

Who Says What to Whom on Twitter

Shaomei Wu*
Cornell University, USA
sw475@cornell.edu

Winter A. Mason
Yahoo! Research, NY, USA
winteram@yahoo-
inc.com

Jake M. Hofman
Yahoo! Research, NY, USA
hofman@yahoo-inc.com

Duncan J. Watts
Yahoo! Research, NY, USA
djw@yahoo-inc.com

ABSTRACT

We study several longstanding questions in media communications research, in the context of the microblogging service Twitter, regarding the production, flow, and consumption of information. To do so, we exploit a recently introduced feature of Twitter known as “lists” to distinguish between elite users—by which we mean celebrities, bloggers, and representatives of media outlets and other formal organizations—and ordinary users. Based on this classification, we find a striking concentration of attention on Twitter, in that roughly 50% of URLs consumed are generated by just 20K elite users, where the media produces the most information, but celebrities are the most followed. We also find significant homophily within categories: celebrities listen to celebrities, while bloggers listen to bloggers etc; however, bloggers in general rebroadcast more information than the other categories. Next we re-examine the classical “two-step flow” theory of communications, finding considerable support for it on Twitter. Third, we find that URLs broadcast by different categories of users or containing different types of content exhibit systematically different lifespans. And finally, we examine the attention paid by the different user categories to different news topics.

Categories and Subject Descriptors

H.1.2 [Models and Principles]: User/Machine Systems;
J.4 [Social and Behavioral Sciences]: Sociology

General Terms

two-step flow, communications, classification

Keywords

Communication networks, Twitter, information flow

1. INTRODUCTION

A longstanding objective of media communications research is encapsulated by what is known as Lasswell’s maxim:

*Part of this research was performed while the author was visiting Yahoo! Research, New York. The author was also supported by NSF grant IIS-0910664.

Copyright is held by the International World Wide Web Conference Committee (IW3C2). Distribution of these papers is limited to classroom use, and personal use by others.

WWW 2011, March 28–April 1, 2011, Hyderabad, India.
ACM 978-1-4503-0637-9/11/03.

“who says what to whom in what channel with what effect” [12], so-named for one of the pioneers of the field, Harold Lasswell. Although simple to state, Lasswell’s maxim has proven difficult to answer in the more-than 60 years since he stated it, in part because it is generally difficult to observe information flows in large populations, and in part because different channels have very different attributes and effects. As a result, theories of communications have tended to focus either on “mass” communication, defined as “one-way message transmissions from one source to a large, relatively undifferentiated and anonymous audience,” or on “interpersonal” communication, meaning a “two-way message exchange between two or more individuals.” [16].

Correspondingly, debates among communication theorists have tended to revolve around the relative importance of these two putative modes of communication. For example, whereas early theories such as the “hypodermic needle” model posited that mass media exerted direct and relatively strong effects on public opinion, mid-century researchers [13, 9, 14, 4] argued that the mass media influenced the public only indirectly, via what they called a two-step flow of communications, where the critical intermediate layer was occupied by a category of media-savvy individuals called *opinion leaders*. The resulting “limited effects” paradigm was then subsequently challenged by a new generation of researchers [6], who claimed that the real importance of the mass media lay in its ability to set the agenda of public discourse. But in recent years rising public skepticism of mass media, along with changes in media and communication technology, have tilted conventional academic wisdom once more in favor of interpersonal communication, which some identify as a “new era” of minimal effects [2].

Recent changes in technology, however, have increasingly undermined the validity of the mass vs. interpersonal dichotomy itself. On the one hand, over the past few decades mass communication has experienced a proliferation of new channels, including cable television, satellite radio, specialist book and magazine publishers, and of course an array of web-based media such as sponsored blogs, online communities, and social news sites. Correspondingly, the traditional mass audience once associated with, say, network television has fragmented into many smaller audiences, each of which increasingly selects the information to which it is exposed, and in some cases generates the information itself [15]. Meanwhile, in the opposite direction interpersonal communication has become increasingly amplified through personal blogs, email lists, and social networking sites to

afford individuals ever-larger audiences. Together, these two trends have greatly obscured the historical distinction between mass and interpersonal communications, leading some scholars to refer instead to “masspersonal” communications [16].

A striking illustration of this erosion of traditional media categories is provided by the micro-blogging platform Twitter. For example, the top ten most-followed users on Twitter are not corporations or media organizations, but individual people, mostly celebrities. Moreover, these individuals communicate directly with their millions of followers via their tweets, often managed by themselves or publicists, thus bypassing the traditional intermediation of the mass media between celebrities and fans. Next, in addition to conventional celebrities, a new class of “semi-public” individuals like bloggers, authors, journalists, and subject matter experts has come to occupy an important niche on Twitter, in some cases becoming more prominent (at least in terms of number of followers) than traditional public figures such as entertainers and elected officials. Third, in spite of these shifts away from centralized media power, media organizations—along with corporations, governments, and NGOs—all remain well represented among highly followed users, and are often extremely active. And finally, Twitter is primarily made up of many millions of users who seem to be ordinary individuals communicating with their friends and acquaintances in a manner largely consistent with traditional notions of interpersonal communication.

Twitter, therefore, represents the full spectrum of communications from personal and private to “masspersonal” to traditional mass media. Consequently it provides an interesting context in which to address Lasswell’s maxim, especially as Twitter—unlike television, radio, and print media—enables one to easily observe information flows among the members of its ecosystem. Unfortunately, however, the kinds of effects that are of most interest to communications theorists, such as changes in behavior, attitudes, etc., remain difficult to measure on Twitter. Therefore in this paper we limit our focus to the “who says what to whom” part of Laswell’s maxim.

To this end, our paper makes three main contributions:

- We introduce a method for classifying users using Twitter Lists into “elite” and “ordinary” users, further classifying elite users into one of four categories of interest—media, celebrities, organizations, and bloggers.
- We investigate the flow of information among these categories, finding that although audience attention is highly concentrated on a minority of elite users, much of the information they produce reaches the masses indirectly via a large population of intermediaries.
- We find that different categories of users emphasize different types of content, and that different content types exhibit dramatically different characteristic lifespans, ranging from less than a day to months.

The remainder of the paper proceeds as follows. In the next section, we review related work. In Section 3 we discuss our data and methods, including Section 3.3 in which we describe how we use Twitter Lists to classify users, outline two different sampling methods, and show that they deliver qualitatively similar results. In Section 4 we analyze the production of information on Twitter, particularly

who pays attention to whom. In section 4.1, we revisit the theory of the two-step flow—arguably the dominant theory of communications for much of the past 50 years—finding considerable support for the theory. In Section 5, we consider “who listens to what”, examining first who shares what kinds of media content, and second the lifespan of URLs as a function of their origin and their content. Finally, in Section 6 we conclude with a brief discussion of future work.

2. RELATED WORK

Aside from the communications literature surveyed above, a number of recent papers have examined information diffusion on Twitter. Kwak et al. [11] studied the topological features of the Twitter follower graph, concluding from the highly skewed nature of the distribution of followers and the low rate of reciprocated ties that Twitter more closely resembled an information sharing network than a social network—a conclusion that is consistent with our own view. In addition, Kwak et al. compared three different measures of influence—number of followers, page-rank, and number of retweets—finding that the ranking of the most influential users differed depending on the measure. In a similar vein, Cha et al. [3] compared three measures of influence—number of followers, number of retweets, and number of mentions—and also found that the most followed users did not necessarily score highest on the other measures. Weng et al. [17] compared number of followers and page rank with a modified page-rank measure which accounted for topic, again finding that ranking depended on the influence measure. Finally, Bakshy et al. [1] studied the distribution of retweet cascades on Twitter, finding that although users with large follower counts and past success in triggering cascades were on average more likely to trigger large cascades in the future, these features are in general poor predictors of future cascade size.

Our paper differs from this earlier work by shifting attention from the ranking of individual users in terms of various influence measures to the flow of information among different categories of users. In this sense, it is related to recent work by Crane and Sornette [5], who posited a mathematical model of social influence to account for observed temporal patterns in the popularity of YouTube videos, and also to Gomez et al [7], who studied the diffusion of information among blogs and online news sources. Here, however, our focus is on identifying specific categories of “elite” users, who we differentiate from “ordinary” users in terms of their visibility, and understanding their role in introducing information into Twitter, as well as how information originating from traditional media sources reaches the masses.

3. DATA AND METHODS

3.1 Twitter Follower Graph

In order to understand how information is transmitted on Twitter, we need to know the channels by which it flows; that is, who is following whom on Twitter. To this end, we used the follower graph studied by Kwak et al. [11], which included 42M users and 1.5B edges. This data represents a crawl of the graph seeded with all users on Twitter as observed by July 31st, 2009, and is publicly available¹. As reported by Kwak et al. [11], the follower graph is a directed

¹The data is free to download from <http://an.kaist.ac.kr/traces/WWW2010.html>

network characterized by highly skewed distributions both of in-degree (# followers) and out-degree (# “friends”, Twitter nomenclature for how many others a user follows); however, the out-degree distribution is even more skewed than the in-degree distribution. In both friend and follower distributions, for example, the median is less than 100, but the maximum # friends is several hundred thousand, while a small number of users have millions of followers. In addition, the follower graph is also characterized by extremely low reciprocity (roughly 20%)—in particular, the most-followed individuals typically do not follow many others. The Twitter follower graph, in other words, does not conform to the usual characteristics of social networks, which exhibit much higher reciprocity and far less skewed degree distributions [10], but instead resembles more the mixture of one-way mass communications and reciprocated interpersonal communications described above.

3.2 Twitter Firehose

In addition to the follower graph, we are interested in the content being shared on Twitter, and so we examined the corpus of all 5B tweets generated over a 223 day period from July 28, 2009 to March 8, 2010 using data from the Twitter “firehose,” the complete stream of all tweets². Because our objective is to understand the flow of information, it is useful for us to restrict attention to tweets containing URLs, for two reasons. First, URLs add easily identifiable tags to individual tweets, allowing us to observe when a particular piece of content is either retweeted or subsequently reintroduced by another user. And second, because URLs point to online content outside of Twitter, they provide a much richer source of variation than is possible in the typical 140 character tweet³. Finally, we note that almost all URLs broadcast on Twitter have been shortened using one of a number of URL shorteners, of which the most popular is <http://bit.ly/>. From the total of 5B tweets recorded during our observation period, therefore, we focus our attention on the subset of 260M containing bit.ly URLs; thus all subsequent counts are implicitly understood to be restricted to this content.

3.3 Twitter Lists

Our method for classifying users exploits a relatively recent feature of Twitter: Twitter Lists. Since its launch on November 2, 2009, Twitter Lists have been used extensively to group sets of users into topical or other categories, and thereby to better organize and/or filter incoming tweets. To create a Twitter List, a user provides a name (required) and description (optional) for the list, and decides whether the new list is public (anyone can view and subscribe to this list) or private (only the list creator can view or subscribe to this list). Once a list is created, the user can add/edit/delete list members. As the purpose of Twitter Lists is to help users organize users they follow, the name of the list can be considered a meaningful label for the listed users. The

²<http://dev.twitter.com/doc/get/statuses/firehose>

³Naturally, this restriction also has downsides, in particular that some users may be more likely to include URLs in their tweets than others, and thus will appear to be relatively more active and/or have more impact than if we were instead to consider all tweets. For our purposes, however, we believe that the practical advantages of the restriction outweigh the potential for bias.

classification of users can therefore effectively exploit the “wisdom of crowds” with these created lists, both in terms of their importance to the community (number of lists on which they appear), and also how they are perceived (e.g. news organization vs. celebrity, etc.).

Before describing our methods for classifying users in terms of the lists on which they appear, we emphasize that we are motivated by a particular set of substantive questions arising out of communications theory. In particular, we are interested in the relative importance of mass communications, as practiced by media and other formal organizations, masspersonal communications as practiced by celebrities and prominent bloggers, and interpersonal communications, as practiced by ordinary individuals communicating with their friends. In addition, we are interested in the relationships between these categories of users, motivated by theoretical arguments such as the theory of the two-step flow [9]. Rather than pursuing a strategy of automatic classification, therefore, our approach depends on defining and identifying certain predetermined classes of theoretical interest, where both approaches have advantages and disadvantages. In particular, we restrict our attention to four classes of what we call “elite” users: media, celebrities, organizations, and bloggers, as well as the relationships between these elite users and the much larger population of “ordinary” users.

Analytically, our approach has some disadvantages. In particular, by determining the categories of interest in advance, we reduce the possibility of discovering unanticipated categories that may be of equal or greater relevance than those we selected. Thus although we believe that for our particular purposes, the advantages of our approach—namely conceptual clarity and ease of interpretation—outweigh the disadvantages, automated classification methods remain an interesting topic for future work. Finally, in addition to these theoretically-imposed constraints, our proposed classification method must also satisfy a practical constraint—namely that the rate limits established by Twitter’s API effectively preclude crawling all lists for all Twitter users⁴. Thus we instead devised two different sampling schemes—a snowball sample and an activity sample—each with some advantages and disadvantages, discussed below.

3.3.1 Snowball sample of Twitter Lists

The first method for identifying elite users employed snowball sampling. For each category, we chose a number u_0 of seed users that were highly representative of the desired category and appeared on many category-related lists. For each of the four categories above, the following seeds were chosen:

- **Celebrities:** Barack Obama, Lady Gaga, Paris Hilton
- **Media:** CNN, New York Times
- **Organizations:** Amnesty International, World Wildlife Foundation, Yahoo! Inc., Whole Foods

⁴The Twitter API allows only 20K calls per hour, where at most 20 lists can be retrieved for each API call. Under the modest assumption of 40M users, where each user is included on at most 20 lists, this would require roughly 11 weeks. Clearly this time could be reduced by deploying multiple accounts, but it also likely underestimates the real time quite significantly, as many users appear on many more than 20 lists (e.g. Lady Gaga appears on nearly 140,000).

- **Blogs**⁵: BoingBoing, FamousBloggers, prologger, mashable. Chrisbrogan, virtuosoblogger, Gizmodo, Ileane, dragonblogger, bbrian017, hishaman, copyblogger, engadget, danielscocco, BlazingMinds, bloggersblog, TycoonBlogger, shoemoney, wchingya, extremejohn, GrowMap, kikolani, smartbloggerz, Element321, brandonacox, remarkablogger, jsinkeyst, seosmarty, NotAProBlog, kbloemendaal, JimiJones, ditesco

After reviewing the lists associated with these seeds, the following keywords were hand-selected based on (a) their representativeness of the desired categories; and (b) their lack of overlap between categories:

- **Celebrities**: star, stars, hollywood, celebs, celebrity, celebrities, celebsverified, celebrity-list, celebrities-on-twitter, celebrity-tweets
- **Media**: news, media, news-media
- **Organizations**: company, companies, organization, organisation, organizations, organisations, corporation, brands, products, charity, charities, causes, cause, ngo
- **Blogs**: blog, blogs, blogger, bloggers

Having selected the seeds and the keywords for each category, we then performed a snowball sample of the bipartite graph of users and lists (see Figure 1). For each seed, we crawled all lists on which that seed appeared. The resulting “list of lists” was then pruned to contain only the l_0 lists whose names matched at least one of the chosen keywords for that category. For instance, Lady Gaga is on lists called “faves”, “celebs”, and “celebrity”, but only the latter two lists would be kept after pruning. We then crawled all u_1 users appearing in the pruned “list of lists” (for instance, finding all users that appeared in the “celebrity” list with Lady Gaga), and then repeated these last two steps to complete the crawl. In total, 524,116 users were obtained, who appeared on 7,000,000 lists; however, many of the more prominent users appeared on lists in more than one category—for example Oprah Winfrey was frequently included in lists of “celebrity” as well as “media.” To resolve this ambiguity, we computed a user i ’s membership score in category c :

$$w_{ic} = \frac{n_{ic}}{N_c},$$

where n_{ic} is the number of lists in category c that contain user i and N_c is the total number of lists in category c . We then assigned each user to the category in which he or she had the highest membership score (i.e., belonged to the highest fraction of the category’s lists). The number of users assigned in this manner to each category is reported in Table 1.

3.3.2 Activity Sample of Twitter Lists

Although the snowball sampling method is convenient and is easily interpretable with respect to our theoretical motivation, it is also potentially biased by our particular choice of seeds. To address this concern, we also generated a sample of users based on their activity. Specifically, we crawled

⁵The blogger category required many more seeds because bloggers are in general lower profile than the seeds for the other categories

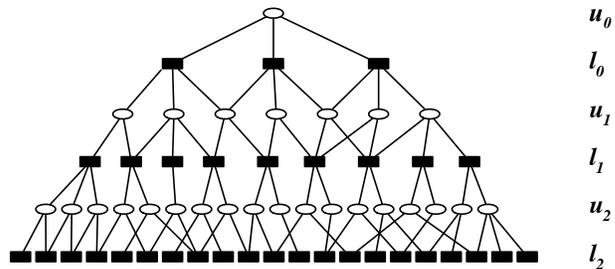


Figure 1: Schematic of the Snowball Sampling Method

Table 1: Distribution of users over categories

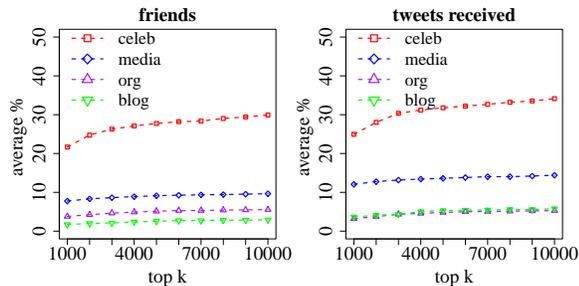
| category | Snowball Sample | | Activity Sample | |
|----------|-----------------|------------|-----------------|------------|
| | # of users | % of users | # of users | % of users |
| celeb | 82,770 | 15.8% | 14,778 | 13.0% |
| media | 216,010 | 41.2% | 40,186 | 35.3% |
| org | 97,853 | 18.7% | 14,891 | 13.1% |
| blog | 127,483 | 24.3% | 43,830 | 38.6% |
| total | 524,116 | 100% | 113,685 | 100% |

all lists associated with all users who tweeted at least once every week for our entire observation period.

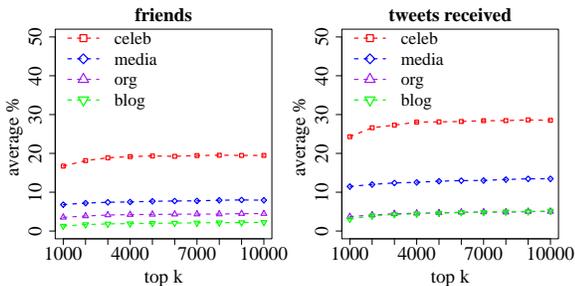
This “activity-based” sampling method is also clearly biased towards users who are consistently active. Importantly, however, the bias is likely to be quite different from any introduced by the snowball sample; despite these differences, the qualitative results that follow are similar for both samples, providing evidence that our findings are not artifacts of the sampling procedures. This method initially yielded 750k users and 5M lists; however, after pruning the lists to those that contained at least one of the keywords above, and assigning users to unique categories (as described above), we obtained a refined sample of 113,685 users, where Table 1 reports the number of users assigned to each category. We note that the number of lists obtained by the activity sampling methods is considerably smaller than that obtained by the snowball sample, and that bloggers are more heavily represented among the activity sample at the expense of the other three categories—consistent with our claim that the two methods introduce different biases. Interestingly, however, 97,614 of the activity sample, or 85%, also appear in the snowball sample, suggesting that the two sampling methods identify similar populations of elite users—as indeed we confirm in the next section.

3.3.3 Classifying Elite Users

Having classified users into the desired categories, we next refined the categories to identify “elite” users within each set. In doing so, we sought to reduce the size of each category while still accounting for a large fraction of content consumed from these categories. In addition, we fixed the four categories to be of the same size, as categories of very different sizes would require us to draw two sets of comparisons—one on the basis of total activity/impact, the other on a per-capita basis—rather than just one. To this end, we first ranked all users in each of category by how frequently they are listed in that category. Next, we measured the flow of information from the top k users in each of the four categories to a random sample of 100K ordinary (i.e. unclassified) users



(a) Snowball sample



(b) Activity sample

Figure 2: Average fraction of # following (blue line) and # tweets (red line) for a random user that are accounted for by the top K elites users crawled

in two ways: the proportion of accounts the user follows in each category, and the proportion of tweets the user received from everyone the user follows in each category.

Figures 2(a) and 2(b) show the fraction of following links (square symbols) and tweets received (diamonds) by an average user from each category, respectively. Although the numerical values differ slightly, the two sets of results are qualitatively similar. In particular, for both sampling methods, celebrities outrank all other categories, followed by the media, organizations, and bloggers. Also in both cases, the bulk of the attention is accounted for by a relatively small number of users within each category, as evidenced by the relatively flat slope of the attention curves, where we note that the curve for celebrities asymptotes more slowly than for the other three categories. Balancing the requirements described above, therefore, we chose $k = 5000$ as a cut-off for the elite categories, where all remaining users are henceforth classified as ordinary. Naturally, imposing categorical distinctions of any kind artificially transforms differences of degree (e.g. more or less prominent users) into differences of kind (“elite” vs. “ordinary”), but again we feel the interpretability gained by this distinction outweighs the costs. Moreover, because the choice of $k = 5000$ is arbitrary, we replicated our analysis with a range of values of k , finding qualitatively indistinguishable results. Thus, from this point on, we restrict our analysis to the top 5,000 users in each category identified by the snowball sampling method, noting that both methods generate similar results.

Based on this definition of elite users, Table 2 shows that although ordinary users collectively introduce by far the highest number of URLs, members of the elite categories are far more active on a per-capita basis. In particular, users classified as “media” easily outproduce all other categories,

Table 2: # of URLs initiated by category

| <i>category</i> | # of URLs | # of URLs per-capita |
|-----------------|-------------|----------------------|
| celeb | 139,058 | 27.81 |
| media | 5,119,739 | 1023.94 |
| org | 523,698 | 104.74 |
| blog | 1,360,131 | 272.03 |
| ordinary | 244,228,364 | 6.10 |

Table 3: Top 5 users in each category

| <i>Celebrity</i> | <i>Media</i> | <i>Org</i> | <i>Blog</i> |
|------------------|--------------|------------|-------------|
| aplusk | cnnbrk | google | mashable |
| ladygaga | nytimes | Starbucks | prologger |
| TheEllenShow | asahi | twitter | kibeloco |
| taylorswift13 | BreakingNews | joinred | naosalvo |
| Oprah | TIME | ollehkt | dooce |

followed by bloggers, organizations, and celebrities. Ordinary users originate on average only about 6 URLs each, compared with over 1,000 for media users. In the rest of this paper, therefore, when we talk about “celebrity”, “media”, “organization”, “blog”, we refer the top 5K users drawn from the snowball sample listed as “celebrity”, “media”, “organization”, “blog”, respectively.

Table 3, which shows the top 5 users in each of the four categories, suggests that the sampling method yields results that are consistent with our objective of identifying users who are prominent exemplars of our target categories. Among the celebrity list, for example, “aplusk,” is the handle for actor Ashton Kuser, one of the first celebrities to embrace Twitter and still one of the most followed users, while the remaining celebrity users—Lady Gaga, Ellen DeGeneres, Oprah Winfrey, and Taylor Swift, are all household names. In the media category, CNN Breaking News and the New York Times are most prominent, followed by Breaking News, Time, and Asahi, a leading Japanese daily newspaper. Among organizations, Google, Starbucks, and Twitter are obviously large and socially prominent corporations, while JoinRed is the charity organization started by Bono of U2, and ollehkt is the Twitter account for KT, formerly Korean Telecom. Finally, among the blogging category, Mashable and ProBlogger are both prominent US blogging sites, while Kibe Loco and Nao Salvo are popular blogs in Brazil, and dooce is the blog of Heather Armstrong, a widely read “mommy blogger” with over 1.5M followers.

4. “WHO LISTENS TO WHOM”

The results of the previous section provide qualified support for the conventional wisdom that audiences have become increasingly fragmented. Clearly, ordinary users on Twitter are receiving their information from many thousands of distinct sources, most of which are not traditional media organizations—even though media outlets are by far the most active users on Twitter, only about 15% of tweets received by ordinary users are received directly from the media. Equally interesting, however, is that in spite of this fragmentation, it remains the case that 20K elite users, comprising less than 0.05% of the user population, attract almost 50% of all attention within Twitter. Thus, while attention that was formerly restricted to mass media channels is now

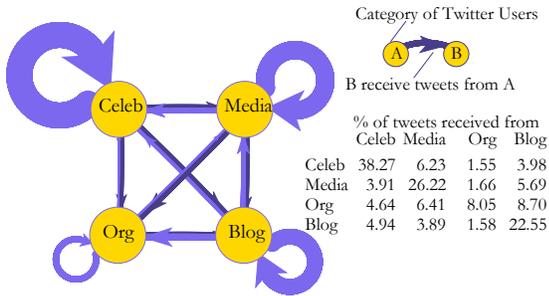


Figure 3: Share of tweets received among elite categories

shared amongst other “elites”, information flows have not become egalitarian by any means.

The prominence of elite users also raises the question of how these different categories listen to each other. To address this issue, we compute the volume of tweets exchanged between elite categories. Specifically, Figure 3 shows the average percentage of tweets that category i receives from category j (indicated by edge thickness), exhibiting noticeable homophily with respect to attention: celebrities overwhelmingly pay attention to other celebrities, media actors pay attention to other media actors, and so on. The one slight exception to this rule is that organizations pay more attention to bloggers than to themselves. In general, in fact, attention paid by organizations is more evenly distributed across categories than for any other category.

Figure 3, it should be noted, shows only how many URLs are received by category i from category j , a particularly weak measure of attention for the simple reason that many tweets go unread. A stronger measure of attention, therefore, is to consider instead only those URLs introduced by category i that are subsequently retweeted by category j . Figure 4 shows how much information originating from each category is retweeted by other categories. As with our previous measure of attention, retweeting is strongly homophilous among elite categories; however, bloggers are disproportionately responsible for retweeting URLs originated by all categories, issuing 93 retweets per person, compared to only 1.1 retweets per person for ordinary users. This result therefore reflects the conventional characterization of bloggers as recyclers and filters of information. Interestingly, however, we also note that the total number of URLs retweeted by bloggers (465k) is vastly outweighed by the number retweeted by ordinary users (46M); thus in spite of the much greater per-capita activity, their overall impact is still relatively small.

4.1 Two-Step Flow of Information

Examining information flow on Twitter also sheds new light on the theory of the two-step flow [8], arguably the theory that has most successfully captured the dueling importance of mass media and interpersonal influence. As we have already noted, on Twitter the flow of information from the media to the masses accounts for only a fraction of the total volume of information. Nevertheless, it is still a substantial fraction, so it is still interesting to ask: for the special case of information originating from media sources, what proportion is broadcast directly to the masses, and what proportion is transmitted indirectly via some population of intermedi-

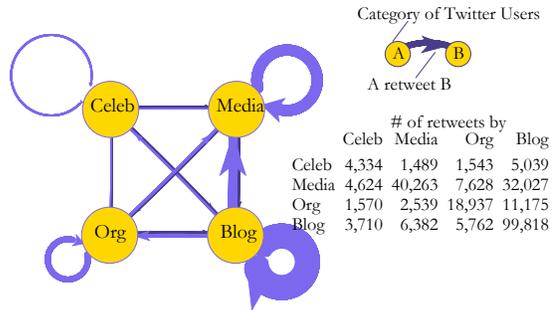


Figure 4: RT behavior among elite categories

aries? In addition, we may inquire whether these intermediaries, to the extent they exist, are drawn from other elite categories or from ordinary users, as claimed by the two-step flow theory; and if the latter, in what respects they differ from other ordinary users.

Before proceeding with this analysis, we note that there are two ways information can pass through an intermediary in Twitter. The first is via retweeting, which occurs when a user explicitly rebroadcasts a URL that he or she has received from a friend, along with an explicit acknowledgement of the source—either using the official retweet functionality provided by Twitter or by making use of an informal convention such as “RT @user” or “via @user.” Alternatively, a user may tweet a URL that has previously been posted, but without acknowledgement of a source; in this case we assume the information was independently rediscovered and label this a “reintroduction” of content. For the purposes of studying when a user receives information directly from the media or indirectly through an intermediary, we treat retweets and reintroductions equivalently. If the first occurrence of a URL in Twitter came from a media user, but a user received the URL from another source, then that source can be considered an intermediary, whether they are citing the source within Twitter by retweeting the URL, or reintroducing it, having discovered the URL outside of Twitter.

To quantify the extent to which ordinary users get their information indirectly versus directly from the media, we sampled 1M random ordinary users⁶, and for each user, counted the number n of bit.ly URLs they had received that had originated from one of our 5K media users, where of the 1M total, 600K had received at least one such URL. For each member of this 600K subset we then counted the number n_2 of these URLs that they received via non-media friends; that is, via a two-step flow. The average fraction $n_2/n = 0.46$ therefore represents the proportion of media-originated content that reaches the masses via an intermediary rather than directly. As Figure 5 shows, however, this average is somewhat misleading. In reality, the population comprises two types—those who receive essentially all of their media-originating information via two-step flows and those who receive virtually all of it directly from the media. Unsurprisingly, the former type is exposed to less total media than the latter. What is surprising, however, is that

⁶As before, performing this analysis for the entire population of over 40M ordinary users proved to be computationally unfeasible.

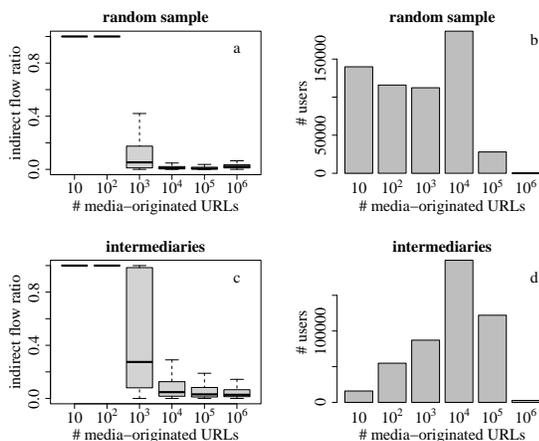


Figure 5: Percentage of information that received via an intermediary as a function of total volume of media content to which a user is exposed

even users who received up to 100 media URLs during our observation period received all of them via intermediaries.

Who are these intermediaries, and how many of them are there? In total, the population of intermediaries is smaller than that of the users who rely on them, but still surprisingly large, roughly 490K, the vast majority of which (484K, or 99%) are classified as ordinary users, not elites. To illustrate the difference, we note that whereas the top 20K elite users collectively account for nearly 50% of attention, the top 10K most-followed ordinary users account for only 5%. Moreover, Figure 5c also shows that at least some intermediaries also receive the bulk of their media content indirectly, just like other ordinary users.

Comparing Figure 5a and 5c, however, we note that intermediaries are not like other ordinary users in that they are exposed to considerably more media than randomly selected users (9165 media-originated URLs on average vs. 1377), hence the number of intermediaries who rely on two-step flows is smaller than for random users. In addition, we find that on average intermediaries have more followers than randomly sampled users (543 followers versus 34) and are also more active (180 tweets on average, versus 7). Finally, Figure 6 shows that although all intermediaries, by definition, pass along media content to at least one other user, a minority satisfies this function for multiple users, where we note that the most prominent intermediaries are disproportionately drawn from the 4% of elite users—Ashton Kucher (aplusk), for example, acts as an intermediary for over 100,000 users.

Interestingly, these results are all broadly consistent with the original conception of the two-step flow, advanced over 50 years ago, which emphasized that opinion leaders were “distributed in all occupational groups, and on every social and economic level,” corresponding to our classification of most intermediaries as ordinary [9]. The original theory also emphasized that opinion leaders, like their followers, also received at least some of their information via two-step flows, but that in general they were more exposed to the media than their followers—just as we find here. Finally, the theory predicted that opinion leadership was not a binary attribute, but rather a continuously varying one, cor-

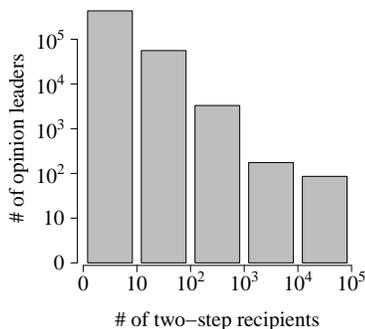


Figure 6: Frequency of intermediaries binned by # randomly sampled users to whom they transmit media content

responding to our finding that intermediaries vary widely in the number of users for whom they act as filters and transmitters of media content. Given the length of time that has elapsed since the theory of the two-step flow was articulated, and the transformational changes that have taken place in communications technology in the interim—given, in fact, that a service like Twitter was likely unimaginable at the time—it is remarkable how well the theory agrees with our observations.

5. WHO LISTENS TO WHAT?

The results in Section 4 demonstrate the “elite” users account for a substantial portion of all of the attention on Twitter, but also show clear differences in how the attention is allocated to the different elite categories. It is therefore interesting to consider what kinds of content is being shared by these categories. Given the large number of URLs in our observation period (260M), and the many different ways one can classify content (video vs. text, news vs. entertainment, political news vs. sports news, etc.), classifying even a small fraction of URLs according to content is an onerous task. Bakshy et al. [1], for example, used Amazon’s Mechanical Turk to classify a stratified sample of 1,000 URLs along a variety of dimensions; however, this method does not scale well to larger sample sizes.

Instead, we restricted attention to URLs originated by the New York Times which, with over 2.5M followers, is the most active and the second-most-followed news organization on Twitter (after CNN Breaking News). To classify NY Times content, we exploited a convenient feature of their format—namely that all NY Times URLs are classified in a consistent way by the section in which they appear (e.g. U.S., World, Sports, Science, Arts, etc.)⁷. Of the 6398 New York Times bit.ly URLs we observed, 6370 could be successfully unshortened and assigned to one of 21 categories. Of these, however, only 9 categories had more than 100 URLs during the observation period, one of which—“NY region”—was highly specific to the New York metropolitan area; thus we focused our attention on the remaining 8 topical categories. Figure 7 shows the proportion of URLs from each New York Times section retweeted or reintroduced by each category. World

⁷<http://www.nytimes.com/year/month/day/category/title.html?ref=category>

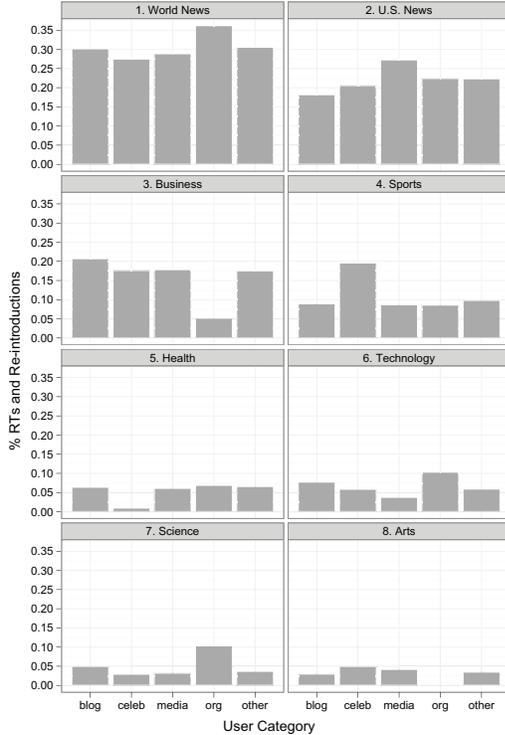


Figure 7: Number of RTs and Reintroductions of New York Times stories by content category

news is the most popular category, followed by U.S. News, Business, and Sports, where increasingly niche categories like Health, Arts, Science, and Technology are less popular still. In general, the overall pattern is replicated for all categories of users, but there are some minor deviations: in particular, organizations show disproportionately little interest in business and arts-related stories, and disproportionately high interest in science, technology, and possibly world news. Celebrities, by contrast, show greater interest in sports and less interest in health, while the media shows somewhat greater interest in U.S. news stories.

5.1 Lifespan of Content

In addition to different types of content, URLs introduced by different types of elite users or ordinary users may exhibit different lifespans, by which we mean the time lag between the first and last appearance of a given URL on Twitter.

Naively, measuring lifespan seems a trivial matter; however, a finite observation period—which results in censoring of our data—complicates this task. In other words, a URL that is last observed towards the end of the observation period may be retweeted or reintroduced after the period ends, while correspondingly, a URL that is first observed toward the beginning of the observation window may in fact have been introduced before the window began. What we observe as the lifespan of a URL, therefore, is in reality a lower bound on the lifespan. Although this limitation does not create much of a problem for short-lived URLs—which account for the vast majority of our observations—it does potentially create large biases for long lived URLs. In particular, URLs that appear towards the end of our observation period will

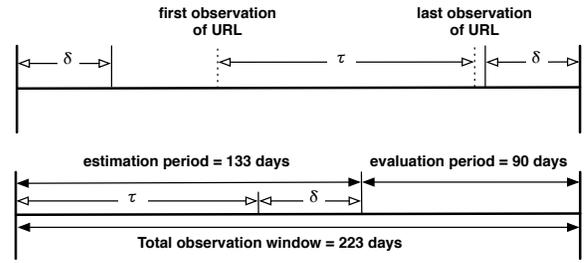


Figure 8: (a) Definition of URL lifespan τ (b) Schematic of lifespan estimation procedure

be systematically classified as shorter-lived than URLs that appear towards the beginning.

To address the censoring problem, we seek to determine a buffer δ at both the beginning and the end of our 223-day period, and only count URLs as having a lifespan of τ if (a) they do not appear in the first δ days, (b) they first appear in the interval between the buffers, and (c) they do not appear in the last δ days, as illustrated in Figure 8(a). To determine δ we first split the 223 day period into two segments—the first 133 day estimation period and the last 90 day evaluation period (see Figure 8(b))—and then ask: if we (a) observe a URL first appear in the first $(133 - \delta)$ days and (b) do not see it in the δ days prior to the onset of the evaluation period, how likely are we see it in the last 90 days? Clearly this depends on the actual lifespan of the URL, as the longer a URL lives, the more likely it will re-appear in the future. Using this estimation/evaluation split, we find an upper-bound on lifespan for which we can determine the actual lifespan with 95% accuracy as a function of δ . Finally, because we require a beginning and ending buffer, and because we can only classify a URL as having lifespan τ if it appears at least τ days before the end of our window, we need to pick τ and δ such that $\tau + 2\delta \leq 223$. We determined that $\tau = 70$ and $\delta = 70$ sufficiently satisfied our constraints; thus for the following analysis, we consider only URLs that have a lifespan $\tau \leq 70$ ⁸.

5.2 Lifespan By Category

Having established a method for estimating URL lifespan, we now explore the lifespan of URLs introduced by different categories of users, as shown in Figure 9(a). URLs initiated by the elite categories exhibit a similar distribution over lifespan to those initiated by ordinary users. As Figure 9(b) shows, however, when looking at the percentage of URLs of different lifespans initiated by each category, we see two additional results: first, URLs originated by media actors generate a large portion of short-lived URLs (especially URLs with $\tau = 0$, those that only appeared once); and second, URLs originated by bloggers are overrepresented among the longer-lived content. Both of these results can be explained by the type of content that originates from different sources: whereas news stories tend to be replaced by updates on a daily or more frequent basis, the sorts of URLs that are picked up by bloggers are of more persistent interest, and so are more likely to be retweeted or reintroduced months

⁸We also performed our analysis with different values of τ , finding very similar results; thus our conclusions are robust with respect to the details of our estimation procedure.

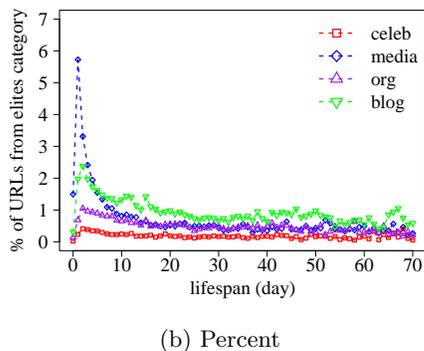
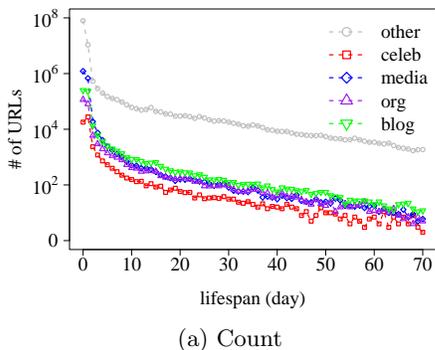


Figure 9: 9(a) Count and 9(b) percentage of URLs initiated by 4 categories, with different lifespans

or even years after their initial introduction. Twitter, in other words, should be viewed as a subset of a much larger media ecosystem in which content exists and is repeatedly rediscovered by Twitter users. Some of this content—such as daily news stories—has a relatively short period of relevance, after which a given story is unlikely to be reintroduced or rebroadcast. At the other extreme, classic music videos, movie clips, and long-format magazine articles have lifespans that are effectively unbounded, and can seemingly be rediscovered by Twitter users indefinitely without losing relevance.

To shed more light on the nature of long-lived content on Twitter, we used the bit.ly API service to unshorten 35K of the most long-lived URLs (URLs that lived at least 200 days), and mapped them into 21034 web domains. As Figure 10 shows, the population of long-lived URLs is dominated by videos, music, and consumer goods. Two related points are illustrated by Figure 11, which shows the average RT rate (the proportion of tweets containing the URL that are retweets of another tweet) of URLs with different lifespans, grouped by the categories that introduced the URL⁹. First, for ordinary users, the majority of appearances of URLs after the initial introduction derives not from retweeting, but rather from reintroduction, where this result is especially pronounced for long-lived URLs. For the vast majority of URLs on Twitter, in other words, longevity is determined not by diffusion, but by many different users independently

⁹Note here that URLs with lifespan = 0 are those URLs that only appeared once in our dataset, thus the RT rate is zero.

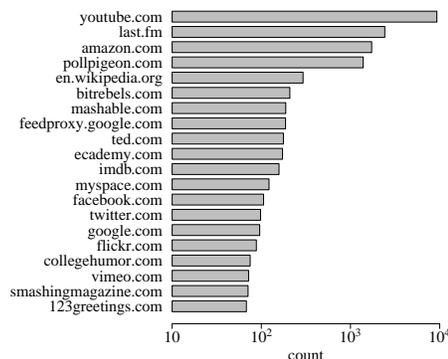


Figure 10: Top 20 domains for URLs that lived more than 200 days

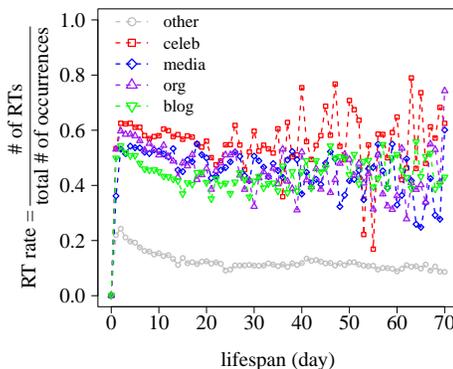


Figure 11: Average RT rate by lifespan for each of the originating categories

rediscovering the same content, consistent with our interpretation above. Second, however, for URLs introduced by elite users, the result is somewhat the opposite—that is, they are more likely to be retweeted than reintroduced, even for URLs that persist for weeks. Although it is unsurprising that elite users generate more retweets than ordinary users, the size of the difference is nevertheless striking, and suggests that in spite of the dominant result above that content lifespan is determined to a large extent by the type of content, the source of its origin also impacts its persistence, at least on average—a result that is consistent with previous findings [1].

6. CONCLUSIONS

In this paper, we investigated a classic problem in media communications research, captured by the first part of Laswell’s maxim—“who says what to whom”—in the context of Twitter. In particular, we find that although audience attention has indeed fragmented among a wider pool of content producers than classical models of mass media, attention remains highly concentrated, where roughly 0.05% of the population accounts for almost half of all posted URLs. Within this population of elite users, moreover, we find that attention is highly homophilous, with celebrities following celebri-

ties, media following media, and bloggers following bloggers. Second, we find considerable support for the two-step flow of information—almost half the information that originates from the media passes to the masses indirectly via a diffuse intermediate layer of opinion leaders, who although classified as ordinary users, are more connected and more exposed to the media than their followers. Third, we find that although all categories devote a roughly similar fraction of their attention to different categories of news (World, U.S., Business, etc), there are some differences—organizations, for example, devote a surprisingly small fraction of their attention to business-related news. We also find that different types of content exhibit very different lifespans: media-originated URLs are disproportionately represented among short-lived URLs while those originated by bloggers tend to be over-represented among long-lived URLs. Finally, we find that the longest-lived URLs are dominated by content such as videos and music, which are continually being rediscovered by Twitter users and appear to persist indefinitely.

By restricting our attention to URLs shared on Twitter, our conclusions are necessarily limited to one narrow cross-section of the media landscape. An interesting direction for future work would therefore be to apply similar methods to quantifying information flow via more traditional channels, such as TV and radio on the one hand, and interpersonal interactions on the other hand. Moreover, although our approach of defining a limited set of predetermined user-categories allowed for relatively convenient analysis and straightforward interpretation, it would be interesting to explore automatic classification schemes from which additional user categories could emerge. Finally, another two areas for future work are first, to extract content information in a more systematic manner—the “what” of Lasswell’s maxim; and second, to focus more on the effects of communication by merging the data regarding information flow on Twitter with other sources of outcome data, such as the opinions or actions of the recipients of the information.

7. REFERENCES

- [1] E. Bakshy, J. M. Hofman, W. A. Mason, and D. J. Watts. Identifying ‘influencers’ on twitter. In *Fourth ACM International Conference on Web Search and Data Mining (WSDM)*, Hong Kong, 2011. ACM.
- [2] W. L. Bennett and S. Iyengar. A new era of minimal effects? the changing foundations of political communication. *Journal of Communication*, 58(4):707–731, 2008.
- [3] M. Cha, H. Haddadi, F. Benevenuto, and K. P. Gummadi. Measuring user influence on twitter: The million follower fallacy. In *4th Int’l AAAI Conference on Weblogs and Social Media*, Washington, DC, 2010.
- [4] J. S. Coleman, E. Katz, and H. Menzel. The diffusion of an innovation among physicians. *Sociometry*, 20(4):253–270, 1957.
- [5] R. Crane and D. Sornette. Robust dynamic classes revealed by measuring the response function of a social system. *Proceedings of the National Academy of Sciences*, 105(41):15649, 2008.
- [6] T. Gitlin. Media sociology: The dominant paradigm. *Theory and Society*, 6(2):205–253, 1978.
- [7] M. Gomez Rodriguez, J. Leskovec, and A. Krause. Inferring networks of diffusion and influence. In *Proceedings of the 16th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 1019–1028. ACM, 2010.
- [8] E. Katz. The two-step flow of communication: An up-to-date report on an hypothesis. *Public Opinion Quarterly*, 21(1):61–78, 1957.
- [9] E. Katz and P. F. Lazarsfeld. *Personal influence; the part played by people in the flow of mass communications*. Free Press, Glencoe, Ill., 1955.
- [10] G. Kossinets and D. J. Watts. Empirical analysis of an evolving social network. *Science*, 311(5757):88–90, 2006.
- [11] H. Kwak, C. Lee, H. Park, and S. Moon. What is twitter, a social network or a news media? In *Proceedings of the 19th international conference on World Wide Web*, pages 591–600. ACM, 2010.
- [12] H. D. Lasswell. The structure and function of communication in society. In L. Bryson, editor, *The Communication of Ideas*, pages 117–130. University of Illinois Press, Urbana, IL, 1948.
- [13] P. F. Lazarsfeld, B. Berelson, and H. Gaudet. *The people’s choice; how the voter makes up his mind in a presidential campaign*. Columbia University Press, New York, 3rd edition, 1968.
- [14] R. K. Merton. Patterns of influence: Local and cosmopolitan influentials. In R. K. Merton, editor, *Social theory and social structure*, pages 441–474. Free Press, New York, 1968.
- [15] C. Sunstein. *Going to extremes: how like minds unite and divide*. Oxford University Press, USA, 2009.
- [16] J. B. Walther, C. T. Carr, S. S. W. Choi, D. C. DeAndrea, J. Kim, S. T. Tong, and B. Van Der Heide. Interaction of interpersonal, peer, and media influence sources online. In Z. Papacharissi, editor, *A Networked Self: Identity, Community, and Culture on Social Network Sites*, pages 17–38. Routledge, 2010.
- [17] J. Weng, E. P. Lim, J. Jiang, and Q. He. Twiterrank: finding topic-sensitive influential twitterers. In *Proceedings of the third ACM international conference on Web search and data mining*, pages 261–270. ACM, 2010.