VISUAL REVIEW | Vol. 18, No. 7, 2026 | ISSN 2695-9631
International Visual Culture Review / Revista Internacional de Cultura Visual
CC BY ND
https://doi.org/10.62161/revvisual.v18.6496

INFLUENCE OF AI ON VIRALITY AND ATTENTION ON SOCIAL NETWORKS

JOSE ÁNGEL MERCADER JIMÉNEZ [1]
josea.mercader@upct.es
VIRGINIA ALARCÓN MARTÍNEZ [2]
virginia.alarcon@unir.net
[1]Universidad Politécnica de Cartagena, España
[2]Universidad Internacional de la Rioja, España

KEYWORDS

Artificial intelligence
Virality
Social networks
Empirical
Model
Diversity
Information

ABSTRACT

The object of study is quality discernment in a stylized model of an online social network, where individual agents prefer high-quality information. An empirical research methodology is employed, aiming to gain knowledge from observation of reality. The result is a successful balance between discriminatory power and information diversity. The findings reveal that the model predicts that low-quality information has the same probability of going viral, thus providing an explanation for the high volume of misinformation observed online.

RECIBIDO: 05/ 08 / 2026
ACEPTADO: 20/ 10 / 2026

1.Introduction

1.1. Backgrounds

Today, social media exerts a profound influence on global society. The cohort commonly referred to as "millennials" spends approximately five hours a day navigating platforms such as TikTok, X (formerly Twitter), and Instagram.

Consequently, numerous companies allocate substantial financial resources to promote their products and services through these modern communication channels. This has given rise to a form of indirect advertising on social media, as marketing content is frequently camouflaged within "entertaining posts, ideas, or memes" shared by so-called "influencers."

Approximately four centuries ago, Milton (1644) posited that in an environment of unrestricted intellectual exchange, truth ultimately asserts itself. Since that time, the notion of a “marketplace of ideas” has been invoked to underpin freedom of expression and has also been applied to the domain of scientific investigation.

The framework is conceptually analogous to evolutionary dynamics, in which differential survival is contingent upon species-specific traits, and to economic market mechanisms, where the intrinsic value of a commodity fundamentally shapes its diffusion, adoption, and long-term viability.

Two central components of this theory are the heterogeneity of ideas accessible to individuals and the market’s “discriminative efficacy,” defined as its capacity to ensure that the most robust or superior ideas attain prominence.

The rise of social media as a predominant mode of communication has profoundly transformed this marketplace by broadening participation and enabling the large-scale circulation and exchange of knowledge, information, and viewpoints. Consequently, following the disruption caused by modern digital media, both corporations and society at large face the daily challenge of identifying formats, tones, and content capable of capturing the attention of users navigating these networks.

In the present study, the terms "entertaining post," "idea," and "meme" will be used interchangeably to refer to a unit of transmissible information. According to the Royal Spanish Academy (2001), a "meme" is defined as a "cultural or behavioral trait transmitted from person to person or generation to generation through imitation."

A meme may take the form of a news link, phrase, hashtag, video, or image. On platforms such as X (Twitter) or Facebook, users encounter a daily flood of memes competing for attention. Nevertheless, cognitive limitations restrict both the number of social interactions an individual can maintain and the amount of information they can process, giving rise to what is commonly referred to as the “attention economy.”

The resulting flows of information from these complex dynamics have increasingly important implications for public policy. Accordingly, considerations of discrimination and diversity have become more crucial than ever in contemporary online information markets.

This study analyzes the discriminative performance and diversity of a simulated information network designed to approximate contemporary social media environments, where memes propagate through interpersonal transmission. Each meme is endowed with an intrinsic quality parameter, and the model evaluates how two key determinants the intensity of informational competition and the finiteness of user attention, shape the system’s capacity to prioritize high-quality content while preserving informational heterogeneity.

The findings indicate the emergence of a balance between quality-based selection and diversity, with both dimensions remaining comparatively robust under conditions of ample attentional resources and moderate information load.

It should be emphasized that calibrating the model with empirical data poses considerable challenges, as attributes such as value, innovation, reliability, and relevance are seldom directly measurable or definable a priori. This complicates the quantification of discriminative power in online information markets. Nevertheless, the distribution of information load and attention can be described using empirical data obtained from social networks.

Regrettably, the data indicate that online information networks often exist in a state where the link between quality and popularity is tenuous, falling short of the ideal balance.

This work draws on databases such as PubMed, Scopus, and Web of Science, incorporating observational, experimental, quasi-experimental, and mixed-methods studies from both economics and the social sciences. These studies explore the relationship between "quality" and "popularity" through both theoretical models and empirical data.

Adler (1985) argued that straightforward rationality considerations, based on the effort required to evaluate quality, can lead to the emergence of "stars" with disproportionate popularity, even when genuine differences in quality are absent. Following this reasoning, the model assumes that the cost of evaluation is negligible: each agent can assess quality, though the selection process remains noisy.

Weng et al. (2012) showed that, under conditions of constrained attention, certain memes attain viral diffusion independently of their intrinsic quality. Their framework, however, does not account for agents' potential preference for higher-quality content.

In cultural industries such as music, literature, and film, exceptional successes frequently surpass the average by a considerable margin, suggesting that top-performing works may be qualitatively differentiated from their competitors. Yet even domain experts consistently struggle to anticipate which products will ultimately achieve widespread success. This apparent contradiction has been examined in controlled experimental settings. Notably, in an artificial music market involving 14.341 participants who downloaded previously unknown songs either with or without information about prior participants' selections greater social influence amplified both outcome inequality and the unpredictability of success.

It is noteworthy that in this experiment, success was only partially determined by quality. While the best songs rarely performed poorly and the worst rarely succeeded, virtually any other outcome remained possible.

In line with these findings, Salganik et al. (2006), in their seminal "music lab" experiment, demonstrated that the relationship between quality and popularity can be distorted when consumers make decisions based on aggregate information. By contrast, the present study assumes a networked system in which new items are continuously introduced, and agents have access only to local information shared by their immediate neighbors.

1.2. Justification

This study analyzes the system's ability to encourage the emergence of high-quality content, investigating how this discriminatory capacity is constrained by the cognitive limitations of individual agents and how it impacts the level of diversity the system can preserve.

A noticeable trade-off emerges between quality discrimination and informational diversity, with both remaining comparatively strong when agents possess sufficient attention and are not subjected to excessive information overload.

It is important to note that calibrating the model with empirical data involves substantial challenges, as characteristics such as value, originality, reliability, and relevance are seldom directly measurable or clearly definable a priori. This difficulty complicates the evaluation of the system's discriminatory performance in online information markets. Nevertheless, the distributions of information load and attention can be empirically characterized using data obtained from social networks.

2. Objectives

A primary objective of this study is to examine the conditions under which the "best ideas" or "highest-quality memes" succeed in capturing the greatest collective attention, and to determine whether this outcome entails a reduction in intellectual diversity.

An additional goal has been the development of a "simple agent-based model" (ABM), grounded in the long-established tradition of conceptualizing the diffusion of ideas as an epidemic-like process, whereby messages spread through the connections of a network.

Within this framework, agents are modeled as nodes in a static network, with links representing social ties. The network dynamics embody the fundamental processes characteristic of prominent social media platforms. Every message or post conveys a "meme" or "idea," conceptualized as the primary unit of information transmitted between individuals.Importantly, several distinct messages may transmit the same meme.

3. Methodology of the proposed model

Building upon the definition provided by the Royal Spanish Academy (2001), a model is described as "an archetype or point of reference to be imitated or reproduced."

Consistent with this definition, it is important to emphasize that although the proposed “simple model” satisfies these criteria, it does not explicitly account for many of the behavioral, social, and technological factors that shape discrimination and diversity within the online marketplace of ideas.

Numerous studies have investigated the role of technological tools in separating reliable information from false or misleading content. It has been suggested that truthful content is inherently easier to disseminate, as original sources can be verified more readily, thereby raising the “cost” associated with sustaining misinformation online. Nevertheless, the rapid diffusion of disinformation across social media platforms may challenge this claim.

Moreover, the “wisdom of the crowd” enabled by social networks is theoretically expected to improve quality-based information discrimination by aggregating the heterogeneous judgments of a large population of individuals.

When individuals communicate, their opinions cease to be independent, which often generates greater confidence but lower accuracy. The mental and behavioral processes triggered by exposure to challenging opinions can weaken one’s ability to discriminate between reliable and less reliable information.

A clear illustration of this phenomenon is confirmation bias, which may have evolved as an adaptive mechanism to mitigate misinformation by comparing new information against pre-existing beliefs and adopting it only when it aligns sufficiently. On social media, however, this bias can undermine effective discrimination. Strategies such as accepting information solely when corroborated by multiple sources become ineffective, as users typically lack knowledge of the network’s structure, which is essential to verify the true independence of these sources.

Furthermore, confirmation bias is exacerbated online by the limited capacity to handle information overload resulting from the continuous influx of messages and the consequent need to rapidly filter out irrelevant content. As a result, the diversity of information within the marketplace can be impaired due to the interplay between cognitive and behavioral limitations and the algorithmic biases inherent to online platforms.

Moreover, mechanisms such as "triadic closure" a social network theory property where strong connections between two individuals and between one of them and a third party likely lead to a connection between the first and the third can be suboptimal for discovering relevant but unknown information. These selection mechanisms can group individuals into a limited number of homogeneous clusters, commonly known as “echo chambers” or “filter bubbles.” Over time, this segregation may foster greater polarization, as members of one group tend to automatically reject the perspectives of others.

Research indicates that, paradoxically, the very behavioral mechanisms evolved to manage information overload may reduce meritocracy and diversity within online information markets. This, in turn, facilitates the proliferation of misinformation and increases societal susceptibility to manipulation.

Evidence of widespread hoaxes, conspiracy theories, and false information on social media has become so extensive that digital misinformation is recognized as a major global threat.

From a modeling perspective, Crane and Sornette (2008) introduced an epidemic-style framework within social networks to capture both exogenous and endogenous surges in attention directed at a video.

The issue of fake news has become a prominent topic of discussion in regions such as the United States and Europe. Consequently, numerous studies have examined how network dynamics shape the popularity of individual memes.

4. Types of models.

4.1. The Epidemiological Model

Within the scope of this study, an epidemic model is defined as a collection of approaches employing mathematical, statistical, and computational techniques to investigate the propagation of transmissible entities across populations.

Ratkiewicz et al. (2010) implemented a model in which stochastic shifts in collective attention driven by exogenous events account for the wide-ranging distribution of popularity surges. Similarly, Bingol

(2008) proposed a dynamic framework in which agents possess the ability to remember and forget, while also leveraging recommendations to discover new agents. This model predicts a linear relationship between an agent's popularity and its memory capacity.

Regarding content engagement, Huberman (2013) examined how content novelty and popularity influence attention acquisition. In a related study, Wu and Huberman (2007) introduced a model in which the novelty of a news item diminishes over time, demonstrating that attention decays following a natural temporal scale.

According to Lerman et al. (2015), the combined influence of competition and positional bias, conceptualized as a manifestation of constrained attention on social media platforms, affects the exposure of memes and thereby restricts contagion dynamics.

The following figure provides a visual illustration of these concepts, depicting a meme diffusion model alongside the predicted distributions of popularity.

Figure 1a. Modeling meme propagation and the resulting popularity distributions

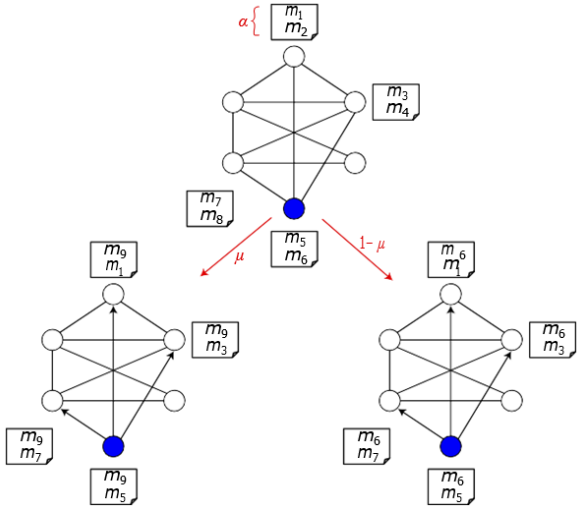

Source(s): Own elaboration, 2026.

Figure 1b. Probability density function.

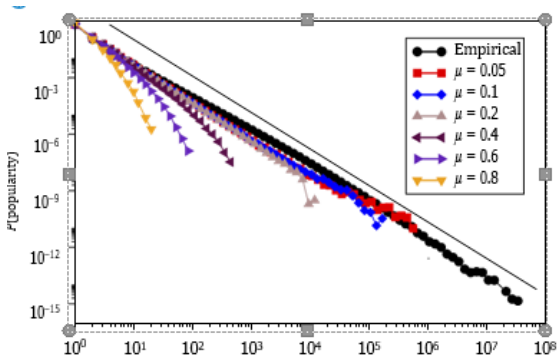

Source(s): Own elaboration based on study data, 2026.

As depicted in Figure 1, Part 1a models the behavior of a focal agent (blue node) at each time step. The agent either creates and shares a novel meme (m_9) with probability μ or, with probability 1 − μ, retransmits one of the α messages currently attended to in their feed (m_6). The chosen meme is then delivered to the agent's neighbors and positioned at the top of their feeds.

Part 1b reports the probability density function (PDF) describing the distribution of meme popularity p.

4.2. Network Definition and General Properties: Homogeneous and Heterogeneous Networks

Formally, a complex network is expressed as C = (V, E), where V represents the set of N nodes (or vertices) in the network, and E is a subset of m pairs of V, indicating the existence of a link or connection between the elements of said pair.

If the links consist of unordered pairs (where the link i → j is indistinguishable from enlace j → ij), the network is termed undirected or bidirectional. Conversely, if the order of the pairs is determinant, the network is considered directed or unidirectional.

These relationships are frequently represented through an adjacency matrix A, with entries a_{ij} = 1if there exists a directed link (i, j) contained in E, and 0 otherwise. As a result of this definition, if a network is undirected, A is symmetric; if the network is directed, A is asymmetric.

Furthermore, a network is connected if all its nodes are directly or indirectly linked to one another. In the case of disconnected networks, a connected component is defined as a subset of nodes that satisfy the aforementioned condition. The largest of these is referred to as the giant component.

A simple network is defined as a structure formed by undirected, unweighted links that lack "self-loops" (links from a node to itself, $a_{ij} = 0$) or multiple links between any two nodes i and j.

Finally, a complete graph is defined as a simple network in which a_{ii}=1 for all i ≠ j, and 0 otherwise; in this case, a node i is linked to all N - 1 remaining nodes.

Among the statistical properties of complex networks, the degree distribution, P(k), is frequently highlighted. In the case of undirected networks, the degree of a node i, denoted as ki, is defined as the number of links connected to i (that is, $ki = \sum^{N}_{j-1} a_{ij}$). Thus, the degree distribution P(k) determines the probability of finding a node with a specific degree k.

By observing P(k), two broad categories of networks emerge: homogeneous networks and heterogeneous networks. These are distinguished based on the "tail" of the P(k) distribution. If the probability decreases exponentially with the degree, the networks are termed homogeneous. A highly representative example is the random graph from the model (Erdős y Rényi ,1959).

In the aforementioned model, one begins with N disconnected nodes and establishes all possible connections with a fixed probability p. Consequently, its degree distribution P(k) is a binomial distribution, which converges to a Poisson distribution in the thermodynamic limit, provided the average degree ⟨k⟩ remains constant:

Formula 1. Degree distribution P(k)

$$P(k) = \binom{N-1}{k} p^k (1-p)^{N-1-k} \to \lim_{\substack{N\to\infty \\ \langle k\rangle \approx \text{cons.}}} P(k) = e^{-\langle k\rangle} \frac{\langle k\rangle^k}{k!}.$$

Source(s): Allen-Perkins, 2018.

The limiting case of homogeneous networks is comprised of regular networks. These are defined as networks where all nodes possess the exact same degree ki; that is, $P(k = k_i) = 1$.

On the other hand, a network is considered heterogeneous when P(k) is heavy tailed. If P(k) follows a power law, $P(k) \sim k^{-\gamma}$, the graph is referred to as a scale-free network.

The Barabási-Albert model (or B-A model) is the paradigm for this latter type of graph (Barabási and Albert, 1999). This model introduced a widely utilized concept within Network Science: preferential attachment. The B-A model is based on a network that grows through the addition of new nodes. Upon joining, each new node connects to a fixed number of existing nodes, m_{BA}, with a linking probability p_i proportional to the degree of each target node i (selected from the pre-existing nodes): $p_i = k_i / \sum^{N}_{j=1} k_j$.

Owing to this mechanism, the process is commonly referred to as preferential attachment. Considering the aforementioned factors, predictions generated by the model (α = 10) were subsequently compared with an empirical distribution derived from the observed frequency of hashtags within a sample of public tweets. The resulting empirical popularity exhibits $P(p) \approx p^{-1,94}$, represented by the black line as a visual reference.

As indicated previously, Figure 1(a) portrays the dynamics of the model devised to examine the market's capacity for discrimination, where each meme is characterized by a specific endogenous quality value. Agents observe the memes propagated by their neighbors, and it is presumed that the likelihood of an agent endorsing a meme and consequently promoting its spread is contingent upon its inherent quality.

Quality may represent different properties that increase the likelihood of a meme being shared, depending on the modeled context; these may validly include the originality of an idea, the aesthetic appeal of an image, or the veracity of a statement.

In contrast to standard epidemiological frameworks, it becomes clear that messages containing new memes are steadily introduced into the system through exogenous channels, without altering the core

premise of the model. The rate parameter μ is employed to regulate agents' information load, defined as the average number of memes an agent encounters per unit time.

Agents both produce messages that introduce new memes and circulate messages that originate either from themselves or from their neighbors, without altering the underlying dynamic of creation and diffusion. During dissemination, an agent can attend to a finite number of messages simultaneously. Assuming that neighbors' messages are presented in reverse chronological order, users typically focus on the most recent posts within a session.

Markets where multiple memes compete for the limited attention of social media users remain relatively underexamined. Weng et al. (2012) provided an agent-based modeling framework demonstrating that the combination of social network topology and finite user attention is sufficient for the emergence of viral memes. Extending this approach, Gleeson et al. (2016) formalized the dynamics as a critical branching process, showing that meme popularity follows a heavy-tailed power-law distribution.

Although these findings indicate that content quality is not strictly required to reproduce popularity patterns in online social networks, the impact of information quality has not been fully addressed. It is reasonable to propose that quality affects individual decisions in information consumption, motivating further theoretical analysis of how such behavior may influence system-level discriminatory dynamics.

This motivates a theoretical investigation aimed at assessing whether quality-based discrimination at the individual level can manifest as discriminatory power at the systemic level, and to what extent this occurs at the expense of the market's capacity to sustain a diverse informational ecosystem.

4.3. Proposed Model: Diffusion Framework and Simulation Specifications

The fundamental structure of the model examined in this study hereafter referred to as "our model" comprises a collection of agents interconnected within a social network. Each agent maintains a feed containing the α most recent messages generated by its neighbors.

The reverse chronological arrangement of the feed serves as a realistic simplifying assumption, characteristic of social media platforms such as X (Twitter). In other platforms, such as Facebook, ranking algorithms additionally incorporate factors including popularity and user interaction. Nonetheless, recency remains a central organizing principle across platforms.

As the system evolves over time, an agent "I" is randomly selected. With probability μ, agent *i* produces a message that introduces a novel meme, the quality of which is sampled from a uniform distribution on the unit interval.

Alternatively, with probability 1− μ , agent "i" chooses one of the messages available in its feed. The likelihood of selecting a given message is proportional to the quality of the meme it carries.

More formally, Mi let denote the feed of agent i , where ($|M_i| = \alpha$). The probability that a message m ∈ Mi is chosen is given by P(m)=f (m) /Σ ∈ f (j) j where f(m) represents the quality of the meme transmitted by message m.. The selected message is subsequently shared with the neighbors of agent i ; if a feed exceeds α messages, the oldest entry is removed.

This mechanism captures the manner in which limited attention is allocated to information disseminated through social ties. The two key model parameters enable us to analyze how the magnitude of the information load (μ) and the scope of attention (α) interact with the intrinsic quality of an idea and influence its probability of attaining success.

If the model's behavior is analyzed by simulating the process of information diffusion and load dynamics in synthetic scale-free networks understood as a specific class of complex networks, it becomes evident that such networks are characterized by a highly heterogeneous degree distribution. In a scale-free network, a small subset of nodes is highly connected, possessing a large number of links to other nodes, whereas the vast majority of nodes exhibit relatively low connectivity.

In the present context, as competition unfolds, certain memes fade rapidly, whereas others persist over extended periods and spread widely enough to capture a significant share of the network. This dynamic persists until the system converges to a steady state, where the mean number of distinct memes stabilizes at an approximately constant level. The equilibrium magnitude depends on the parameter μ. A meme's popularity is measured by the total amount of attention it garners throughout the network.

In practice, this is quantified by tracking the number of times a meme is shared or reshared. All measurements are carried out once the system has stabilized.

For each experiment and parameter setting, the model dynamics are simulated on a synthetic undirected network. While many real-world social networks are directed, the prevalence of reciprocal connections makes this approximation reasonable.

The findings presented here rely on scale-free networks constructed via the preferential attachment mechanism, comprising N=103 nodes and an average degree of K=20. These results extend to larger and more sparsely connected networks exhibiting higher clustering coefficients.

Once the system reaches stationarity, meme performance is assessed. Only memes introduced after the stationary state is attained are considered. Each meme is tracked from its initial sharing until it disappears entirely from the network, with both its intrinsic quality and cumulative popularity systematically recorded.

Across simulations, roughly 100.000 memes are introduced and subsequently forgotten following stationarity, yielding a comprehensive dataset for analysis is repeated 20 times, such that the popularity analyses encompass a total of 2×10^6 memes. Measurements of discriminative power and diversity are averaged across independent simulation runs.

4. Attention-Shifting Model

Instead of presuming a fixed feed depth and limiting activity to a single post or repost per session, we conceptualize users as browsing through their feeds, engaging with content and resharing it sequentially until they choose to conclude the session. During each session, with probability ρ, a user generates a post introducing a new meme and then exits. Conversely, with probability 1-p, the user enters a scrolling phase.

During a scrolling session, after sharing a post, the user ceases activity with probability q . If the user does not stop, with probability 1−q they continue scrolling, potentially sharing the next post, and this process iterates until the session concludes.

Furthermore, for each session, the stopping probability q is drawn uniformly from the interval [(q) + σ], where the parameter σ captures heterogeneity in attention across sessions, reflecting variations in user engagement and persistence.

This attention-shifting framework provides a more realistic representation of user behavior on social networks, capturing sequential interactions with content rather than assuming static or uniform engagement. When σ is small, the distribution of α is well approximated by an exponential decay; in the limit $\sigma \to 0$, $P(\alpha) \approx e\lambda(\alpha-1)$ con $\lambda = \ln(1 - (q)) < 0$. As σ approaches q, $(\sigma \to (q))$ for large α, the distribution converges to a power-law, $P(\alpha) \approx \alpha-2$; the resulting heavy tail reflects that users occasionally scroll through a large number of posts, as illustrated in Figure 5.

These behavioral patterns are consistently represented throughout this study, notably in Figure 3c. The parameters ρ, q, and σ can be calibrated to match the empirical distributions of μ and α observed on X(Twitter) and Tumblr (Figures 4b and 4c). Both distributions are attainable with approximately similar values of q and σ, whereas ρ varies across platforms.

To investigate the effect of attention heterogeneity σ on discriminative power, the scrolling mechanism is incorporated into the model to generate α. It is observed that high σ leads to a substantial reduction in discriminative power. Although α increases on average with σ, large values provide minimal additional benefit for discrimination, whereas low values result in significant losses in discriminative capability.

5. Analysis and Discussion

5.1. Connection between Meme Quality and Success

This section examines how information load influences the association between meme quality and observed performance.

The connection is not straightforward, since diffusion does not depend exclusively on intrinsic quality. Put differently, even lower-quality content can spread widely if it is disproportionately present in an agent's feed. The probability that a meme is reshared therefore emerges from the interaction of multiple elements, including its current level of popularity, the structure of the network, and agents' constrained attention.

Meme success may be defined in various ways, for instance in terms of persistence over time. In this analysis, however, success is measured through popularity, understood as the cumulative number of times a meme is redistributed across the network from its introduction until it vanishes entirely.

Figure 2 represents this relationship. In both panels (a) and (b), popularity (total reshares) is plotted against meme quality. Panel 2a reports outcomes for different values of the information load parameter μ while keeping attention fixed at α = 10. Panel 2b, in turn, explores how the pattern varies with attention α when μ is set to 0.1.

Figure 2: Average Meme Popularity

Figure 2a

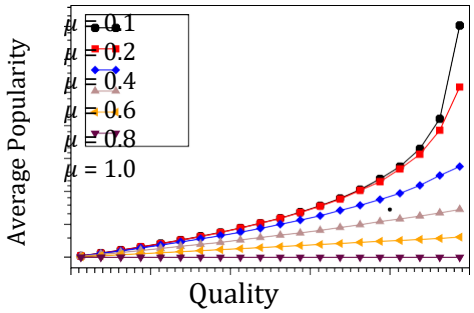

Source(s): Own elaboration based on study data, 2026.

Figure 2b

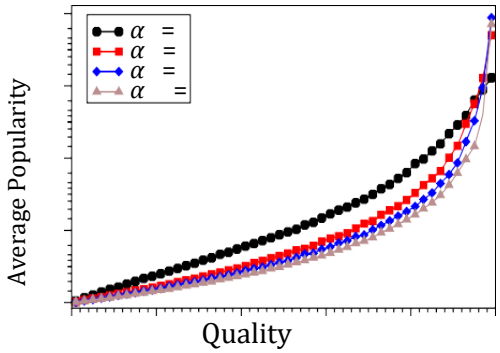

Source(s): Own elaboration based on study data, 2026.

Figure 2b shows that meme popularity is strongly influenced by the information load μ. At high μ values, the distribution is tightly concentrated, indicating that no individual meme attains viral status. As μ decreases below 0.2, the distribution broadens, reproducing the heavy-tailed pattern observed in empirical datasets and suggesting that some memes are capable of reaching a large fraction of the network.

Assuming no variation in quality, the popularity distribution would theoretically follow a power-law with exponent β = 1.5. When the empirical distribution is adjusted to the data, a higher exponent is obtained (β ≈ 1.94). This value aligns closely with the dynamics anticipated by a branching process under uniform fitness, which theoretically yields β = 2. This deviation reflects the influence of quality on meme propagation within the network.

Analysis of Figure 2a indicates that, on average, higher-quality memes have an increased probability of survival and propagation, yet the extent of this effect is strongly modulated by μ. In the limit of very low μ, only one meme generally endures, and its retention does not necessarily correspond to high quality.

For low values of μ ($0 < \mu \ll 1$), memes with very high quality are disproportionately more likely to succeed. As μ increases ($0 \ll \mu < 1$), the relative benefit associated with superior quality becomes less pronounced. In the limiting case of μ = 1, new memes are introduced continuously, preventing any meme from propagating, regardless of its intrinsic quality.

In essence, higher information load reduces the system's discriminative capacity, as the influence of quality on popularity diminishes. The role of limited attention in shaping the link between quality and

popularity is depicted in Figure 2b. Assuming $\alpha > 1$, agents have some ability to select which memes to share.

Consistent with expectations, average popularity rises with quality, indicating that higher-quality memes are considerably more likely to spread. As individual attention α increases, the resulting curves become more concave: average popularity grows more slowly for most memes, with the exception of those of the highest quality, highlighting stronger selective pressures favoring the most valuable content.

The overall relationship between quality and success can be summarized using a single metric of discriminative power, defined as the correlation between meme quality and its network-wide popularity.

Because neither variable follows a normal distribution, Kendall's rank correlation coefficient (τ) is employed. This coefficient is calculated by ranking memes according to both criteria and determining the number of concordant and discordant pairs, while appropriately accounting for tied ranks.

A higher τ indicates that superior-quality memes are more likely to propagate, reflecting the network's discriminative capacity. In the extreme case of $\tau = 1$, the rankings are fully aligned. Conversely, a lower τ reflects weak or negligible discrimination based on quality.

Figure 3 provides a visual representation of both discriminative power and content diversity. The figure consists of three panels: 3a, 3b, and 3c. Panel 3a shows τ (color-coded) as a function of information load and attention span. Panel 3b depicts normalized diversity, $H/H_{(\mu = 1)}$, as a function of information load and α. Panel 3c illustrates the trade-off between the system's ability to preferentially transmit high-quality memes and the overall diversity of content within the network.

In the network visualization, nodes represent agents, with node color corresponding to the most recently shared meme and node size representing the quality of that meme (larger nodes denote higher quality). Edges indicate social links, such as followers on X (Twitter) or friends on Facebook.

When the information load (μ) is low, the content circulating in the network is largely composed of high-quality memes, which naturally constrains overall diversity. As μ rises, a broader range of memes enters the system, leading to greater diversity, though at the expense of a weaker capacity to distinguish and promote the highest-quality content.

The example network presented includes $N = 128$ nodes with $\alpha = 10$.

Figure 3: Discriminative Capacity and Informational Diversity

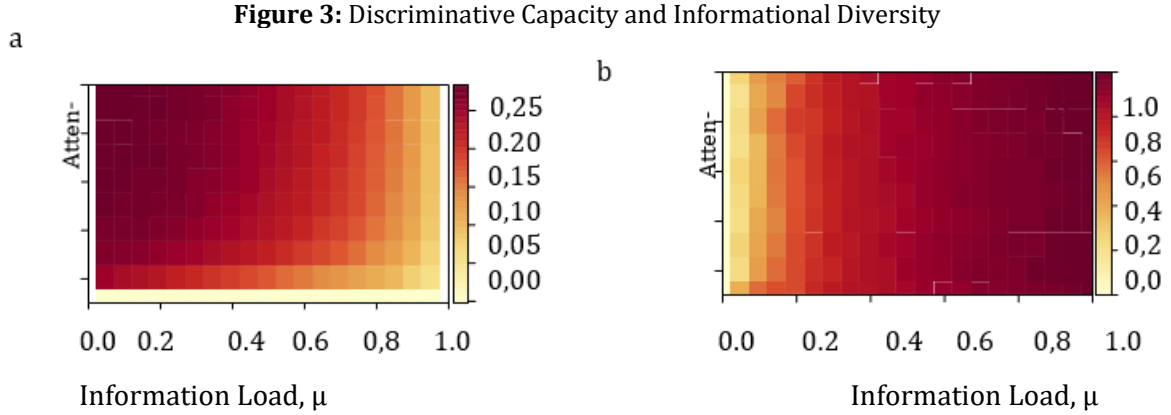

Source(s): Own elaboration based on study data, 2026

Figure 3c: Simultaneously Active Memes

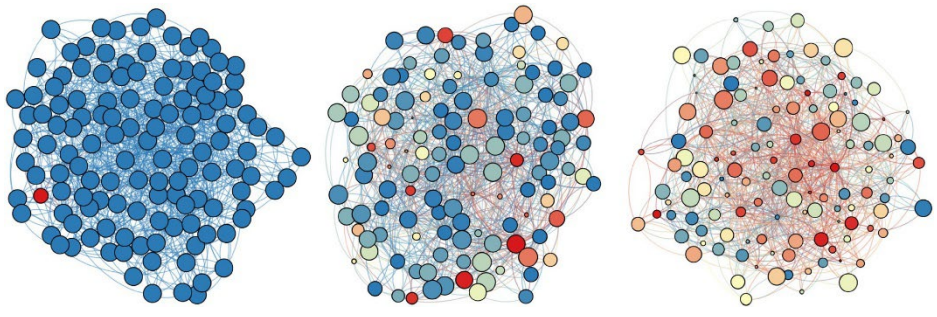

Source(s): Own elaboration based on study data, 2026.

Figure 3a indicates that the network's ability to distinguish between contents of differing quality weakens as information load increases and as individual attention becomes more constrained. Comparable outcomes emerge when mutual information is employed instead of Kendall's τ as an indicator of discriminative capacity.

As noted earlier, an effective social network should be capable of preferentially promoting higher-quality content. At the same time, an efficient communication environment must also safeguard informational diversity, allowing a plurality of distinct content sources to coexist within the system.

In Figure 3c, titled "Simultaneously Active Memes," it is shown that these two objectives are contradictory: the cost associated with the network's capacity to allow when a single high-quality meme comes to dominate, overall diversity inevitably declines, since numerous alternative memes despite possessing substantial intrinsic value receive only marginal attention. For this reason, we examine the trade-off that emerges from competition for limited attention within the network.

Diversity in the stationary regime is quantified through the Shannon entropy $H = -\Sigma_m P(m)\log P(m)$, where P(m)represents the share of total attention captured by meme m, measured as the proportion of messages containing that meme across all user feeds. The summation extends over all memes present at a given time and is averaged over a sufficiently long time window once the stationary state has been attained.

Entropy reaches its lower bound, H=0, when every node displays the same meme (μ = 0). Its upper bound occurs in the limiting case μ = 1, although this maximum depends on the attention parameter α. To control for this dependence, diversity is expressed as the normalized entropy H/H (μ = 1). .

As illustrated in panel 3b, this normalization renders diversity largely independent of α. In line with intuition, diversity increases with the intensity of information inflow and attains its highest levels for large μ. The interplay between discriminative capacity, diversity, and empirical calibration is presented in Figure 4.

Figure 4. Interplay between Discriminatory Capacity and Information Diversity: An Empirically Calibrated Approach

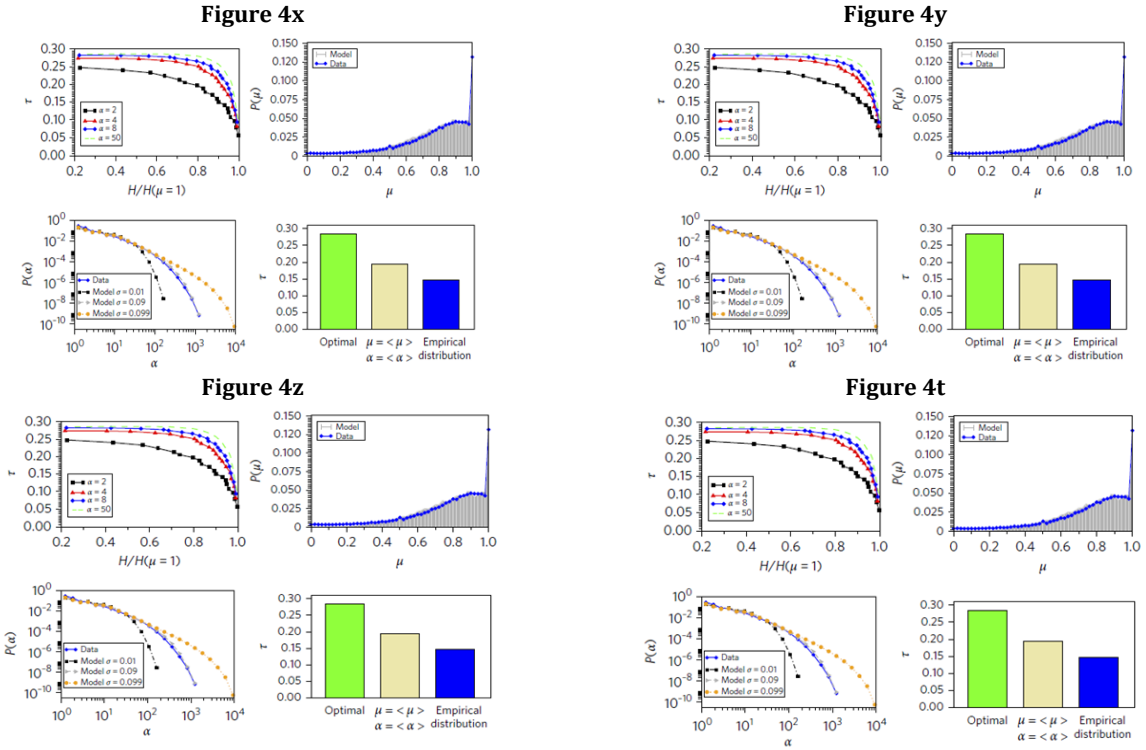

Source(s): Own elaboration based on study data, 2026.

As depicted in Figure 4, for any finite attention level (α > 1), the network transitions from a state of relatively high discriminative ability and low diversity under light information load to one of elevated diversity accompanied by reduced discriminative capacity as the information load grows.

The attentional capacity α significantly influences this trade-off. For a given degree of diversity, larger attention spans enable the network to better favor high-quality memes. Conversely, higher α allows the system to sustain greater diversity without substantially compromising discriminative capacity. In

particular, when α is sufficiently large, the network can achieve very high diversity with only a modest reduction in discriminative power.

The model is governed by two primary parameters: the information load μ and individual attention α. For simplicity, we assume that all agents introduce new memes at the same rate μ and possess identical attention α.

In practice, some users create new memes, while others primarily share content encountered through their social connections. Similarly, attention spans vary: some individuals engage with only a few messages, whereas others browse social platforms for prolonged periods.

When calibrating the model with empirical data, it is essential to consider multiple posts and reshares from a large sample of social media users to estimate each user's proportion of original versus shared memes. Panel 4y presents the resulting empirical distribution of new meme introduction rates, corresponding to the parameter μ in the model. Information load covers the entire range ($\mu \in [0,1]$) but is skewed toward higher values, peaking at $\mu = 1$, which represents users who post original content without sharing. The mean rate is relatively high ($\langle\mu\rangle \approx 0.75$).

The amount of attention a user allocates to evaluating information, ideas, and opinions on social networks depends not only on individual differences but also on contextual factors at the time of interaction. A single user may behave hurriedly at one moment and more carefully at another. In this study, finite attention was estimated based on the number of times a user paused on a post while browsing a social blogging platform. This measure corresponds to the parameter α in the current model.

Figure 4z illustrates the observed distribution of attention, α, which spans a wide range and exhibits a mean value of $\langle\alpha\rangle \approx 14$. An initial step in calibrating the idea-market model involves setting the parameters for information load and attention to their respective empirical averages, $\langle\mu\rangle \approx 0.75$ and $\langle\alpha\rangle \approx 14$. Under this configuration, the resulting discriminative power, τ, is shown in Figure 4t, reaching approximately 70% of the theoretical maximum τ achievable under conditions of high attention and low information load.

The assumption of uniform μ and α can be relaxed by sampling directly from the empirical distributions. Accordingly, a second model calibration draws μ and α values from the observed data, and the analysis is repeated under these more realistic conditions.

Panel 4x, labeled "Trade-off between discriminative power τ and diversity H/H ($\mu = 1$)," shows the trade-off for different attention levels α as μ varies from 0 to 1. The dashed line indicates convergence for large α.

Meanwhile, panel 4z, labeled "Empirical distribution of attention α derived from Tumblr," presents curves produced by the same scrolling model, calibrated with $\rho = 0.05$ and $\langle q\rangle = 0.1$.

Panel 4y, labeled "Empirical distribution of μ measured on X (Twitter)," illustrates the shaded region produced by the scrolling model described in the Methods section, using the parameters $\rho = 0.978$, $\langle q\rangle = 0.1$, and $\sigma = 0.09$.

Panel 4t, "System discriminative power τ for different empirical calibrations of μ and α," demonstrates that the ideal combination occurs at low information load ($\mu \leq 0.5$) and high attention ($\alpha \geq 50$). Under empirical heterogeneity, however, the network achieves a much lower discriminative power ($\tau \approx 0.15$), indicating that variation in information load and attention reduces the system's ability to preferentially propagate high-quality memes.

Given the significance of heterogeneity in user behavior, we further explore whether the empirical distributions of μ and α can be reproduced as emergent outcomes of the information market mechanisms, considering users' scrolling behavior.

Finally, Figure 5, titled "Popularity Distributions," depicts the frequency with which memes from two quality categories are shared within the model. Memes falling within the top two-thirds of the quality distribution ($f(m) \geq 0.34$) are designated as high quality, whereas the remaining memes are classified as low quality.

The inset of Figure 5 displays empirical Facebook data, contrasting articles that spread false claims or challenge true ones (low-quality) with those that fact-check false claims or reinforce verified information (high-quality).

Figure 5. Popularity Distributions

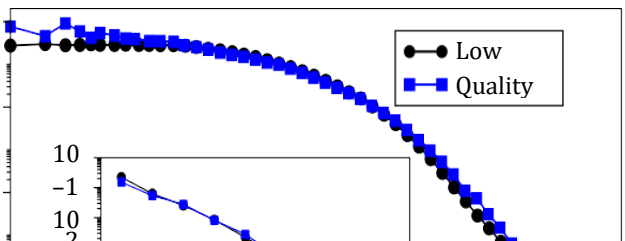



Source(s): Own elaboration based on study data, 2026

Figure 5 presents the popularity distributions for two categories of memes, corresponding to low and high-quality content. The results indicate that high-quality memes do not exhibit a consistent advantage in terms of their likelihood of achieving viral spread.

Empirically evaluating this prediction demands a quantifiable proxy for information quality, which is frequently challenging to establish. In this study, we analyzed data from **Emergent**, concentrating on social media posts referencing news articles categorized into two groups. The first group comprises articles that either endorse claims debunked by fact-checkers or contradict verified information, while the second group includes articles that fact-check false claims or corroborate verified facts. It is reasonable to consider the second group as indicative of higher-quality content.

The number of times articles from each group were shared was then compared. As shown in the inset of Figure 5, both low and high-quality groups display similar viral potential, in line with the predictions of the model.

6. Results

The results obtained confirm the scope and relevance of the study. "X" (Twitter) data were used to measure hashtag popularity. In addition to these primary data, several considerations were taken into account.

The analyzed dataset originates from a subset comprising approximately 15% of publicly accessible tweets, collected via X's (Twitter) streaming API in 2025. Less frequently used hashtags were consequently underrepresented within this subset.

The empirical metric μ was calculated for 150 X (Twitter) users by sampling their posts at random. For each individual, considering the number of original tweets (nt) and retweets (nr), μ was defined as nt divided by the sum of nt and nr ($\mu = nt / (nt + nr)$). Although users with lower activity levels were less likely to appear in the sample, this limitation is not expected to significantly distort the μ estimate.

Attention-related behavioral data were additionally obtained from roughly 160 mobile browsing sessions on Tumblr over a fifteen-day period in 2025. The feed interface of Tumblr shares general characteristics with those found on other social media platforms. A session was deemed to have ended if no interaction occurred for a minimum of 40 minutes. During each session, the number of times a user scrolled at least 500 pixels and paused for at least one second was recorded, providing a measure of α.

Data on Facebook posts linking to articles that either verified or challenged true and false claims were obtained from Emergent, a now-defunct rumor-tracking initiative. The meticulous work of Emergent's journalists was pivotal for this study, as they manually curated and assessed claims presented in articles shared on social media between October 2024 and May 2025. Through the Emergent API, the dataset included 842 high-quality articles (584 supporting 66 true claims and 258 fact-checking 82 false claims) and 483 low-quality articles (15 debunking 4 true claims and 468 propagating 82 false claims).

The model applies the same 2:1 ratio of high- to low-quality memes to set a quality threshold and produce the distributions illustrated in Figure 5.

Although relatively simple and governed by only a few parameters, the proposed model successfully reproduces fundamental behavioral dynamics of information diffusion across social networks. This approach enables the examination of how both information overload and constrained attention influence the network's ability to discriminate high-quality content, specifically the probability that the most valuable memes or ideas reach a wide audience.

7. Conclusions

The main insight from this study is that "survival of the fittest" does not automatically apply to information spread in online networks. Factors such as excessive information flow and limited user attention substantially diminish the system's ability to discriminate, reducing the likelihood that the highest-quality memes consistently gain prominence.

It is possible for networks to maintain both high meme diversity and meaningful discriminative power when attention is abundant and the system is not overwhelmed with content. An important direction for future work involves examining how network structure shapes both discriminative capability and the diversity of ideas. Although the present results are stable across variations in network size, density, and clustering, the model could be extended to account for additional features observed in real-world networks, such as segregated or modular communities. Investigating how these structural characteristics influence model predictions offers a fruitful area for further research.

Validating these predictions empirically remains challenging, largely because it is difficult to quantify the intrinsic quality of memes in practice. Nevertheless, it is possible to derive estimates for key parameters, such as the rate at which new memes are introduced and the level of user attention. According to these calibrations, real social networks operate under heterogeneous conditions of information load and attention, which places them in a regime of limited discriminative power.

If earlier studies had indicated that intrinsic quality is not essential to account for the broad variability in meme popularity, these results would be less alarming. Nonetheless, the evidence suggests that high-quality memes are not significantly more prone to achieving widespread adoption than low-quality ones. In this light, the rapid diffusion of hoaxes and false information illustrated in Figure 5 is consistent with the model's predictions.

The findings also suggest that one approach to enhance discriminative capacity in social networks is to reduce information overload by controlling the volume of content circulating online. At present, a significant share of accounts are automated bots, many of which flood platforms with low-quality content in attempts to manipulate discourse. Effectively curbing such practices could greatly improve the overall quality of information presented to users, thereby increasing the network's capacity to elevate the most valuable content.

8. Acknowledgments

This text emerged from the collaboration between two faculty members with extensive academic careers at two different universities in Spain. As Daniel Cassany (2012, p. 19) aptly states, "Let us remember that research is a collective endeavor."

First and foremost, this study would not have been possible without the contribution of numerous individuals and groups who participated as subjects of analysis; we wish to dedicate these lines to them as an expression of our sincere gratitude.

Secondly, we extend our appreciation to the students to whom we have had the opportunity to offer guidance and support throughout this study and over the course of our teaching careers.

Finally, as one of the authors of this work, I wish to extend my sincere gratitude to my colleague José Ángel Mercader Jiménez for his professionalism, unwavering support, and enduring presence throughout the many years we have worked together. Thank you for being an exceptional colleague.

References

Adler, M. (1985). Estrellato y Talento. *Economía* ,75, 208–12.

Barabási, A. Y Albert, R. (1999). Emergence of scaling in random networks. *Science*,286,509-512. DOI: 10.1126/science.286.5439.509

Bingol, H. (2008)**.** La fama surge como resultado de la poca memoria. *Física*, 77.

Crane, R. y Sornette, D. (2008). Clases dinámicas robustas reveladas mediante la medición de la función de respuesta de un sistema social. *Física*, 105, 15649–15653.

Erdös, P. y Rényi. A. (1959) Sobre gráficos aleatorios. *Matemáticas,* 6,90-297.

Gleeson, J. P., O'Sullivan, K. P., Baños, R. A. y Moreno, Y. (2016). Efectos de la estructura de red, la competencia y el tiempo de memoria en los fenómenos de difusión social. Física, X (6), 021019.

Huberman, B. A. (2013). Computación social y la economía de la atención. *Física,*151, 329–339.

Kang, J.-H. y Lerman, K. (2015). Computación Social, Modelado del Comportamiento y Predicción. *Ciencias de la Computación*, 9021,101–110.

Milton, J. (2026,20 de diciembre). Areopagitica.http://www.dartmouth.edu/milton/reading_ room/areopagitica/text.html

Ratkiewicz, J., Fortunato, S., Flammini, A., Menczer, F. y Vespignani, A. (2010). Caracterización y modelado de la dinámica de la popularidad en línea. Phys, 105.

Real Academia Española. (2001). *Diccionario de la lengua española* (22.a ed.). Madrid, España.

Periodismo ciudadano (2025,20 de diciembre).https://www.periodismociudadano.com/emergent-un-rastreador-de-rumores-para-verificar-informacion/.

Un rastreador de rumores en tiempo real. (2025,20 de diciembre). http://www.emergent.info/

Salganik, M. J., Dodds, P. S. y Watts, D. J. (2006). Estudio experimental de la desigualdad y la imprevisibilidad en un mercado cultural artificial. *Ciencia*, 311, 854–856.

Weng, L., Flammini, A., Vespignani, A. y Menczer, F. (2012). Competencia entre memes en un mundo con atención limitada. *Fisica*,2, 335.

Wu, F. y Huberman, B. A. (2007). Novedad y atención colectiva. Fisica,104, 17599–17601.