

Information Evolution in Social Networks

Lada A. Adamic
Facebook Inc.
Menlo Park, CA
ladamic@fb.com

Thomas M. Lento
Facebook Inc.
Menlo Park, CA
tlento@fb.com

Eytan Adar
U. Michigan
Ann Arbor, MI
eadar@umich.edu

Pauline C. Ng
Genome Institute of Singapore
Singapore, Singapore
pauline.c.ng@gmail.com

ABSTRACT

Social networks readily transmit information, albeit with less than perfect fidelity. We present a large-scale measurement of this imperfect information copying mechanism by examining the dissemination and evolution of thousands of memes, collectively replicated hundreds of millions of times in the online social network Facebook. The information undergoes an evolutionary process that exhibits several regularities. A meme's mutation rate characterizes the population distribution of its variants, in accordance with the Yule process. Variants further apart in the diffusion cascade have greater edit distance, as would be expected in an iterative, imperfect replication process. Some text sequences can confer a replicative advantage; these sequences are abundant and transfer "laterally" between different memes. Subpopulations of the social network can preferentially transmit a specific variant of a meme if the variant matches their beliefs or culture. Understanding the mechanism driving change in diffusing information has important implications for how we interpret and harness the information that reaches us through our social networks.

1. INTRODUCTION

Richard Dawkins coined the word “meme” to designate ideas or messages that spread and evolve analogously to genes through communication [12]. Determining the extent to which the gene analogy applies to memes has been hampered by the lack of large-scale data containing the evolution histories of many memes, where one can pinpoint when and where information was modified and how the modifications were subsequently diffused. In this study we are able to gather precisely such data, consisting of near-complete traces of thousands of memes, collectively comprising over 460 million individual instances propagated via Facebook. Facebook is an online social network that enables individuals to communicate with their friends. The primary mechanism for this communication are *status updates* which appear in news feeds of friends. A message or idea that is appealing can quickly propagate, as those friends can choose to post the message as their own status update, thus exposing their own friends. In this way Facebook acts as a large petri dish in which memes can mutate and replicate over the

substrate of the network of friendship ties. While there are other environments where memes flourish, those memes that do enter Facebook can be examined in detail, uncovering mechanisms previously difficult—or impossible—to study.

Prior studies have either analyzed the overall popularity of memes [11,24,33,35,36,40,41], or examined a small sample of meme variants in detail [2, 7, 17, 25, 37, 38]. However, neither the large- nor small-scale studies have been able to formulate a model by which new meme variants arise. This has left a gap in our understanding of the mechanism by which social networks can not only cause information to spread [5, 10], but can also cause it to evolve, a mechanism that carries important implications for the fidelity of all socially transmitted information, as well as for specific types of information such as political campaign messages [33]. The extent to which the biological analogy of genetic evolution carries over into the evolution of memes is also unclear. We will use the terminology from genetics in describing the processes memes undergo, but defer evaluation of the biological analogy to the Discussion.

2. DATA

We analyzed de-identified data gathered over a time span of 18 months, from April 2009 to October 2011, during which many memes were propagated as textual status updates on Facebook. This period and medium represents a unique opportunity to study the evolution of memes, one which has not presented itself before or since. To understand how memes mutate naturally as they are transmitted from individual to individual, an environment is needed where individuals can freely both share information with their friends and always have the opportunity to modify the information. Prior to April 2009 character limits on status updates curtailed the length, modifiability and replication ability of memes. However, once the character limit had been increased, allowing one to not only express a complete idea or story, but also add replication instructions, e.g. “copy and paste” or “repost”, memes immediately started to flourish. Widely propagated memes would typically carry these instructions because there was yet no “share” functionality, which could at the click of a button exactly replicate the information as one’s own update. As a result, memes propagating via a manual copy and paste mechanism can be exact, or they might contain a “mutation”, an accidental or intentional modification. Since the notion of a meme was first based on the parallels between genes and ideas, we consider how textual status updates match genes in their structure and mechanics. First, they encode information, whether it is a joke, a warning, or a call to action, in a way that parallels genetic information: a string of characters is “transcribed”, allowing for some characters to be added, deleted, or substituted. Second, the replication (accurate or inaccurate) can be performed by anyone exposed to the meme. There are other forms

ACM acknowledges that this contribution was authored or co-authored by an employee, contractor or affiliate of a national government. As such, the Government retains a nonexclusive, royalty-free right to publish or reproduce this article, or to allow others to do so, for Government purposes only.

WSDM’16, February 22 - 25, 2016, San Francisco, CA, USA

© 2016 Copyright held by the owner/author(s). Publication rights licensed to ACM. ISBN 978-1-4503-3716-8/16/02...\$15.00

DOI: <http://dx.doi.org/10.1145/2835776.2835827>

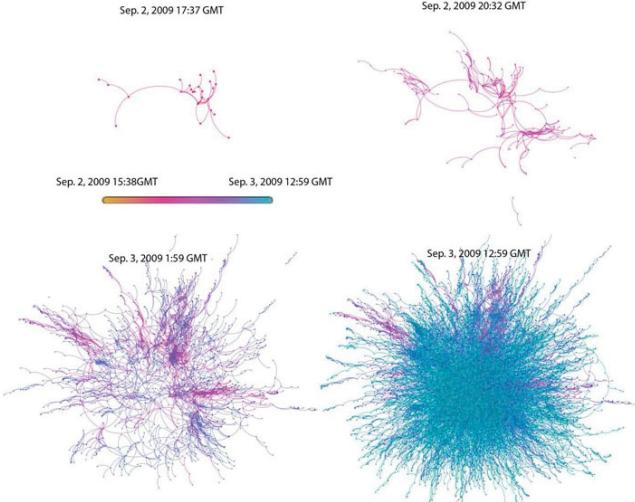


Figure 1: Early diffusion of the “no one should” meme. In a span of two days the meme diffused through many individuals and generated diffusion cascades of non-negligible depth.

in which memes occur in online environments, e.g. photographs or videos. However, the information in these is not as readily analyzed, and they cannot be as easily modified by anyone as spoken or written ideas can. We therefore focus our attention on textual status updates on Facebook to understand how information evolves when anyone can easily modify and retell it.

In order to generate a set of candidate memes, we identified status updates that had at least 100 exact copies. Nearly all such status updates contained replication instructions such as ‘copy’, ‘paste’, and ‘repost’. The few exceptions included updates generated automatically by Facebook applications and some ubiquitous memes: jokes and “wise” sayings. We therefore narrowed our scope to status updates containing replication terms such as ‘copy’, ‘paste’, etc., which included the vast majority of memes propagating via Facebook, but excluded ubiquitous text whose origin would be difficult to discern. Since the replication instructions we searched for were in English, the process captured primarily English language variants of memes.

Prior to clustering the variants into memes, we removed non-alphanumeric characters and converted the remainder to lowercase. Each distinct variant was shingled into overlapping 4-word-grams, creating a term frequency vector from the 4-grams. Sorting the meme variants by month, then by frequency, we created a new cluster, i.e. a meme, if the cosine similarity of the 4-gram vector was below 0.2 to all prior clusters. Otherwise, we added the status update to the cluster it matched most closely and adjusted the term-frequency vector of the matching cluster to incorporate the additional variant. We modified the term frequency vector of existing clusters, or created a new cluster, only if the variant frequency exceeded 100 within a month. In a post-processing step we aggregated clusters whose term vectors had converged to a unigram cosine similarity exceeding 0.4. We then gathered all variants for these 4,087 most significant memes by assigning status updates to them if their cosine similarity exceeded 0.05 using 4-grams and 0.1 using unigrams. The unigram threshold assured that an unrelated status update wasn’t erroneously included in a cluster simply for containing a relatively rare substring. The number and size of memes was insensitive to moderate adjustment of thresholds.

99.49% of posts identified as memes were made by individuals, allowing us to potentially trace the meme across the friendship graph. The remainder were made primarily by Facebook pages: group entities which enable one-to-many communication. For each meme, we induce the subgraph of directed friendship connections from people who posted the meme to their friends who posted before. We add edges between the posts of the same meme that are from the same person. We sort edges first by Levenshtein edit distance and then by time elapsed, picking the closest match as the parent in the diffusion tree. The process generates a diffusion forest rather than a tree because some origins are unknown, e.g. if a post has been deleted, or if it was propagated through a page rather than an individual’s status update. Picking the textually and temporally closest source gives us a conservative estimate of how much mutation is occurring in a meme.

We define the popularity of a variant to be the number of copies of that variant posted as a Facebook status update, and the popularity of the meme as a whole to be the sum of the popularity of all of its variants. Our sample includes some memes which never proliferated in large numbers, yet nevertheless produced at least one variant with 100 or more instances. This presented a challenge to estimating the distribution of variant popularity, especially for smaller memes of low mutation rates, where very few distinct variants are generated. We therefore appropriately limited ourselves to memes with a sufficient number of observations to yield accurate statistics. To estimate power-law exponents, which require observations over several orders of magnitude, we included memes with upwards of 100,000 variants. To estimate the required number of variants to generate accurate Gini coefficients, we simulated the Yule process and contrasted the asymptotic Gini coefficient for a meme that had evolved for a long time period, with the G during the early evolution of the meme. We found that the two values matched closely once the meme had grown to over 1,000 variants and so set this as the lower bound for the number of variants for the empirical measurements of G .

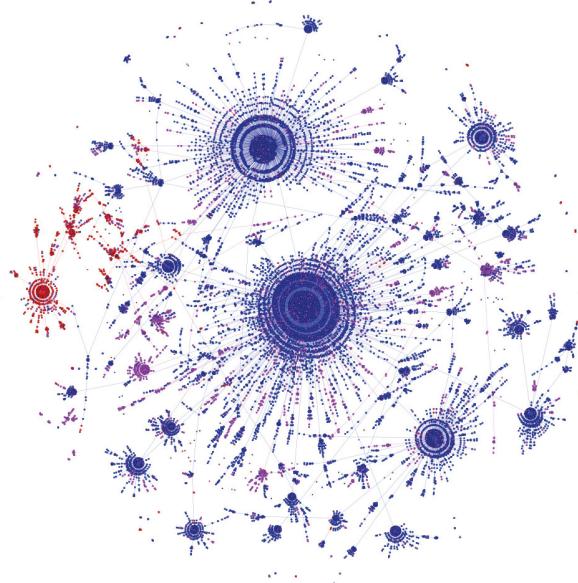


Figure 2: Approximate phylogenetic forest of the “no one should” meme. Each node is a variant, and each edge connects a variant to the most likely ancestor variant. Nodes are colored by timing prompt: rest of the day (blue), next 24 hours (red), or other (purple), showing that mutations in the timing prompt are preserved along the branches of the tree.

2.1 Mapping individuals' political leanings

To understand selection based on political inclinations of individuals in the United States, we used political affiliations specified by individuals in their Facebook profiles. Among active Facebook users in the United States in May of 2012, approximately 17.6% entered a political affiliation in their profiles. Because of the freeform nature of affiliation entry, there is a long tail of political affiliations, and 1.6 million distinct terms used, from *Democratic Party* to *Politics, n: [Poly "many" + tics "blood-sucking parasites"]*. The top 100 designations, listed in Table 1, accounted for 86.6% of all entries. We limited our analysis to those among the top 100 affiliations that could be mapped on a liberal to conservative scale, -2 being very liberal, 0 being moderate or independent, and 2 being very conservative. We excluded responses such as "Other", "I don't care", and "Libertarian (Party)". For example, we labeled the Tea Party as very conservative, and the Green Party as very liberal, the rest are specified in Table 1. These 51 designations comprise 41.2% of all people who had entered something into the political affiliation field on their profiles. Prior to December 2010, the responses were limited to a choice of (Very Liberal, Liberal, Moderate, Conservative, Very Conservative, Apathetic, Libertarian and Other). "Other" remained a popular choice and accounts for fully 41.4% of all responses. "Liberal" and "Democratic Party" accounted for 14.8%, while "Republican Party" and "Conservative" accounted for 15.5%.

Table 1: Most popular political leanings for US users in 2012 mapped on a -2 to 2 Liberal to Conservative scale.

valence	political affiliation specified
-2	Very Liberal, Liberal Democratic Party, Green Party, Peace and Freedom Party, Socialist, Communism, Communist, Socialist Party USA
-1	Democratic Party, Liberal, Barack Obama, Democratic, Democrat, OBAMA, Conservative Democrat, Progressive, Moderate Liberal, Democratic Party of the Virgin Islands, Conservative liberal, Democrat
0	Moderate, Independent, Independence Party of America, Neutral, Independent, Middle of the road, Independent Party, Moderate Party, Centrist
1	Republican Party, Conservative, Conservative Party, Republican, Independent Citizens Movement, Conservatives, Conservative Independent, Re却publican, Reform Party of the United States of America, Liberal conservative, Fiscal Conservative, Moderate conservative, NOBAMA, GOP, Liberal Republican, OBAMA SUCKS!
2	Very Conservative, Tea Party, Constitution Party, Conservative Republicans, Conservative Republican, America First Party

3. RESULTS

3.1 Replication

In this study, we are interested in memes that are replicated when a person copies a piece of text from a friend. As the process of friend copying from friend is repeated, large diffusion cascades may occur as shown in Figure 1. Typically, one replicates the meme by selecting the text of a friend's update, copying it, pasting it into one's own status update box, and sometimes editing the text further before clicking 'post'. However, the abundance of typos in some variants indicates that they were retyped character-by-character as opposed to replicated via a block copy-and-paste action. We define the mutation rate μ as the proportion of copies which introduce new edits as opposed to creating exact replicas. We further link

Table 2: Normalized text of the most common variants of the appreciation meme M1190. The pairwise Levenshtein distances are: $l(n1, n2) = 65$, $l(n2, n3) = 3$, $l(n1, n3) = 65$, $l(n1, f1) = 283$, $l(n1, c1) = 235$, $l(c1, f1) = 185$, $l(c1, c2) = 25$.

rank	copies	variant
n1	71831	somewhere right now a nurse is getting yelled at for being late with pain meds while holding her bladder because she doesn't have time to pee starving because she missed her break being pooped peed bled on and is missing her family while taking care of yours in the minute you took to read this nurses all over the world are saving lives re post this if you are a nurse love a nurse or appreciate one
f2	59352	a firefighter is being yelled at for taking too long to get there while trying desperately to save the life of a total stranger holding his bladder because he didn't have time to pee when the alarm sounded starving because he missed one of his three meals tired because the alarm sounded just as he closed his eyes it's now 4 in the morning and missing his family while taking care of yours re post if you are a firefighter love a firefighter or appreciate a firefighter
c1	26649	somewhere a cop is being yelled at for taking too long to get there trying to save the life of a total stranger holding his bladder because his mdt just went off starving because he missed a meal having to keep a stone face at an accident scene where there's 2 dead kids and now it's 4 am and he's missing his family while taking care of yours re post if you are a cop love a cop or appreciate a cop
n2	18068	somewhere right now a nurse is getting yelled at for being late with meds holding their bladder because they don't have time to pee starving because they missed lunch being peed on puked on pooped on bled on and is missing their family while taking care of yours in the minute it took you to read this nurses all over the world are saving lives repost if you are a nurse love a nurse or appreciate a nurse
n3	17954	somewhere right now a nurse is getting yelled at for being late with meds holding their bladder because they don't have time to pee starving because they missed lunch being peed on puked on pooped on bled on and is missing their family while taking care of yours in the minute it took you to read this nurses all over the world are saving lives repost if you are a nurse love a nurse or appreciate nurses
c2	15175	somewhere a cop is being yelled at for taking too long to get there trying to save the life of a total stranger holding his bladder because his mdt just went off starving because he missed a meal having to keep a stone face at an accident scene where there's 2 dead kids and now it's 4 am and he's missing his family while taking care of yours re post if you are a cop love a cop or appreciate a cop you all know who you are

each new variant to the variant it was most likely derived from, as described below.

To see concretely how mutation occurs we consider the meme "No one should die because they cannot afford health care and no one should go broke because they get sick. If you agree please post this as your status for the rest of the day." In this form, which emerged early on in the meme's evolution, it was copied over 470,000 times. A variant prepending "thinks that" (which would follow the individual's name), was copied 60,000 times. The third most popular variant inserted "We are only as strong as the weakest among us" in the middle. Some copies changed "the rest of the day" to "the next 24 hours." Figure 2 shows how this particular timing trait

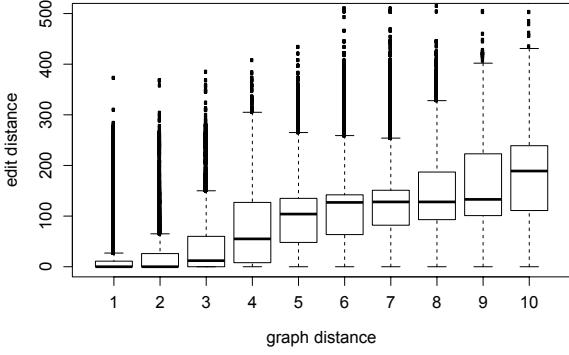


Figure 3: Edit distance from the root typically increases as a node is further down in the copy chain, i.e. the greater the number of hops in the diffusion tree. Here edit distance is shown for a sample of memes where copy chains reached a length of 10 or greater.

is passed down through the genetic lineage. The high number of exact copies is consistent with this meme having a modest and typical mutation rate $\mu = 0.11$. That is, 89% of the copies were exact, while 11% introduced a mutation.

In order to trace the lineage of particular variants, we use the edit distance between the variant an individual posted and the variants previously posted by that person’s friends. We define the edit distance as the number of character additions and deletions that must be performed in order to obtain one variant of the meme from another. For example, changing the phrase “thinks that” to “agrees that” creates two variants separated by edit distance 10, since the 5 characters ‘t h i n k’ are deleted and ‘a g r e e’ are added. Further examples are given in Table 2. As might be expected, there is a compound effect of such mutations over several generations, with edit distance increasing as copy is made from copy is made from copy (Figure 3).

3.2 Population Genetics of Memes

The first striking pattern in the data is the uneven popularity of the variants, shown in Figure 4. For 121 of the 123 largest memes with over 100,000 distinct variants each, a Kolmogorov-Smirnov test confirmed a power-law fit ($D < 0.05$), with maximum likelihood-fitted exponents of 2.01 ± 0.15 . Expanding to the 435 memes with an excess of 10,000 variants yielded similar exponents: 1.99 ± 0.21 .

Although a variety of processes can produce power laws [4, 31], we show that the evolution of memes in an online environment can be modeled by the Yule process [42]. This simple process, previously used to model the number of species per genus, bacterial populations [27] and protein domain family sizes [23, 32], contains just two simple components: replication and mutation. The process starts with a single variant of a meme. Each individual instance of the meme has an equal probability of generating a new copy per unit time. The probability of a copy containing a mutation is μ , and we let $r = \frac{\mu}{1-\mu}$ be the ratio between the probability of a mutated and non-mutated copy.

For very large times, the probability that the number of copies of a variant exceeds y is given by [42]:

$$Pr(Y \geq y) = r \frac{\Gamma(1+r)\Gamma(y)}{\Gamma(y+1+r)} \quad (1)$$

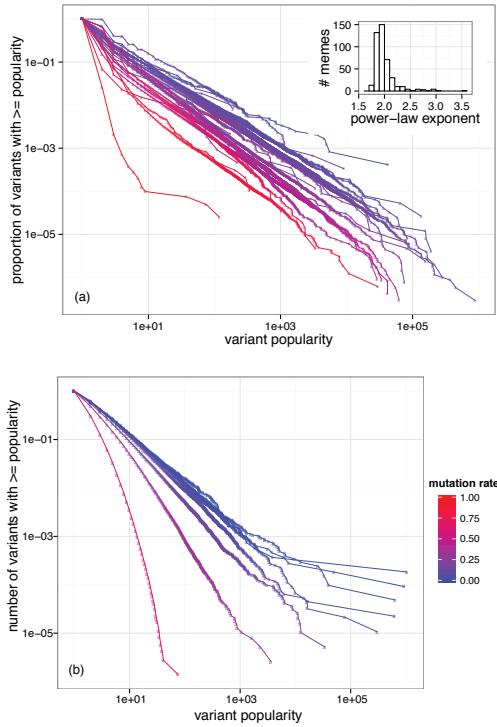


Figure 4: Distributions of variant frequency in (a) a sample of observed memes, (b) a simulated Yule process. The inset shows the distribution of fitted power-law exponents for all memes with $\geq 10,000$ variants.

where Γ is the gamma function. As the mutation rate $\mu \rightarrow 0$, $r \rightarrow 0$. This simplifies to $Pr(Y = y) = \frac{1}{y(y+1)}$, i.e. a power-law distribution with exponent 2. For higher mutation rates, the tail of the distribution, as $y \rightarrow \infty$, is given by

$$Pr(Y \geq y) \sim r\Gamma(1+r)y^{-(1+r)} \quad (2)$$

Thus the Yule model predicts that memes with a low mutation rate will have variants distributed according to a power-law distribution with an exponent close to 2, which is what we observed with our data (Figure 4(a), inset). The Yule model also predicts that memes with a high mutation rate will deviate from a power law because frequent mutation prevents any single variant from achieving an extremely high number of identical copies. Simulations of the Yule model for a range of mutation rates produce variant popularity distributions, shown in Figure 4(b), which bear a close resemblance to the empirically observed distributions in Figure 4(a).

The observed power-law distributions for low mutation rates (and their absence for high mutation rates), make the Yule process a plausible mechanism of variant creation and replication. However, many different mechanisms can produce similar distributions, and we want to validate the model predictions further, even for less popular memes where the data is insufficient to confirm a power-law fit. We use the Gini coefficient G to measure the inequality in variant frequency within a meme and to compare it against the G predicted by the Yule model for μ observed for that meme. $G = 0$ if all variants are equally popular (e.g., if every copy is a mutation), and $G = 1$ if only one variant is present (i.e., $\mu = 0$). If mutation can occur ($\mu > 0$), but none of the mutated variants are able to produce additional copies, then $G = 1 - \mu$. For a power-law distribution where $Pr(Y \geq y) \sim y^{-\alpha}$, $G = \frac{1}{2\alpha-1}$ [29]. Using Eq. 2, we

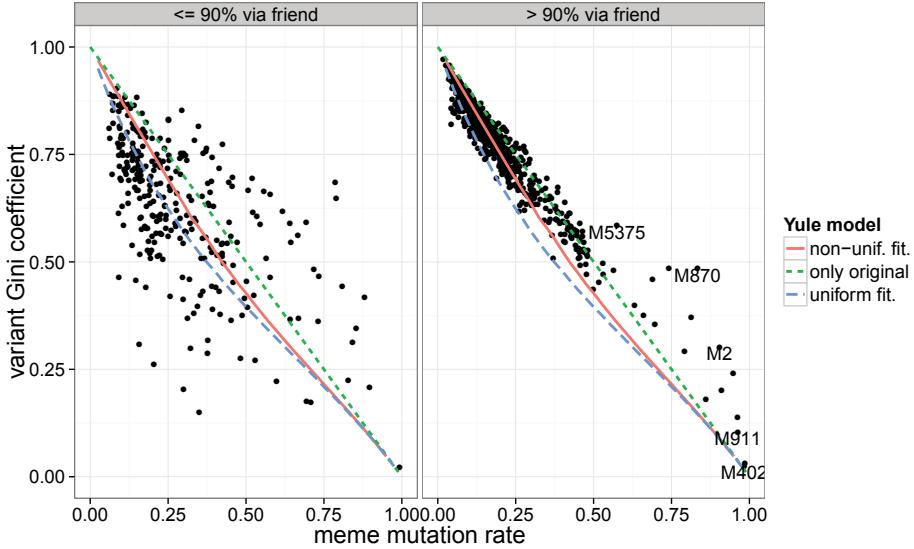


Figure 5: Correspondence between the Gini coefficient G and the mutation rate for memes with over 1000 distinct variants: (a) 287 memes with $\leq 90\%$ of instances being posted following a friend’s post, and (b) 876 memes with $> 90\%$ instances posted after a friend. $1 - \mu$ is the expected relationship if none of the mutations but the original are viable. Two simulations, with uniform and variable fitness, are shown for comparison. The identified memes are: M2 (copy sentence from p. 56 of closest book), M402 (child’s birthweight in honor of mother’s day), M870 (place of birth), M911 (first concert you attended), M5375 (zodiac sign).

can derive the correspondence $G = \frac{1-\mu}{1+\mu}$, valid for μ close to 0 in the Yule model. This allows us to form a direct prediction from the Yule model of the expected G for a given μ . Since the theoretical prediction holds only for small μ and uniform fitness, we also simulate the Yule process with both uniform and normally distributed variant fitness, to derive relationships between G and μ over the full range $0 < \mu < 1$, as shown in Figure 5.

Empirically, we need to be able to determine whether an individual copy is a mutation or not. We do this by connecting meme instances through the friendship graph. If someone posts a new variant distinct from those previously posted by their friends, the instance is counted as a mutation. If they post an exact copy of one of their friends’ previous posts, then it is counted as a copy. Figure 5 shows a close correspondence ($\rho = -0.97$) between the mutation rate and G , for memes with over 1,000 variants and over 90% of instances diffusing through the social network, i.e. having at least 90% of posts occurring after a friend’s post. The curves delineate theoretical and simulated predictions while each point represents an entire meme, with the vast majority of memes falling between these two curves. This lends additional support that meme propagation follows the Yule model. Note that the correspondence between mutation rate and Gini coefficient is no coincidence and depends crucially on the ability to measure the mutation rate precisely, as demonstrated by the contrast between Figures 5(a) and (b). If fewer than 50% of a meme’s posts follow friends’ posts of the same meme, it is likely that meme replication is being driven at least partly outside of the social network. This adds uncertainty to the replication path and μ , reducing the correlation between μ and G significantly ($\rho = -0.60$).

Only a handful of memes appeared to mutate at rates above $\mu > 0.5$. Such rates are remarkably high, with more than one out of every two copies resulting in mutation. Indeed, unlike the vast majority of copy and paste memes which simply instruct the reader to replicate them, memes with what appeared to be very high mutation rates specifically instructed the person transmitting them to

modify the text of the meme, typically by adding custom information: M402 and M4416 (see [1] for meme text) asked for the birthweights of one’s children, M6265 for oldest friends, and M431 for billboard charts from one’s birthday. The replication mechanics of these *customizable* memes are different. Rather than copying the entire text of the meme, a portion of the text is copied, while the remainder can be changed. In order to analogously treat customizable memes one would need to identify and remove the substitutable text, which is beyond the scope of our method.

Even though the replication mechanics are different for customizable memes, one might still be able to make a prediction about the distribution of copies among variants based only on the rate of change of the entire text, including customization. For example, one might assume that the main variant, containing the replication instructions, but no customization, would be the only variant with non-zero fitness and hence multiple copies, since any customization would make the variant irrelevant to the next person copying the text. In that case the prediction would be $G = 1 - \mu$. While several memes fall roughly on this line, customizable memes typically have higher Gini coefficient than expected, meaning that even though these memes prompt the text to be modified, these modifications tend to coincidentally produce the same variant. This is especially pronounced for memes where such coincidences are more likely to happen. For example, M870 asks its host to copy it, and add their place of birth. This gives it a very high mutation rate, but also a high G , related to the power-law distribution of population center sizes [15,34], i.e. individuals are likely to independently generate identical strings even while mutating the meme. The same elevated Gini coefficient due to many identical customizations occurs to a lesser extent in memes asking the host to copy a sentence from a book (M2), list their zodiac sign (M5375), and mention the first concert they attended (M911). On the other hand, the meme asking for children’s names and birthweights (M402), has both a μ close to 1, and G close to 0. Therefore, while the Yule model is a good fit for the great majority memes which are ideas evolving

“naturally” under low mutation rates, it is not appropriate for the handful of textual games inducing high mutation rates and coincidental copies.

Finally, we examine whether mutation rate affects the successful spread of a meme, but find no overall effect of mutation rate on the total number of copies of the meme ($\rho(N_{meme}, \mu_{meme}) < 0.1$). We further find that meme mutation rates are constant over time ($\rho(t, \mu_t) < 0.04$), and that the most popular variants are just as likely to generate mutated copies as less popular ones ($\rho(\text{num mutated copies}, \text{popularity}) > .95$), indicating that a meme does not converge to an optimal consensus sequence.

3.3 Mutation characteristics

While the Yule model explains the overall distributions of different variants, a more detailed look reveals several interesting patterns in the mechanics of how new variants are created: where edits occur, whether the length matters, and whether an entire meme or just parts of it are inserted into other memes.

First, there is a non-uniform probability of mutation along the length of the meme (see Figure 6). This is likely in part due to the copy and paste mechanism for replication, where most people copy and paste the text rather than typing it anew. After pasting, they might prepend, append, or modify the text. To identify the locations where text is most likely to be mutated, we segmented each variant into overlapping sequences of 4 words (4-grams). For each 4-gram within the most popular variant of a meme, we computed the proportion of other variants that contained it. On average, a 4-gram occurring in the middle of the most popular variant is preserved in 70.2% of other variants. This falls to 59.3% for the first 4-gram, and 50.3% for the last. Edits at the margins are likely to occur because the selection of text to be copied is incomplete, e.g. it misses the start or end. Similarly, pre/post-paste modifications will be at the beginning/end of the meme, unless the person moves the cursor.

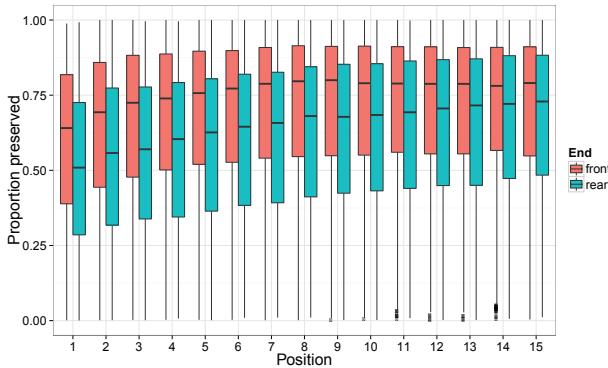


Figure 6: Probability that a segment of a meme is preserved as a function of its location within the most popular variant of that meme.

If memes are being replicated through a simple mouse swipe, one would expect a weak selection pressure for shorter variants, reflective primarily of the additional effort of reading and copying a longer meme. Figure 7 shows just such a weak within-meme selection pressure toward slightly shorter variants. Variants that are too short, and potentially have lost some of the message of the meme, are again not as popular. Consistent with the selection pressure being weak, we observe a high variation both in the average string length between memes (mean length = 356.7 ± 303.95), and length of individual variants within a meme (average within-meme length difference = 69.90 ± 195.85).

Many people are exposed to and even participate in replications of several memes. This gives them the opportunity to transfer textual sequences from one meme to another. Just four 4-grams occurred in over 200 memes with a prevalence of between 5 and 95% of the variants. They contained individual words which were initially used to identify memes in the data, but arranged them in a specific way: “put this as your”, “copy and paste this”, “this as your status”, “re[-]post if you”. We computed the ratio r_g of the average frequency of all variants containing a given 4-gram against the average frequency of variants of the same meme that do not have the exact 4-gram. For the four most widespread 4-grams, $r_g > 1.4$, meaning that even though all variants contained replication keywords, these *specific* formulations of replication instructions were advantageous.

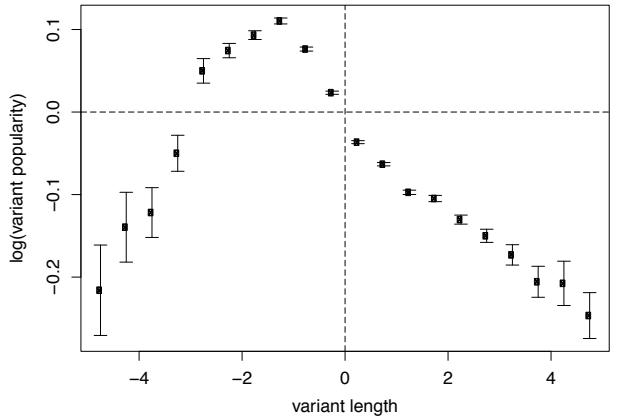


Figure 7: Relative frequency of a variant within a meme as a function of its length relative to other variants.

In Table 3 we list the top 4-grams with the highest r_g and occurring in over 20 memes. As mentioned, although all strings clustered into memes had some substring associated with replication, e.g. “post” or “copy”, the most successful variants put these together in clear replication instructions, e.g. “copy and paste”, “as your status”, “into your status”, which can facilitate the meme’s expression and transmission [18]. Other successful substrings include encouragement and allusions to competition (“see how many people”), persistence cues (“status for at least”), or conditions that are easy to match or identify with (“if you love your”, “if you know someone”, “paste if you agree”, “proud to be a”). A specific pattern, “of you won’t”, occurred in prompts such as “95% of you won’t copy this, but the 5% who [have a positive attribute] will”. 144 memes contained at least one variant matching “won’t [...] will”. These variants had significantly higher likelihood of being copied, 10.98 copies on average, relative to an average of 7.05 overall. In 102 of the 144 cases the individual who introduced the phrasing into the meme had either used it previously in another meme, or had a friend who had, suggesting that in a large number of cases the substring could have been transferred from one meme to another. On the other hand, some copy errors, occurring in significant numbers, e.g. “ago comment like unlike write”, which indicate spurious selection of Facebook boilerplate text below the update, enjoyed only one fourth of the popularity of variants which did not contain them, possibly indicating a notion of *grammatical fitness*. This is further supported by a mild negative correlation between the within-meme relative number of English spelling errors s for a variant and its popularity y ($\rho(s, \log(y)) = -0.133, p < 10^{-15}$).

Sometimes, although infrequently, more than a short substring will be shared between separate memes. This occurs when one

Table 3: 4-grams which conferred replicative advantage on the variants which contained them, within memes where they represented between 5 and 95% of the population.

4 gram	r_g	# memes
plain copy and paste		
please post this as	2.59	23
it as your status	2.23	28
copy and paste and	2.08	20
to copy and paste	2.04	27
this into your status	2.03	29
condition/identification/agreement		
paste if you agree	2.84	25
and paste if you	2.02	35
if you love your	2.00	21
status if you know	1.98	20
your status to show	1.94	28
status to show your	1.92	20
proud to be a	1.79	20
post if you agree	1.79	67
status if you are	1.79	33
if you know someone	1.78	24
other		
see how many people	2.85	25
status for at least	1.99	29

meme mutates so drastically that it becomes a separate meme, but also when memes recombine at later points, such as when two separate memes are copied and pasted together, thereby merging into a single fused meme. One such example is a variant of an anti-bullying meme (M4038A) with over 300,000 copies, which at one point was pasted together with a meme about a sick child (M4038B). The hybrid, containing most of the text of the two original memes, was exactly copied over 10,000 times. Similarly, these fused memes contain both messages of the original meme. The combination might not occur for months after the memes are individually introduced. For example, M954, a chain letter meme [20], was present in fragments since November 2008. The status update character limit imposed at that time did not permit it to be copied entirely, and the fragments posted were likely transferred from copies spreading via email. M26, a similar but distinct chain letter, first appeared in April of 2009. They combined in an awkward copy-and-paste two years later, in October 2011, with over 300 copies made of the combination, including attempts at correcting the copy error in the transition from one meme to the other.

3.4 Adaptation of Memes to Niches

So far we have presented aggregate characteristics of meme evolution. Examining the evolution of a meme in detail can give insight into how a meme adapts to specific niches within the social network environment. Prior work on cultural transmission has examined how language [22] and music [26] evolve in experimental settings. These lacked the social network structure where different variants are adopted by populations having different preferences. The niche can be time or location-specific; A meme asking mothers to describe their childrens' features, in some variants was prepended by "In honor of mother's day" (M4416) and an Amber alert (M246) appeared with the same make of car and license plate but different locations where the child was abducted: Edmonton, KY, Edmonton, Canada, Quebec, "Washington", and sometimes multiple locations simultaneously.

Here we explore in more detail a meme—"No one should die because they can't afford health insurance..."—that initially appeared

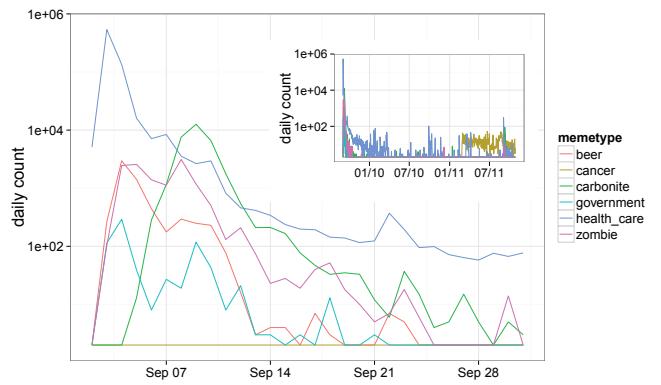


Figure 8: Different variants of the "no one should..." meme (see Table 4) peak in popularity at different times. The inset shows that the meme persisted in low numbers for two years.

as a political meme but then its variants evolved to appeal to other demographics, including the opposing political party. The meme goes on to say "no one should go broke because they get sick", which is something President Obama stated in his weekly address on August 14th, 2009, shortly before the meme appeared, and again when addressing the US Congress on healthcare reform on Sept. 9, 2009, when the meme had already taken hold. The meme would appeal primarily to liberals, which we were able to verify by computing the average political leaning for those people in the US who both filled this field out in their profile and posted the meme. Although the most prominent variants of the meme propagated primarily from liberal to liberal, some (humorous) variants appealed across the political spectrum, e.g. "no one should be without a beer because they cannot afford one...". Humor has been noted to play an important role in online discussions [6], and political satire and memes [39]. Other variants, criticizing the bill (e.g., "no one should die because the government is involved with health care...") appealed primarily to conservatives (see Table 4). These secondary variants peaked at different times (see Figure 8).

Once a person posts a meme, they are unlikely to repost it again. That is, exposure to one variant of a meme usually confers immunity to others. However, a small fraction of people post a meme more than once, sometimes years apart, as they identify that the message is still important. In the case of the "no one should" meme, 4.36% of posts were repeat posts by the same individual. Of those posting the meme a second time with a non-healthcare related variant, 40.4% had posted a healthcare variant previously. Another, more subtle, interaction is that exposure to one variant of a meme can make an individual more susceptible to other variants parodying the original variant. This is evident for the health care meme, where among the 61,979 people posting one of the non-healthcare variants, a full 93.0% had at least one friend who had previously posted a healthcare variant. This collaborative success of memes has also been observed in a population of image memes [11].

3.5 Network diffusion dynamics

Just as different biological species experience population booms and declines, memes experience bursts of activity with different variants being dominant in different periods. Figure 8 shows that the above mentioned variants of a single meme occurred proximately in time, but peaked at different points. Variants of the appreciation meme, listed in Table 2, peaked weeks apart. Simulations of a meme diffusing and mutating according to the Yule

Table 4: Average political bias of individuals sharing variants of the ‘no one should’ meme (-2: very liberal, +2: very conservative).

keyword	average bias	example
health care	-0.87	die because they cannot afford health care...
carbonite	-0.37	be frozen in carbonite because... they couldn't pay Jabba the Hut
zombie	-0.30	die because of zombies if they cannot afford a shotgun...
cancer	-0.02	have to worry about dying tomorrow but cancer patients do...
wine	0.15	go thirsty because they cannot afford wine...
beer	0.22	be without a beer because they cannot afford one...
government	0.88	die because the government is involved with health care...
Obamacare	0.96	die because Obamacare rations their healthcare...
taxes	0.97	go broke because government taxes and spends...

process in a fixed and evenly mixed population, excluding network topology, can easily reproduce the distributions of variants within a meme (see Figure 4(b)), but not the bursts in popularity of individual variants. The ingredients missing are network diffusion [3, 16, 35, 41], combined with network growth and variable user activity levels [19].

In order to simulate meme diffusion in a realistic, yet tractable network environment, we selected a small country with a large proportion of the population being Facebook users, and over 2 million having 100 or more Facebook friends. We used this actual friendship network within the country for the memes to diffuse over. In the simulation, we infected a single individual with one variant of a meme. The simulation proceeds as follows. Every friend of an infected individual is flagged as having the potential to be infected the next time they log in. Login frequencies vary between users, with each user being assigned a uniform (0,1) random variable (0.1 corresponds to logging in every 10th day on average, 1.0 to logging in every day). Once users log in, they have a probability of becoming infected that decays exponentially with each subsequent login since their first exposure. This is intended to capture both the meme being pushed down further in their Facebook feed, and also the fact that someone who chooses not to post at first exposure is less likely to post at subsequent exposures.

Such a model typically produces a single-peaked distribution in overall meme popularity (see Figure 9(a)), with different variants peaking at different times. The peak widens when individuals join the network while the meme is diffusing, but the distribution is still single-peaked Figure 9(b). Interestingly, at the point at which only several tens of thousands of status updates are posted, 1 million users, constituting over half of the network, are potentially exposed because their friends have posted. This leaves little room for untouched areas of the network that could contribute to secondary peaks.

Finally, when we further modify the simulation such that variants have variable fitness (mean=1, sd=0.5), we observe multiple peaks (see Figure 9(c)). As mentioned, memes’ fitness can vary due to the particular sequence they carry and how that message aligns with susceptibilities in the underlying network transmitting them. Variants

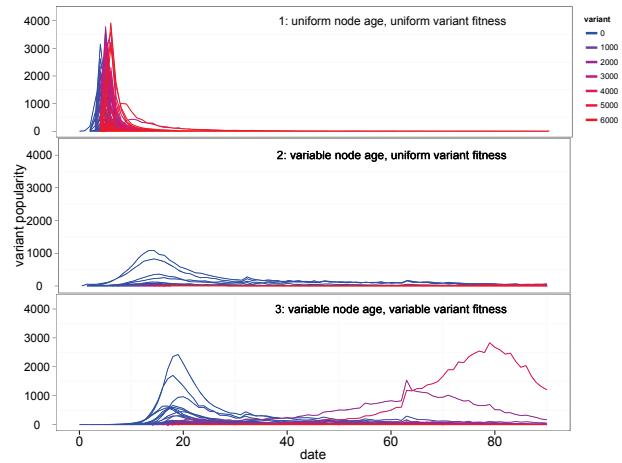


Figure 9: Simulated popularity of variants in a network simulation model, under different conditions.

in fitness can also be due to external factors, e.g. versions of an anti-bullying meme mentioning that it’s anti-bullying week, or all similar versions receiving a boost when Facebook altered its feed-ranking algorithm to cluster status updates containing similar text and placed them together near the top of a user’s feed. Thus a Yule process incorporating variable fitness, and overlaid on a growing network of variably active nodes, can explain the observed bursty meme variant popularity.

4. DISCUSSION

There are convenient analogies between the language of genes and biological evolution that are helpful to us in discussing and analyzing memes. At the most basic level, memes contain two key ingredients that are reminiscent of biological populations: replication and mutation [13]. When Richard Dawkins introduced the concept of a meme, he suggested that “memes should be regarded as living structures, not just metaphorically but technically” (Richard Dawkins, Selfish Gene p.192.) In this study, we apply this concept more rigorously to Facebook memes and find similarities between genes and memes. In biology, a gene’s genotype is the DNA sequence, and the gene’s phenotype is the resulting function that the gene carries out. For textual memes, the string is the information that is being passed, or the ‘genotype.’ The ‘phenotype’ is what is expressed by the meme, which can include the meme’s message and replication instructions.

We define a meme according to biological terms. A gene is a hereditary unit that is transmitted to offspring. Thus, since the meme is transmitted when copied and pasted, the meme itself can be considered to be a gene. Genotype is the description of variants within the gene. For example, a bacterium with a mutation in the Lac gene that knocks out Lac’s gene function is described to have the genotype Lac-. Genotypic variants can be simple text edits that don’t change the meme’s meaning (e.g. a misspelling), or they can have larger effects such as a change in meaning (“healthcare” versus “beer”). The genotype is a representation of differences, for example, “post for 24 hours” versus “post for a day” are two possible genotypes for a particular meme. The genotype is independent of any effects of the environment and is heritable [9].

Phenotypes are observable characteristics which include morphology and behavior, and are influenced by both environment and genotype. One possible phenotype then, is the expression of the

meme itself, because the meme's entity defines its morphology. This is similar to RNA viruses where the genetic information is contained within the RNA (the genotype), but when the RNA is expressed, this becomes a phenotypic trait [9]. Similarly, when a meme is posted (expressed), this becomes a phenotype.

There are many other meme phenotypes; we provide some examples of behavior phenotypes. One behavior phenotype is the average time a meme is posted. This phenotype depends on the meme's genotype (whether the meme says “post for a day”), and the environment (whether people using Facebook follow directions, remember to log in and update, etc.) Another meme phenotype is a meme's appeal to a certain demographic. This will depend on which variant of the meme is being used (whether the genotype is “healthcare” or “taxes”), and how much the intended target audience places importance on the issue being discussed (environment).

In biology, mutations occur randomly and blindly. Most mutations are neutral, but occasionally, some mutations are deleterious or advantageous [21]. Neutral mutations do not affect gene function, deleterious mutations are evolutionarily disadvantageous, while advantageous mutations provide a positive benefit to its carriers. Extending the analogy, neutral edits in memes are minor edits to the string that do not change the meaning of the meme. The Yule process is a neutral model where each meme variant has the same probability of copying itself, and possibly mutating again. Notably, most memes follow the Yule process (Figure 5) which suggests that most meme variants are neutral with respect to the parent's meme. We also find some ‘advantageous’ mutations that increase the likelihood of a meme being copied (Table 3). There is one critical difference between biology and social media: evolution is a blind process in biology, but in social media, there can be a conscious effort to create mutations that will intentionally spread a meme's presence (i.e. marketing campaigns). We discuss further similarities, as well as some mechanistic differences, pertaining to the more detailed replication mechanism, below.

The first is the slight preference for shorter variants of the same meme, analogous to the observations of bacteria favoring small genomes for fast replication [30]. Both memes and genes are shaped by the same two factors: their length has to be sufficient to encompass information that helps them to replicate, but excessive length slows down replication.

Even more interesting mutation mechanisms also have parallels across both genes and memes. For example, advantageous sequences can occur across multiple memes, likely transferred by a single individual from one meme to another. This process is analogous to lateral gene transfer in bacteria wherein useful genetic code, e.g. plasmids conferring antibiotic resistance, can be transferred through mechanisms other than replication. Yet another example is that of fusion. In genomes, two single genes are combined in a fusion event to create a single functional gene that retains both functionalities of the two single genes [14]. We observed several fusion events between different memes, some giving rise to popular variants.

Some similarities in the replication mechanism might be spurious. For example, mutations of both genes [8] and memes occur preferentially at the boundaries. Many mutations occur at the beginning and end of a gene (untranslated regions) because most of the function is in the middle of the gene and so mutations at the gene boundaries do not affect gene function [8]. Similarly, we observe that meme strings tend to have text edits at the beginning and end of a meme (Figure 6). The same factor could be at play, e.g. even the most common variants frequently start with a non-essential or generic address: “everyone”, “attention”, and end with words also not essential to the message: “if you agree”, “takes a

second”, “so your fb friends are aware”. However, we cannot rule out that the beginnings and ends of memes are simply more vulnerable because of the block copy-and-paste mechanism.

There are several limitations to this study. Narrowly defining the success of a meme to be the number of copies posted, while keeping the analogy of a gene's replicative success, does not account for all of the memes' effect, which might encompass the number of exposures to the meme and the resulting reaction e.g. in the form of likes and comments. Furthermore, while copy-and-paste memes provide an easy means to study unit of transmission needed in memetics, they do not address other mechanisms in cultural evolution, such as the adoption and evolution of norms, which may also hold interesting biological analogies [28].

In summary, we have described how information evolves as it is passed from individual to individual in a social network, sometimes exactly, and sometimes with a modification which produces a new variant. This process is well-described by the Yule model, with the mutation rate predicting the distribution of popularity among variants. Although many variants appear to emerge from neutral drift, there is evidence of some selection, as successful subsequences were found across memes, and individual meme variants were found to match preferences of individuals transmitting them. A clear limitation of the present study is that the information analyzed was of but one type (copy-and-paste memes) and was observed in a single environment. However, this same environment was the first opportunity to study the evolution of information with precision and on a large scale. We believe our findings are likely to be applicable to a range of environments where information diffuses through multiple steps with a potential for modification, and that the observed evolutionary process carries important implications for the reach and fidelity of information diffusing through social networks.

5. ACKNOWLEDGMENTS

We thank Brian Karrer, Eytan Bakshy, Jon Kleinberg, Jure Leskovec, Alex Dow and Mark Newman for comments.

6. REFERENCES

- [1] www.ladamic.com/research/data/memetext.html.
Supplementary table: Normalized text of the most common variant of memes, 2015.
- [2] L. Adamic, T. M. Lento, and A. T. Fiore. How you met me. *ICWSM'12*, 2012.
- [3] E. Adar and L. Adamic. Tracking information epidemics in blogspace. In *Web Intelligence'05*, pages 207–214. IEEE, 2005.
- [4] P. Andriani and B. McKelvey. Perspective-from Gaussian to Pareto thinking: causes and implications of power laws in organizations. *Organization Science*, 20(6):1053–1071, 2009.
- [5] E. Bakshy, I. Rosenn, C. Marlow, and L. Adamic. The role of social networks in information diffusion. In *Proc. WWW'12*, pages 519–528. ACM, 2012.
- [6] N. K. Baym. The performance of humor in computer-mediated communication. *JCMC*, 1(2):0–0, 1995.
- [7] C. H. Bennett, M. Li, and B. Ma. Chain letters and evolutionary histories. *Scientific American*, 288(6):76–81, 2003.
- [8] A. Chakravarti et al. Population genetics – making sense out of sequence. *Nature genetics*, 21(Suppl 1):56–60, 1999.
- [9] B. Charlesworth and D. Charlesworth. *Elements of evolutionary genetics*. Roberts and Company Publishers Greenwood Village, 2010.

- [10] J. Cheng, L. Adamic, P. A. Dow, J. M. Kleinberg, and J. Leskovec. Can cascades be predicted? In *WWW'14*, pages 925–936, 2014.
- [11] M. Coscia. Competition and success in the meme pool: a case study on Quickmeme.com. *ICWSM 2013*, 2013.
- [12] R. Dawkins. *The selfish gene*. Oxford University Press, USA, 1976.
- [13] S. F. Elena and R. E. Lenski. Evolution experiments with microorganisms: the dynamics and genetic bases of adaptation. *Nature Reviews Genetics*, 4(6):457–469, 2003.
- [14] A. J. Enright, I. Iliopoulos, N. C. Kyriakis, and C. A. Ouzounis. Protein interaction maps for complete genomes based on gene fusion events. *Nature*, 402(6757):86–90, 1999.
- [15] X. Gabaix. Zipf's law for cities: an explanation. *The Quarterly Journal of Economics*, 114(3):739–767, 1999.
- [16] M. Girvan, D. S. Callaway, M. E. J. Newman, and S. H. Strogatz. Simple model of epidemics with pathogen mutation. *Phys. Rev. E*, 65:031915, Mar 2002.
- [17] C. Heath, C. Bell, and E. Sternberg. Emotional selection in memes: The case of urban legends. *Journal of Personality and Social Psychology*, 81(6):1028, 2001.
- [18] F. Heylighen. What makes a meme successful? Selection criteria for cultural evolution. *Proc. 16th Int. Congress on Cybernetics*, 1998.
- [19] J. L. Iribarren and E. Moro. Impact of human activity patterns on the dynamics of information diffusion. *Phys. Rev. Lett.*, 103:038702, Jul 2009.
- [20] M. D. Kirby. Email forwardables: folklore in the age of the internet. *New Media & Society*, 7(6):770–790, 2005.
- [21] M. Kimura. *The neutral theory of molecular evolution*. Cambridge University Press, 1984.
- [22] S. Kirby, H. Cornish, and K. Smith. Cumulative cultural evolution in the laboratory: An experimental approach to the origins of structure in human language. *Proceedings of the National Academy of Sciences*, 105(31):10681–10686, 2008.
- [23] E. Koonin, Y. Wolf, G. Karev, et al. The structure of the protein universe and genome evolution. *Nature*, 420(6912):218–223, 2002.
- [24] J. Leskovec, L. Backstrom, and J. Kleinberg. Meme-tracking and the dynamics of the news cycle. In *KDD'09*, pages 497–506. ACM, 2009.
- [25] D. Liben-Nowell and J. Kleinberg. Tracing information flow on a global scale using internet chain-letter data. *PNAS*, 105(12):4633, 2008.
- [26] R. M. MacCallum, M. Mauch, A. Burt, and A. M. Leroi. Evolution of music by public choice. *PNAS*, 109(30):12081–12086, 2012.
- [27] B. Mandelbrot. A population birth-and-mutation process, I: Explicit distributions for the number of mutants in an old culture of bacteria. *Journal of Applied Probability*, pages 437–444, 1974.
- [28] A. Mesoudi, A. Whiten, and K. N. Laland. Towards a unified science of cultural evolution. *Behavioral and Brain Sciences*, 29(04):329–347, 2006.
- [29] T. Moothathu. Sampling distributions of Lorenz curve and Gini index of the Pareto distribution. *Sankhyā: The Indian Journal of Statistics, Series B*, pages 247–258, 1985.
- [30] N. A. Moran, A. Mira, et al. The process of genome shrinkage in the obligate symbiont Buchnera aphidicola. *Genome Biol*, 2(12):1–0054, 2001.
- [31] M. Newman. Power laws, Pareto distributions and Zipf's law. *Contemporary physics*, 46(5):323–351, 2005.
- [32] J. Qian, N. Luscombe, and M. Gerstein. Protein family and fold occurrence in genomes: power-law behaviour and evolutionary model. *Journal of Molecular Biology*, 313(4):673–681, 2001.
- [33] J. Ratkiewicz, M. Conover, M. Meiss, B. Gonçalves, S. Patil, A. Flammini, and F. Menczer.Truthy: mapping the spread of astroturf in microblog streams. In *WWW'11*, pages 249–252, New York, NY, USA, 2011. ACM.
- [34] W. Reed. The Pareto, Zipf and other power laws. *Economics Letters*, 74(1):15–19, 2001.
- [35] D. Romero, B. Meeder, and J. Kleinberg. Differences in the mechanics of information diffusion across topics: Idioms, political hashtags, and complex contagion on twitter. In *WWW'11*, pages 695–704. ACM, 2011.
- [36] E. Segev, N. A., N. Stolero, and L. Shifman. Families and networks of internet memes: The relationship between cohesiveness, uniqueness, and quiddity concreteness. *JCMC*, 2015.
- [37] L. Shifman. An anatomy of a YouTube meme. *New Media & Society*, 2011.
- [38] L. Shifman and M. Thelwall. Assessing global diffusion with Web memetics: The spread and evolution of a popular joke. *JASIST*, 60(12):2567–2576, 2009.
- [39] P. D. d. Silva and J. L. Garcia. YouTubers as satirists: Humour and remix in online video. *JedDEM*, 4(1):89–114, 2012.
- [40] M. Simmons, L. Adamic, and E. Adar. Memes online: Extracted, subtracted, injected, and recollected. *ICWSM 2011*, 2011.
- [41] L. Weng, A. Flammini, A. Vespignani, and F. Menczer. Competition among memes in a world with limited attention. *Scientific Reports*, 2, 2012.
- [42] G. Yule. A mathematical theory of evolution, based on the conclusions of Dr. JC Willis, FRS. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 213(402-410):21–87, 1925.