim that these statistical models accurately model choice for a single individual, they allow us to model the population level variability in adoption behavior. Another key feature of our approach is that we do not rely on the calculation of empirical adoption thresholds as in (Romero, Meeder, and Kleinberg 2011; Barash 2011). In real world settings, due to intermittent attention and activity, the observed numbers of exposures and sources at the time of adoption are likely to be much higher than the actual number that was needed for the individual to adopt (Cameron 2016). Using our probabilistic contagion models and random activity models we simulate these delayed adoptions and show the bias in observed thresholds. Because of this large bias in estimated adoption thresholds, we show that distributions of observed adoption thresholds are not predictive of whether a complex or simple contagion process is being observed.

Our approach, in summary, focuses on what we can observe - the empirical adoption curve for a hashtag - and determines which contagion model best fits the data.

Related Work

Granovetter (Granovetter 1978) describes a class of aggregate group behaviors that are dependent on people choosing to engage in the behavior only after a certain number of others have chosen to engage in that behavior. Slight changes in the distribution of adoption thresholds across a population can lead to quite different aggregate outcomes, and some of this heterogeneity of thresholds depends on local differences in network structure as well as temporal and spatial factors. Riots, strikes, leaving a party that appears to be winding down, the spread of rumors, and the diffusion of innovations are all given as examples of collective behaviors that depend on adoption thresholds. Centola and Macy (Centola and Macy 2007) cast Granovetter's threshold model of collective behavior as describing complex contagions, and investigate how they spread across networks. In responding to work by Granovetter (Granovetter 1973) about the importance of long range ties in a network for the spread of ideas and behaviors, they use simulations to demonstrate that while complex contagions will saturate

dense regions of a network with many redundant ties, their spread to uninfected regions of a network are impeded unless there are a sufficient number of redundant long range ties. Centola (Centola 2010) designed an experiment using a constructed real-world online network and tested threshold behavior on highly clustered lattice networks and randomly rewired lattice networks. All nodes had the same degree and the experiment was repeated for different node degrees and network sizes. He found that the contagions spread faster and further on the clustered than on the random networks giving some credence to his results in (Centola and Macy 2007) about the importance of redundant ties for the spread of complex contagions. In related work, Barash, et al. (Barash, Cameron, and Macy 2012) simulated complex contagions on lattice and scale free networks and found that complex contagions require a critical mass of adopters before there were sufficient redundant long ties for the contagion to spread to remote regions of the network. The theoretical work on complex contagion demonstrates that, except for highly connected dense networks, there are significant barriers for the success of complex contagions, illustrating Granovetter's original descriptions of these being fragile processes.

Given the growth in user generated content on the Web in recent years, researchers have turned to social media data to look for empirical support for theories of complex contagion. Barash (Barash 2011) looked at tags of photos on Flickr, a photo sharing service, which allows users to comment on and tag photographs. A few tags relating to offline behavior that entailed some risk ("jailbreaking" iPhones was one example) showed some evidence of complex contagion. Romero, et al. (Romero, Meeder, and Kleinberg 2011) used a set of eight million tweets from 2009-2010, selected the top 500 hashtags in terms of user mentions, and categorized the tags by topic categories (politics, celebrities, sports, music, and Twitter idioms). From the tweets, they constructed a network based on mentions of other users. For the initial use of a hashtag by a user they calculated $p(k)$: the percentage of users with k or more exposures that adopted a tag with exactly k exposures. They found that the $p(k)$ curves (effectively, the range of probabilities of calculated threshold values) were consistently different for political tags in that $p(k)$ was higher and remained high as k increased. They also found that the initial adopters of political tags formed denser networks, which they argue is consistent with complex contagion and the need for social reinforcement for the adoption of controversial ideas or behaviors. Using the calculated $p(k)$ values for users, they simulated cascades on the constructed networks using the original seed sets of adopting users, and also using random seeds. They found that the empirical cascades of political and idiom tags grew at a faster rate than cascades based on random seeds.

Goel et al. (Goel et al. 2015) looked at Twitter at a much larger scale, using a collection of 1.2 billion tweets from 2011-2012, each tweet containing a URL pointing to one of a set of specific Web sites. Their definition of adoption was a user authoring a tweet containing one of the 622 million unique URLs they identified. They defined a measure of "structural virality" for a cascade of adoptions that bal-

ances the width of the cascade due to broadcasts from a single node, and the depth of the cascade. Most cascades were small, so they focused on the 0.025% of cascades with 100 or more nodes. They found a diversity of cascade structures, including large broadcasts, deep tree structures, and many hybrids of both. The URLs were categorized as links to videos, news, or online petitions. They discovered that 25% of petitions had high structural virality compared to small percentages for other categories, and that the structural virality of petitions was high across different cascade sizes. This result is interesting and in line with the finding in (Romero, Meeder, and Kleinberg 2011) that the initial density of adopters was higher for political hashtags. One interpretation is that, similar to using a political hashtag, linking to an online petition requires social reinforcement, and denser local network structure facilitates this reinforcement through redundant ties. Finally, they ran simulations on a simulated preferential attachment graph using a susceptible-infected (SI) infection model and a constant infection rate. The range of cascades generated showed many of the distributional aspects of the empirical cascades; a small number of large cascades, a similar average virality measure, and invariance between size and structural virality. The results of their simulations suggest that a simple contagion model can explain the diversity of the real-world cascades found in their Twitter data set.

State and Adamic (State and Adamic 2015) looked for evidence for complex contagion on Facebook in the context of people registering support for same-sex marriage by overlaying their profile pictures with the equal sign symbol associated with that social movement. Working from anonymized data from three million Facebook profiles, they used the volunteered demographic information associated with the profiles to investigate the link between these factors and the number of friends that had to change their photo before a user made the decision to change their own photo. They found, for example, that users with larger numbers of friends who indicated same-sex attraction in their profile were more likely to adopt (i.e. change their profile pictures) than those with few or no friends who indicated same-sex attraction. Adoptions increased with the number of friends adopting, and this was even more apparent in cases where a user registered an interest in their friend's adoption either by clicking on a posted link or "liking" a post relating to same sex marriage. They looked at adoption thresholds across other content and found for "cut-and-paste memes," in cases where a user may be exposing themselves to questions or ridicule, a higher adoption threshold was required. Finally, simulations were run on an empirical network of 800 thousand users. They modeled contagion based on user susceptibility (drawn at random from an exponential distribution), and an infection probability that decreases logarithmically with the number of adopting friends. For higher probabilities of infection, adoption thresholds increased with the number of friends adopting, as was seen in the empirical data.

As we mentioned in the introduction, measuring the empirical adoption threshold for a user (i.e. k) is problematic because users check their timelines asynchronously, ob-

scuring our ability to observe an individual's true adoption threshold. The problem also exists in simulations with synchronous user activity when multiple neighbors adopt in the same time-step. Of the four studies cited, only State (State and Adamic 2015) addresses this confound by restricting adoption behaviors to cases where there is evidence that the adopting user was paying attention to the source by looking at click-throughs on posted links and users clicking the "like" button. In this work, we demonstrate the bias in observed k for both simple and complex adoption models using randomized asynchronous activity. We seek to infer that a contagion is complex, as opposed to simple, based on its potential to be explained by complex versus simple models, rather than by direct estimation of users' k thresholds. Our work is also similar to State and Adamic in that we carry out case studies on the diffusion of specific hashtags, similar to their focus on adoption of the equal sign overlay of profile pictures. Both Goel (Goel et al. 2015) and Romero (Romero, Meeder, and Kleinberg 2011) instead focused on categories of content and gross statistics on the number of adoptions.

Data

In this paper, we focus on Twitter data from Nigeria. Nigeria, with over 170 million people, is the most populous country in Africa and its largest economy. Internet penetration in the country is in line with other major countries in Africa (International Telecommunications Union 2015) with most people connecting to the Web via mobile devices (Odufuwa 2012). English and Nigerian Pidgin English are the languages primarily used online in Nigeria, with some representation of local languages such as Hausa, Igbo, and Yoruba. Code switching between English, Pidgin English, and the local languages is common as well.

We collected tweets using the public Twitter API, which allows for geographic queries based on a pair of geographic coordinates and a radius. The queries return tweets that the API determines are from within the defined area based on either the user's stated location or, if their tweets are geo-tagged, their latitude and longitude. We also collected tweets from the timelines of users whose tweets were in the results of the geographic queries in order to collect as complete a sample as possible from the region. We ran persistent crawls of tweets from around 45 Nigerian cities with populations of 100,000 or more using a radius varying from 25 to 40 miles. Data collection began in 2010. In this work, we focus on data collected from January 2014 through November 2014, resulting in a data set containing 147 million tweets from 1.9 million users.

We extracted hashtags from the 2014 tweets that had not occurred in the complete collection prior to January 15, 2014. This was to ensure that we were only looking at new hashtags that were not part of previous diffusion events. We then selected tags that had 1,000 or more adopters. Of these 82 tags, 42 tags had 1,000 or more adopters in the first three weeks of their lifespan. For this paper, we focus on 20 of these 42 as a case study, restricting our attention to adoptions during the first 24 hours of hashtag usage. These 20 hashtags analyzed were the tags in which more than 50 non-seed adoptions were observed in that first 24 hour period.

Hashtag	Users	Description
AmericaWillKnow	4390	Reference to quote by Pres. Jonathan about missing US funds
BRAvsGER	2791	World cup game
BringBackOurDaughters	6963	Kidnapping of schoolgirls by Boko Haram in April 2014
BringBackOurGirls	94513	Kidnapping of schoolgirls by Boko Haram in April 2014
BringBackOurGirlsAlive	3622	Kidnapping of schoolgirls by Boko Haram in April 2014
BringBackOurSisters	1600	Kidnapping of schoolgirls by Boko Haram in April 2014
Ebolafacts	4375	Campaign for awareness about Ebola
Makeuptransformation	1609	Viral trend showing steps to applying makeup
MH17	5308	Malaysian airliner shot down over Ukraine
MH370	8127	Malaysian airliner lost over Indian Ocean
MTNTcheleteGoodLife	3590	Promotion by MTN using the song "Tchelete"
Nyanyablast	2340	Boko Haram bombing in Abuja
RIPRobinWilliams	2122	Related to passing of actor Robin Williams
SaveYakubuYusuf	1885	Raising money for Nigerian student's cancer treatment
TheChibokGirls	1971	Kidnapping of schoolgirls by Boko Haram in April 2014
WeAreAllMonkeys	2435	Campaign against racist hooligans at soccer games
WelcomeDiMaria	1603	Player for Manchester United
WetinBeLove	1968	Pidgin for "what is love?"
WhatJayZSaidToSolange	2129	Jay Z attacked by Solange
YesAllWomen	1615	Related to Santa Barbara shooting in June 2014

Table 1: The 20 hashtags for use in this case study with total number of adopting users in our Nigerian data set and a description.

The remaining popular hashtags will require either longer simulation periods or alternative methods for selecting the time period to begin simulation analysis, and so we leave these to future work. The 20 tags are listed in Table 1. Five of the 20 tags are related to the kidnapping of over 200 schoolgirls in April 2014 by the Islamist group, Boko Haram, an event that eventually gained worldwide visibility.

The follow graphs of the 52,689 users that were in the first 1,000 adopters of any of the 42 tags were then crawled using the public Twitter API. These crawls were carried out in the Spring of 2015, so the extracted social network ties were not guaranteed to be the same as those present during the diffusion events analyzed. A follow relationship exists if a user decides to click the follow button for another user, resulting in that user's tweets being posted to the following user's timeline.² The set of users a person follows gives us an estimate of the number of potential sources of information that the user will be exposed to on Twitter via their timeline. In cases where a followed user retweets a third user that the following user does not follow, the follow graph does not include that user as a potential source. Thus, the follow graph may underestimate the total number of potential sources. The follow graph used for analysis is a closed-world network and restricted to only the 52,689 users for which we retrieved follow graph information, defining a network with 3.8 million edges, in which the average degree is approximately 114, representing about 30% of the follow ties from the full non-closed network.

To ensure that we were primarily analyzing the activity of Nigerians, we matched the geographic metadata returned by the API with each tweet against a gazetteer. We were able to geo-locate 58% of these users to Nigeria; 13% of the 52,689 users had no location information. To detect the presence of

automated Twitter accounts, we used heuristics inspired by Lee, et al. (Lee, Eoff, and Caverlee 2011). These measures identified 427 suspect accounts, less than 1% of the total 52,689 users. By inspection, many of the more active accounts were for news sources or news aggregators. We retain these accounts in our set of analysis users, however, since such accounts are a common part of the Twitter ecosystem and may play an important role in actual diffusion events.

Methods

The aim of our work is to test for statistical separability of adoption curves generated by simple and complex models using actual Twitter networks and contagion events. We model random adoptions according to probabilistic simple or complex models. In addition, we account for random Twitter login activity by using a Poisson model for user logins. We do not model the creation of seed adopters in our network. Seeds include users whose adoptions appear to precede any influence from neighbors in the follow graph. In our dataset, seeds could be observed either due to a totally novel adoption or to influence external to our subgraph.³ In our tests, the true schedule of seed adoptions and uses of the hashtag is replayed during the simulated contagion. In addition, we do not model a user's decision to repeat the use of the hashtag. Rather than characterize that process with another statistical model for behavior, since it is outside the scope of our investigation, we use the actual schedule of subsequent hashtag use for each user. In the simulation, once that user adopts, he or she re-exposes his or her followers in the network according to this fixed time-delay schedule.

²Note that the follow relationship is not reciprocal.

³External influence could be from other online networks or offline channels, as well as missing data in our user or follow data.

Random Timeline Checking

The distribution of logins is modeled as Poisson. Despite the fact that activity rates vary from individual to individual, for simplicity, we model a single average rate parameter, λ , for the entire population. The rate of timeline checking governed by the parameter was assumed to be uniform across the population. We use the full set of tweets collected for our analyzed users - excluding the accounts we flagged as automated - to calculate the frequency of login activity. Rather than counting each tweet as a separate login, since some users can tweet very rapidly, we employ a refractory period of 15 minutes, during which subsequent tweets do not count as additional logins. Using this method of counting logins, we then compute the average λ per user for the entire 11 month period, resulting in the average λ of 0.3846 logins per hour.

For all of our contagion simulations, we first pick a random schedule for logins of every user, using a stationary Poisson model and the rate parameter above. We note that assuming that login activity is uniform across users and stationary with time is unrealistic. Existing research suggests that the inter-arrival times of online human activity are better modeled by a power-law distribution (Barabasi 2005) or as a non-stationary Poisson process (Malmgren et al. 2008). Assuming a stationary model ignores diurnal activity and the intermittent - and sometimes long - hiatus people may take from using social media. We use these assumptions of uniform and stationary login activity for the purpose of simplicity, acknowledging that since we are looking at only the first 24 hours of diffusion events, the inaccuracy in the model is limited to whatever diurnal effects exist over the modeled time period.

Probabilistic Adoption Models

We distinguish between two types of adoption models: *simple* and *complex*. We model user adoptions as a Bernoulli process where the likelihood of adoption after each exposure to a neighbor's use of the hashtag is a function of the number of distinct adopting neighbors. In the simple model, the adoption likelihood is constant with number of adopting neighbors (also called sources). The complex model approximates threshold behavior by using a sigmoid curve of single-exposure adoption likelihood as a function of the number of sources. We take each use of a hashtag by a followed user as an exposure event, so 10 uses of the hashtag by one neighbor translates to the same number of exposures as single uses by 10 different adopting neighbors. After each exposure, the user has a probability of transitioning to an adopted state. Once a user has adopted, they cannot revert to being a non-adopter, so this could be thought of as an SI model.

In simulation, we draw the random login schedule for all users as a first step to simulating the contagion event. Then we step through the sorted full set of login events, checking for activity that occurred since that user's last login and drawing from the probabilistic adoption model to decide if the user will adopt or not. If the user adopts, their deterministic schedule for reuses of the hashtag is added to the

exposure schedule (which already contains the schedule of seed exposures). The simulations are run for a fixed period of time; 24 hours from the first adoption time.

Simple Contagion Model Each exposure is treated as an independent Bernoulli trial. After each exposure, the user has a probability p of becoming an adopter (infected). Hence, the likelihood of having adopted after exactly x exposures follows a geometric distribution, shown in (1). The probability of being infected when a node has been exposed x times is then given by the negative binomial distribution.

$$p(A) = (1 - p)^{x-1} p \quad (1)$$

The probability of being adopting at some point in a sequence of x exposures is given by,

$$p_{\text{simple}}(A; x) = 1 - (1 - p)^{f(x)} \quad (2)$$

where $f(x)$ is a function of the number of exposures seen by the user. In this work, we use the function $f(x) = x$, but one could envision using a function where the user becomes immune to additional exposure after x becomes large enough. In simulation, the above model is accomplished by drawing independent Bernoulli trials with constant, p , after each exposure. At each point when the user logs in, we simply count the number of exposures that occurred between last login and the current time, flipping a biased coin for each exposure.

Complex Contagion Model In the complex contagion model, the user requires (or benefits from) multiple sources of influence. That is, subsequent exposures, when they are from new adopting neighbors, have increased likelihoods of influencing the user to adopt. Traditionally, complex contagion has been modeled as a deterministic threshold model. The user has a fixed threshold, k_{th} or γ_{th} , which is the number or fraction, respectively, of adopting neighbors required to trigger a user's adoption. We translated this threshold model to a probabilistic adoption model where the likelihood of adoption increases with the number of distinct adopting neighbors. In this work, we focus on the contagion model as a function of the absolute number of adopting neighbors. These models could also be adapted to depend on the fraction of adopting neighbors.

The probabilistic complex contagion adoption model that we propose is also Bernoulli for each trial, but where the success probability is nonstationary. In essence, the likelihood of adoption at each exposure is modulated by the number of adopting neighbors. We note that each exposure is evaluated in sequence, regardless of when the user last logged in (so that waiting to check one's timeline doesn't increase one's likelihood of adoption for the exact same sequence of neighbor actions). This is akin to modeling a user who logs in and scrolls through all past history in order of least to most recent.

The model needs to be aware of more than just the total number of times a user's neighbors engaged in the behavior. We also need to know how many of the neighbors were adopters at the time of each new exposure. We denote this sequence of exposures for user i at time t as a sequence of

ordered pairs,

$$X_i^t = [(t_1, 1), (t_2, 1), (t_3, 2), \dots, (t_n, k)] \quad \text{for } t_1 \leq t_2 \leq t_3 \leq \dots \leq t_n < t \quad (3)$$

where $t_1, t_2, t_3, \dots, t_n$ are n nondecreasing (i.e., ordered) timestamps. The second value of each pair is an integer representing the number of adopting neighbors. They are also nondecreasing in the sequence and constrained to be between 1 and the total number of neighbors, F . The sequence can then be summarized by a list of length k , where k is the number of distinct adopting neighbors at the end of the sequence. This summary list, L_i^t , with length k , contains the number of exposures that occurred while the user had each number of adopting neighbors, from 1 to k . In the example, the summary list for user i at time t could be,

$$L_i^t = (2, 1, \dots, 1) \quad (4)$$

indicating that there were 2 exposures when the user had 1 adopting neighbor, 1 exposure while there was 2 adopting neighbors, and 1 exposure when the user had the total length of L_i^t , equal to k , adopting neighbors. Note that each value in L_i^t must be at least 1, since there must be at least one exposure when the number of sources increases. The sum of the values of L_i^t equals the total number of exposures and is also equal to the length of the sequence X_i^t .

The probability of adopting at exactly the last exposure in the sequence L_i^t is given by,

$$p(A; L_i) = \dots \left(\prod_{k=1}^{|L_i|-1} (1 - p_k)^{L_i(k)} \right) (1 - p_{|L_i|})^{L_i(|L_i|)-1} p_{|L_i|} \quad (5)$$

where $|L_i|$ denotes the length of the sequence L_i and we have omitted the t for brevity.

As an example, if the list of exposures were $(2, 1, 3, 1)$, then the probability of adopting at exactly the last exposure would be

$$(1 - p_1)^2 (1 - p_2)^1 (1 - p_3)^3 p_4 \quad (6)$$

Similarly, the probability of adopting at some point in the sequence, L_i^t , of exposures is given by,

$$p_{\text{complex}}(A; L_i) = 1 - \left[\prod_{k=1}^{|L_i|} (1 - p_k)^{L_i(k)} \right] \quad (7)$$

We have not yet defined the likelihood of adoption p_k , which is an increasing function of k (the number of sources). To approximate threshold-like behavior, we use a modified version of the logistic sigmoid function,

$$p_k = \epsilon_{\text{lo}} + \frac{\epsilon_{\text{hi}} - \epsilon_{\text{lo}}}{1 + e^{-g(k-k_0)}} \quad (8)$$

for $k > 0$. We do not model the process of adoption when $k = 0$, instead using the true schedule of seed adoptions and hashtag uses. The 4 parameters of the sigmoid model are ϵ_{lo} , ϵ_{hi} , g , and k_0 . The ϵ parameters are simply set to 0.001 and 0.25, and represent the minimum and maximum




Figure 1: Comparison of simple versus complex model for per-exposure adoption likelihood as a function of the number of adopting neighbors. The simple model has a constant likelihood, 0.006; the complex model is a sigmoid described in (8) with $k_0 = 18$, $\epsilon_{\text{lo}} = 0.001$, $\epsilon_{\text{hi}} = 0.25$, and $g = 1.0$.

single exposure probability of adoption for all $k \geq 1$. The shape parameter g , set to 1, governs the rate of change of the sigmoid; i.e., how sharp the threshold in likelihood is. The parameter k_0 , shifts the location of the threshold in the sigmoid function, approximating the effect of the deterministic threshold in non-probabilistic complex contagion models.

Figure 1 shows the comparison of the per exposure likelihood of adoption as a function of the number of sources for the simple and complex models we are using. The simple model is constant with the number of sources while the complex model rapidly increases the adoption likelihood when the number of source approaches k_0 , here shown at 18. For the simulations in this study, we fix all parameters of the probabilistic complex contagion model to those shown in Figure 1 except for the threshold location, k_0 .

Simulation of this complex contagion model involves looking at the new exposures that have occurred since the user's last login, recording the previous number of adopting neighbors at the time of each new hashtag usage, and incrementing the number of adopting neighbors when the new usage is due to a new source. From this information, the sequence L_i is constructed and the probability of adoption is simulated by drawing from a sequence of Bernoulli trials (potentially with changing probability, p_k).

Similar to the simple model with saturation, when the number of exposures becomes large enough, we could adapt this complex contagion model such that p_k falls to 0 when k gets large enough or when the total number of exposures exceeds some threshold. We do not test models of immunity in this study, leaving that to future work.

Experiments and Results

We tested the probabilistic contagion models against the real world adoption events of 20 hashtags that were popular in Nigerian Twitter users in 2014. Our approach was to optimize the unknown parameters of the complex and simple contagion models, picking values that minimize the mean squared error of predicted adoption curves. For each of a set

Hashtag	Best-Fit Param.'s		Mean Relative ℓ_2 Err.		Log χ^2 Probability	
	k_0	p	Complex	Simple	Complex	Simple
AmericaWillKnow	18.00	0.0064	33.9%	39.2%	-31.4	-Inf
BRAvsGER	8.00	0.0159	38.6%	25.7%	-Inf	-Inf
BringBackOurDaughters	5.33	0.0119	208.5%	49.1%	-0.28	-0.30
BringBackOurGirls	5.67	0.0105	221.1%	67.3%	-16.4	-Inf
BringBackOurGirlsAlive	4.33	0.0168	67.1%	54.9%	-Inf	-Inf
BringBackOurSisters	4.67	0.0186	43.5%	19.4%	-4.7	-29.2
Ebolafacts	20.00	0.0012	14.5%	12.0%	-16.92	-16.85
Makeuptransformation	11.33	0.0127	20.1%	17.6%	-Inf	-23.1
MH17	16.00	0.0067	23.8%	32.0%	-Inf	-Inf
MH370	8.67	0.0076	17.9%	20.2%	-22.8	-11.7
MTNTcheleteGoodLife	7.00	0.0110	77.5%	28.6%	-Inf	-4.7
Nyanyablast	5.67	0.0105	33.0%	28.1%	-8.0	-7.5
RIPRobinWilliams	5.00	0.0132	15.4%	8.9%	-18.4	-17.3
SaveYakubuYusuf	23.00	0.0304	106.2%	16.8%	-15.3	-2.3
TheChibokGirls	15.33	0.0062	18.0%	18.1%	-Inf	-Inf
WeAreAllMonkeys	7.00	0.0089	31.2%	27.3%	-Inf	-30.1
WelcomeDiMaria	7.67	0.0105	37.3%	29.6%	-Inf	-Inf
WetinBeLove	12.33	0.0087	12.4%	19.4%	-35.6	-5.1
WhatJayZSaidToSolange	12.33	0.0119	24.7%	16.9%	-Inf	-4.6
YesAllWomen	12.00	0.0055	28.6%	19.2%	-20.0	-3.1

Table 2: Best-fit (in ℓ_2 error) model parameters and average relative ℓ_2 fit error for each hashtag over 100 simulations with the best fitting complex and simple model parameters. All errors are computed for 10 evenly spaced time samples from 2.18 to 24 hours. The log probability of the chi-squared (χ^2) residual is based on the mean and full covariance estimated from simulation values at the 10 sample times; where -Inf represents a probability too small to be represented in MATLAB.

of 10 random simulations, we compute the predicted adoption curve, denoted $\hat{\mathbf{a}}$, for 24 hours after the first seed adoption and compare it to the actual adoption curve, \mathbf{a} . We compute the average squared error at 10 evenly spaced times between 2.18 and 24 hours since the time of the first adoption. Table 2 shows the average ℓ_2 fit error, normalized by the ℓ_2 norm of the true adoption curve. This average relative ℓ_2 fit error is given by $E[\|\hat{\mathbf{a}} - \mathbf{a}\|_{\ell_2}]/\|\mathbf{a}\|_{\ell_2}$.

We use a two stage parameter search to find model parameters for both the complex and simple models. In the first stage, a binary search is conducted where each tested parameter value is evaluated over 10 simulations. The linearly interpolated number of simulated adoptions is compared to the true number of adoptions at 4, 8, 12, 16, and 20 hours from the first adoption time. If the majority of simulation values for all 10 simulations and sample times are above the true adoption curve, then the parameter k_0 is increased for the complex model or the parameter p is lowered for the simple model. During the binary search phase, 31 evenly spaced samples of k_0 are considered between 2 and 62, starting initially with the middle value, $k_0 = 32$. For the binary search phase when fitting the simple model, 31 evenly spaced samples of p between 0.0005 and 0.015 are searched, starting with the middle value, $p = 0.00775$. After upper and lower bounding the parameter in the binary search phase, a refined search is conducted over 5 additional evenly spaced samples of the model parameter that are between the lower and upper estimates of k_0 or p . The final estimate of the model parameter is chosen among these refined samples, including the upper and lower bounds found in the binary search, to be the parameter that minimizes average mean squared er-

ror over the 10 simulations at 10 time samples between 2.18 and 24 hours. After choosing best fitting model parameters, we generate 100 random simulations, where a new timeline checking schedule is sampled after every 10 simulations, using the best complex and simple parameters. The results of these simulations are shown in Table 2 and Figure 2.

Simulations in Python of 10 samples per model parameter, performing the two stages of search over the 24 hour period with timeline checking rates of 0.3846 logins per hour per each of the 52,689 users, currently take approximately 1 hour per model on a computer with a 2.8 GHz Intel Core i7 processor. This binary search method allows us to cover a wide range of parameter values in 5 or fewer simulation iterations, narrowing down the space before performing a local grid search over 5 additional values. While we are currently searching over exclusively the threshold parameter in simulated complex contagions, the predicted adoption curves are also highly sensitive to the other parameters, especially those governing the minimum and maximum per-exposure adoption likelihoods. We leave optimization of those parameters jointly with the threshold value to future work as this high dimensional search is challenging to efficiently perform.

Aside from the relative ℓ_2 fit error, we report the log of the probability of observing a chi-squared (χ^2) residual between the true and simulated adoption curves. The mean, μ , and full covariance matrix, Σ , of the 10 interpolated time sample values at 2.18 to 24 hours are computed over 100 simulations. Then the χ^2 value is computed as $(\mathbf{a} - \mu)^T \Sigma^{-1} (\mathbf{a} - \mu)$. We use MATLAB's chi2cdf function to compute the likelihood that the residuals should be

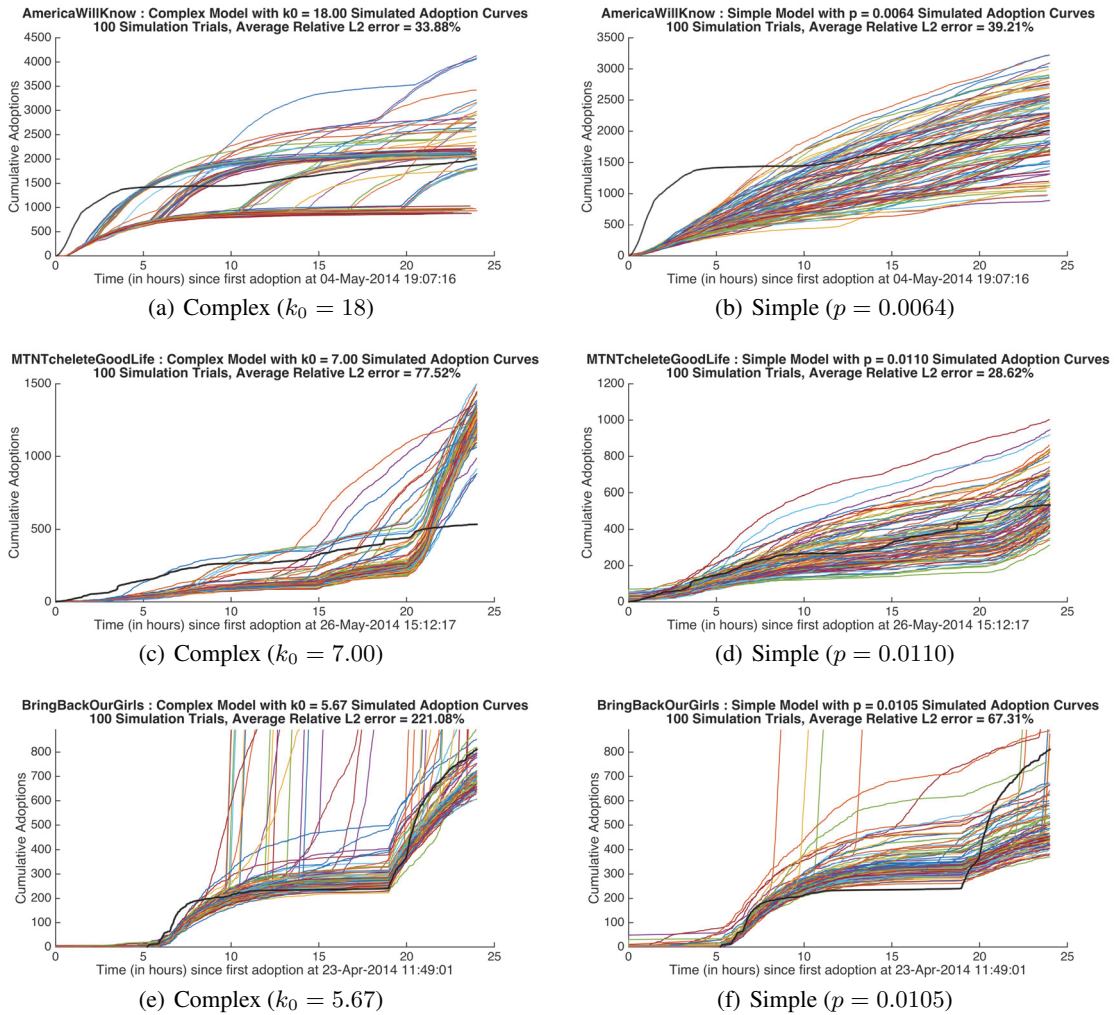


Figure 2: 100 simulated cumulative adoption curves (shown in color) using the best ℓ_2 -error fit parameters to each model versus the true adoption curve (shown in black) for #AmericaWillKnow, #MTNTcheleteGoodLife, and #BringBackOurGirls.

less than the observed value. The log of the probability of observing a residual at least as large as the one measured is shown in Table 2 for comparison between the models. The highest likelihood of any model observed was the simple model on the tag BringBackOurDaughters. However, for some hashtags, large variations in predicted adoption curves lead to the high reported chi-squared fit probabilities, as indicated by large mean relative ℓ_2 errors.

The results of Table 2 are mostly in line with expectations that the first hashtag, AmericaWillKnow, may represent a complex contagion, while many of the other tags are well-explained by simple exposure models. AmericaWillKnow has lower normalized average fit error under the complex model, while MH370 and MTNTcheleteGoodlife have lower fit error under the simple contagion model. In addition, we see the best fitting values of k_0 are higher for AmericaWillKnow than for most of the other hashtags. There are also examples of hashtags where neither of the probabilistic adoption models generated adoption curves whose variation

could explain the observed adoption curve. The tags BRAVGER and WelcomeDiMaria were examples where the adoptions were so concentrated in time that they were likely not due to a contagion phenomenon, but rather due to influence from an external event, a soccer game, which correlates neighboring user behavior. The tag MH17 had low relative ℓ_2 fit error, but adoption curves that were consistently underestimating the initial growth in adoptions. This could be related to not optimizing over the parameters ϵ_{lo} and ϵ_{hi} , or due to external influence from news sources causing adoptions that were again not the result of influence between neighbors.

Also shown in Table 2, are fit results for the important set of hashtags regarding the April 2014 kidnapping of Nigerian schoolgirls by Boko Haram: BringBackOurGirls and its variants. The models fit very poorly in an ℓ_2 error sense when the full set of 100 simulations were run. This was due to a tendency of both models, but especially the complex models, to predict a take-off event where the rate of adop-

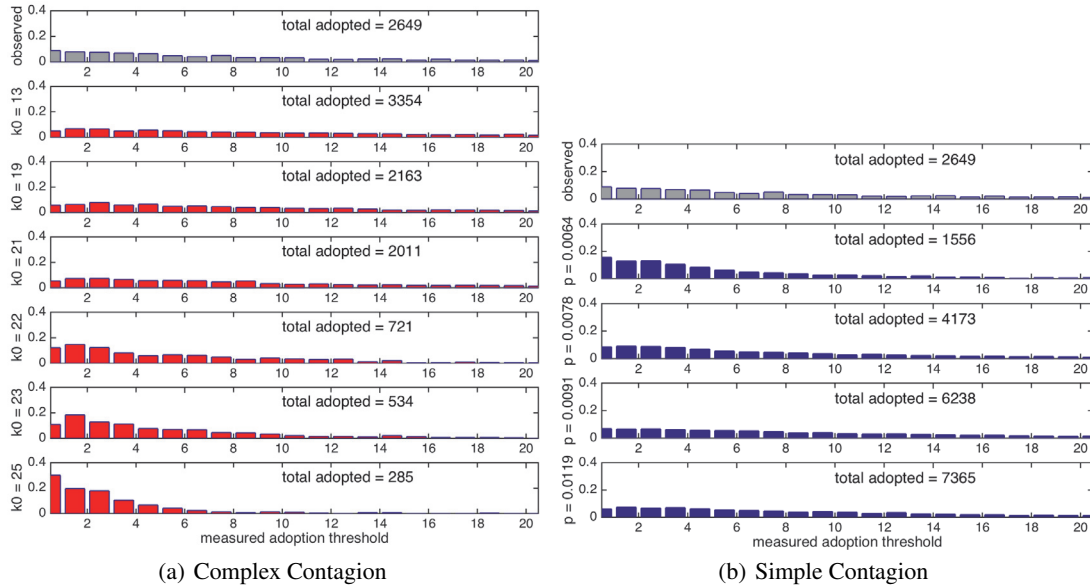


Figure 3: Comparison of the distributions of observed adoption thresholds, k , in the true #AmericaWillKnow contagion (top series) versus single-simulation measured k values for complex and simple models with varying model parameters. For complex contagion simulations, the centering parameter, k_0 was varied between 13 and 25, with other parameters fixed ($\epsilon_{lo} = 0.001$, $\epsilon_{hi} = 0.25$, $g = 1.0$). For the simple contagion, the single parameter p was varied between 0.006429 and 0.011857.




Figure 4: Comparison of simple (blue) and complex (red) simulations with the observed (black) CDF of the number of adopting friends at time of adoption, k , for #AmericaWillKnow. Axes are $CDF(k = 2)$ and $1 - CDF(k = 5)$. Three random simulations per parameter are shown.

tions jumps drastically and the contagion quickly spreads widely in the network of users. This caused wide variation among our simulated curves making analysis of the statistical chi-squared goodness-of-fit a potentially unreliable metric. Figure 2 shows the simulated adoption curves for BringBackOurGirls, zoomed in on the true curve. We see approximate 30 simulations where early take-off is predicted with the complex model and about 6 simulations where the simple model also predicts early take-off. We are investigating

the reasons behind these predictions and believe that more realistic models for fatigue or nonhomogeneous login activity may mitigate the large variability observed.

Figure 2 shows comparisons of the simulated and true adoption curves for three hashtags under each model with the best-fitting parameters. The left panels show the complex contagion simulations while the right panels show simulations using the simple model. We see, in panel (a), the complex contagion model is most capable of modeling the leveling off of adoptions observed in the AmericaWillKnow contagion, while, in panel (b), we see that the simple model predicts steady growth for this event and in this social network. In contrast, an adoption curve with steady growth does a good job of fitting MTNTcheleteGoodLife, panel (d), which is a tag related to a sales promotion. The bottom panels show fits for the tag, BringBackOurGirls, which was related to a social movement surrounding the kidnapping of the school-girls. These curves are zoomed in on the true adoption curve. There are simulations that predict drastic increases in adoption rates that are not seen in the empirical data. We are investigating the reasons behind this variability in simulation predictions, but note that in cases where early take-off did not occur, the complex model did a better job than the simple model of predicting the increase in adoptions that occurred at approximately 18 hours.

Next we examined the relationship of observed adoption threshold values, k , for just the AmericaWillKnow tag. Figure 3 shows the measured probability densities of each k threshold, from 1 to 20, compared to the uppermost distribution, which was the empirical k distribution. Each of these distributions is derived from a single random simula-

tion. Under the complex contagion models, in panel (a), we show the changes in measured k distributions with the model parameter, k_0 . We see that as k_0 increases, we actually see a drop in the average observed k threshold. We note that when k_0 increases, there are fewer adopters in general, so the likelihood that any individual will have a large number of adopting neighbors falls. Even so, for the very small values of k_0 , we hypothesize that the large upward bias in observed adoption thresholds is due to the rapid spread. Since random user logins have a constant average rate, users are more likely to get behind in checking their neighbor's activity in these cases. In panel (b) of Figure 3, we see a similar trend with increasing simple parameter, p . As p increases, the distribution of k thresholds moves further from 1, also likely to be due to the rapid spread of the contagion. This indicates that traditional methods measuring adoption thresholds to assess whether complex or simple contagion is at work would not lead to reliable detectors of complex phenomena.

Figure 4 shows the scatter plot of the cumulative density function (CDF) of users with observed thresholds 2 or lower versus users with observed thresholds greater than 5. Values are shown for 3 random simulations per complex or simple contagion parameter. The red markers for complex contagions with varying k_0 are spread throughout the entire range of blue markers associated with simple contagion. As a result, we expect there to be little hope of inferring whether an empirical contagion is complex or simple based on observed adoption thresholds. This difficulty is part of our motivation for testing observability of model type by directly fitting simulations to the adoption curve.

Conclusions and Future Work

We have proposed probabilistic contagion models for both simple and complex contagion phenomena, as well as a Poisson model for random delays in user activity. We tested these models on 20 contagion events in a 2014 Nigerian Twitter data set of over 50,000 users. By optimizing over unknown parameters of the contagion models; i.e., the threshold for the complex model and the single probability of adoption for the simple model, we compute the fit error of the empirical adoption curve to 100 random simulations under optimized model parameters. We show that for the hashtag #AmericaWillKnow, which is politically themed, the complex contagion model produces a superior fit than the simple contagion model. In contrast, for the majority of hashtags, such as #MTNTcheleteGoodLife, the simple contagion model produces lower average fit error, matching our expectations about many of those contagions.

We also investigated the empirical as well as simulation-based observations of the number of adopting friends at the time of user adoption (i.e., the complex threshold k). We found that the distribution of user thresholds over k is uncorrelated with the threshold parameter of the complex model. In addition, for wide ranges of unknown simple and complex parameters, the cumulative distributions of thresholds do not distinguish between complex and simple contagions.

This initial case study of 20 popular hashtags demonstrates the potential of these generative models for explaining the time evolution of contagion, as well as distinguishing

complex from simple phenomena. In future work, we plan to study the cause of early take-offs observed in some simulations, model varying login rates as well as nonstationarity in user activity, and test the impact of modeling immunity after a large number of exposures, as it may be necessary to reduce variability in predicted adoption curves.

Acknowledgments

This work was funded by the Minerva Initiative through the United States Air Force Office of Scientific Research (AFOSR) under grant FA9550-15-1-0036. We thank the anonymous reviewers for their helpful feedback.

References

- Barabasi, A.-L. 2005. The origin of bursts and heavy tails in human dynamics. *Nature* 435(7039):207–211.
- Barash, V.; Cameron, C.; and Macy, M. 2012. Critical phenomena in complex contagions. *Social Networks* 34(4):451 – 461.
- Barash, V. 2011. *The Dynamics Of Social Contagion*. Ph.D. Dissertation, Cornell University.
- Cameron, C. 2016. *Complex Contagion in Social Networks*. Ph.D. Dissertation, Cornell University.
- Centola, D., and Macy, M. 2007. Complex contagions and the weakness of long ties. *American Journal of Sociology* 113(3):702–734.
- Centola, D. 2010. The spread of behavior in an online social network experiment. *Science* 329(5996):1194–1197.
- Coleman, J. S.; Katz, E.; and Menzel, H. 1966. *Medical innovation: A diffusion study*. Indianapolis: The Bobbs-Merrill Company.
- Goel, S.; Anderson, A.; Hofman, J.; and Watts, D. J. 2015. The structural virality of online diffusion. *Management Science* 62(1):180–196.
- Granovetter, M. S. 1973. The strength of weak ties. *American journal of sociology* 1360–1380.
- Granovetter, M. 1978. Threshold models of collective behavior. *American Journal of Sociology* 8(6):1420–1443.
- International Telecommunications Union. 2015. Percentage of individuals using the internet.
- Lee, K.; Eoff, B. D.; and Caverlee, J. 2011. Seven months with the devils: A long-term study of content polluters on twitter. *Proc. of the 5th International AAAI Conf. on Weblogs and Social Media (ICWSM)* 185–192.
- Malmgren, R. D.; Stouffer, D. B.; Motter, A. E.; and Amaral, L. A. 2008. A poissonian explanation for heavy tails in e-mail communication. *Proc. of the National Academy of Sciences* 105(47):18153–18158.
- Odufuwa, F. 2012. Understanding what is happening in ict in nigeria. Technical report, Research ICT Africa.
- 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 the 20th Internat. Conf. on World Wide Web*, 695–704. ACM.
- State, B., and Adamic, L. 2015. The diffusion of support in an online social movement: Evidence from the adoption of equal-sign profile pictures. In *Proc. of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing*, 1741–1750. ACM.
- Valente, T. W. 1996. Social network thresholds in the diffusion of innovations. *Social Networks* 18:69–89.