

# SCIENTIFIC REPORTS



OPEN

## How algorithmic popularity bias hinders or promotes quality

Giovanni Luca Ciampaglia<sup>1</sup> , Azadeh Nematzadeh<sup>2</sup>, Filippo Menczer<sup>1,2</sup>  & Alessandro Flammini<sup>1,2</sup>

Algorithms that favor popular items are used to help us select among many choices, from top-ranked search engine results to highly-cited scientific papers. The goal of these algorithms is to identify high-quality items such as reliable news, credible information sources, and important discoveries—in short, high-quality content should rank at the top. Prior work has shown that choosing what is popular may amplify random fluctuations and lead to sub-optimal rankings. Nonetheless, it is often assumed that recommending what is popular will help high-quality content “bubble up” in practice. Here we identify the conditions in which popularity may be a viable proxy for quality content by studying a simple model of a cultural market endowed with an intrinsic notion of quality. A parameter representing the cognitive cost of exploration controls the trade-off between quality and popularity. Below and above a critical exploration cost, popularity bias is more likely to hinder quality. But we find a narrow intermediate regime of user attention where an optimal balance exists: choosing what is popular can help promote high-quality items to the top. These findings clarify the effects of algorithmic popularity bias on quality outcomes, and may inform the design of more principled mechanisms for techno-social cultural markets.

Cultural markets, such as social media, the entertainment industry, and the world of fashion are known for their continuous rate of innovation and inherent unpredictability. Success of individual actors (e.g., artists) or products (e.g., songs, movies, memes) is in fact hard to predict in these systems<sup>1–3</sup>, mainly due to the presence of strong social reinforcement, information cascades, and the fact that quality is ultimately predicated on intangible or highly subjective notions, such as beauty, novelty, or virality.

In the absence of objective and readily measurable notions of quality, easily accessible metrics of success – such as the number of downloads of a song, or the number of social media followers of an individual – are often taken as input for future recommendations to potential consumers. Popularity and engagement metrics are intuitive and scalable proxies for quality in predictive analytics algorithms. As a result, we are exposed daily to content that is biased to some degree toward popularity, from bestseller lists to search engine results and from trending videos to engaging social media posts<sup>4</sup>.

The usefulness of such rankings is predicated on the *wisdom of the crowd*<sup>5</sup>: high-quality choices will gain early popularity, and in turn become more likely to be selected because they are more visible. Furthermore, knowledge of what is popular can be construed as a form of social influence; an individual’s behavior may be guided by choices of peers or neighbors<sup>6–12</sup>. These mechanisms imply that, in a system where users have access to popularity or engagement cues (such as ratings, number of views, likes, and so on), high-quality content will “bubble up” and allow for a more cost-efficient exploration of the space of choices. This is such a widely shared expectation that it has become routine for social media and e-commerce platforms to highlight popular and trending items. This is also reminiscent of herd behavior in financial markets, where the actions of previous investors carry information, and thus it is rational to adopt them<sup>13</sup>.

Popularity based metrics, however, can bias future success in ways that do not reflect, or worse, that hinder quality. This can happen in different ways. First, lack of independence and social influence among members of the crowd – as that implicitly induced by the availability of rankings – severely undermines the reliability of the popularity signals<sup>10</sup>. Second, engagement and popularity metrics are subject to manipulation, for example by fake reviews, social bots, and astroturf<sup>14,15</sup>.

Popularity bias can have more subtle effects. In search engines, the use of popularity in ranking algorithms was alleged to impede novel content from rising to the top, although such an entrenchment effect was shown to be mitigated by diverse user queries<sup>16</sup>. In social media, some memes inevitably achieve viral popularity in the

<sup>1</sup>Indiana University Network Science Institute, Bloomington, Indiana, USA. <sup>2</sup>School of Informatics, Computing, and Engineering, Indiana University, Bloomington, Indiana, USA. Correspondence and requests for materials should be addressed to G.L.C. (email: [glciampagl@gmail.com](mailto:glciampagl@gmail.com))

Received: 14 August 2017

Accepted: 1 October 2018

Published online: 29 October 2018

presence of competition among networked agents with limited attention, irrespective of quality<sup>2</sup>, and the popularity of memes follows a power-law distribution with very heavy tails<sup>17</sup>. Mechanisms such as unfriending and triadic closure facilitate the formation of homogeneous “echo chambers”<sup>18</sup> or “filter bubbles”<sup>19</sup> that may further distort engagement metrics due to selective exposure.

Even in the absence of engineered manipulation or social distortion, quality is not necessarily correlated with popularity. Consumers face a trade-off between performing cognitively expensive but accurate assessments based on quality and cognitively cheaper but less accurate choices based on popularity. Adler has shown that the cost of learning about quality will lead to “stars” with disproportionate popularity irrespective of differences in quality<sup>20</sup>. Such trade-offs are common in social learning environments<sup>21</sup>. Salganik *et al.* created a music-sharing platform to determine under which conditions one can predict popular musical tracks<sup>12</sup>. The experiments showed that in the absence of popularity cues, a reliable proxy for quality could be determined by aggregate consumption patterns. However, popularity bias – for example when users were given cues about previous downloads of each track – prevented the quality ranking from being recovered. By influencing choices, popularity bias can reinforce initial fluctuations and crystallize a ranking that is not necessarily related to the inherent quality of the choices<sup>22</sup>. This can happen even in the absence of explicit social signals, if the observed ranking is biased by popularity<sup>23</sup>. Similar results have been found in other studies<sup>8,24–26</sup> and have spurred a renewed interest in the topic of predictability in cultural markets. Idealized multinomial logit models have been used to understand the behavior of social influence. Van Hentenryck *et al.*<sup>27</sup> studied a model of trial-offer markets to analyze the effect of social influence on market predictability. In this model, users chose from a list of items ranked by quality rather than popularity; this modification makes the market predictable and aligns popularity and quality. Empirical tests of specific presentation policies combining quality and popularity do suggest that uncertainty can be reduced in this way<sup>28,29</sup>. Finally, Becker *et al.*<sup>30</sup> addressed the question of which network structure is most conducive to the wisdom of the crowd when people are influenced by others.

Another line of research that bears some connection to the present work is that of network growth models. Bianconi and Barabási<sup>31</sup> incorporated a notion of fitness in the preferential attachment model<sup>32</sup>. This approach provides one way to combine choices based on popularity and quality, but not a way to explore the effect of different mixtures between the two ingredients.

The conditions in which popularity bias promotes or hinders quality content have not been systematically explored. Here we do so by studying an idealized cultural market model in which agents select among competing items, each with a given quality value. A parameter regulates the degree to which items are selected on the basis of their popularity rather than quality. We find that this popularity bias yields a rich behavior when combined with the cognitive cost of exploring less popular items. As we shall see, popularity bias tends to hinder quality in general; but for a critical exploration cost, some popularity bias results in maximal average quality.

## Results

Our model considers a fixed number  $N$  of items. These represent transmissible units of information, sometimes referred to as *memes*<sup>33</sup>, such as music tracks, videos, books, fashion products, or links to news articles. Items are selected sequentially at discrete times. Each item  $i$  has an *intrinsic quality* value  $q_i$  drawn uniformly at random from  $[0, 1]$ . Quality is operationally defined as the probability that an item is selected by a user when not exposed to the popularity of the item. The *popularity* of item  $i$  at time  $t$ ,  $p_i(t)$ , is simply the number of times  $i$  has been selected until  $t$ . At the beginning each item is equally popular:  $p_i(0) = 1, i = 1 \dots N$ .

At each time step, with probability  $\beta$ , an item is selected based on its popularity. All items are first ranked by their popularity, and then an item is drawn with probability proportional to its rank raised to some power:

$$P_i^{pop}(t) = \frac{r_i(t)^{-\alpha}}{\sum_{j=1}^N r_j^{-\alpha}(t)} \quad (1)$$


where the rank  $r_i(t)$  is the number of items that, at time  $t$ , have been selected at least as many times as  $i$ . The exponent  $\alpha$  regulates the decay of selection probability for lower-ranked items. This schema is inspired by the *ranking model*, which allows for the emergence of scale-free popularity distributions with arbitrary power-law exponents<sup>34</sup>; it is consistent with empirical data about how people click search engine results<sup>16</sup> and scroll through social media feeds<sup>35</sup>. This model could accurately capture aggregate behavior even if individuals followed different selection schemes<sup>36</sup>.

Alternatively, with probability  $1 - \beta$ , an item is drawn with probability proportional to its quality:

$$P_i^{qual} = \frac{q_i}{\sum_{j=1}^N q_j} \quad (2)$$

After an item  $i$  has been selected at time  $t$ , we update its popularity ( $p_i(t+1) = p_i(t) + 1$ ) and the ranking. Two items have the same rank  $r$  if they have been selected the same number of times. If  $n$  items are all at the same rank  $r$ , then the next rank is  $r + n$ .

The model has two parameters:  $\beta$  regulates the importance of popularity over quality and thus represents the *popularity bias* of the algorithm. When  $\beta = 0$ , choices are entirely driven by quality (no popularity bias). When  $\beta = 1$ , only popularity choices are allowed, yielding a type of Polya urn model<sup>37</sup>. The parameter  $\alpha$  can be thought of as an *exploration cost*. A large  $\alpha$  implies that users are likely to consider only one or a few most popular items, whereas a small  $\alpha$  allows users to explore less popular choices. In the limit  $\alpha \rightarrow 0$ , the selection no longer depends on popularity, yielding the uniform probability across the discrete set of  $N$  items. Another way to think about the parameter  $\alpha$  is as inversely related to *attention*: low  $\alpha$  means that users have sufficient attention to consider all items (high attention), while high  $\alpha$  means that users have attention for only a limited number of items (low attention).



**Figure 1.** Predicted average quality at equilibrium. Heatmap of average quality  $\bar{q}$  as a function of  $\alpha$  and  $\beta$ , based on numerical solution of equilibrium condition derived in the text.

At equilibrium, after a large number of selection steps  $T$ , we characterize two properties of the distribution of popularity  $\{p_i\}_{i=1}^N$  with respect to the intrinsic quality distribution  $\{q_i\}_{i=1}^N$ . For brevity, we pose  $p_i = p_i(T)$  here. The first quantity we measure is the *average quality*  $\bar{q} = \sum_{i=1}^N p_i q_i / \sum_{i=1}^N p_i$  and the second property  $\tau$  is the *faithfulness* of the algorithm, i.e., the degree to which quality is reflected in popularity. We quantify faithfulness using Kendall’s rank correlation between popularity and quality<sup>38</sup>. The question we ask is whether it is possible to leverage some popularity bias to obtain a higher average quality, even at the cost of decreasing the algorithm’s faithfulness.

We can derive the values of both properties in the extreme cases of popularity bias. When  $\beta = 0$ , selections are made exclusively on the basis of quality and therefore one expects  $p_i \rightarrow q_i$  as  $T \rightarrow \infty$ . The rankings by quality and popularity are therefore perfectly aligned, and  $\tau = 1$ . In the limit of large  $N$  we can make a continuous approximation  $\bar{q} \rightarrow \int_0^1 q^2 dq / \int_0^1 q dq = 2/3$ . When  $\beta = 1$ , quality never enters the picture and any permutation of the items is an equally likely popularity ranking, which translates into  $\tau = 0$ . Also  $p_i \rightarrow 1/N$  and in the continuous approximation  $\bar{q} \rightarrow \int_0^1 q dq = 1/2$ .

What happens for intermediate values of popularity bias is harder to predict due to the role played by ranking, especially for high values of  $\beta$  where initial fluctuations can be strongly enhanced by the popularity-based ranking algorithm. Let us nevertheless try to derive  $\bar{q}$  at equilibrium by discretizing  $q \approx k/N$  for  $k = 1 \dots N$ :

$$\bar{q} \approx \frac{\sum_{k=1}^N \frac{k}{N} [(1 - \beta)P_k^{qual} + \beta P_k^{pop}]}{\sum_{k=1}^N [(1 - \beta)P_k^{qual} + \beta P_k^{pop}]} \tag{3}$$

Using Eq. 2 we can write the first probability term as

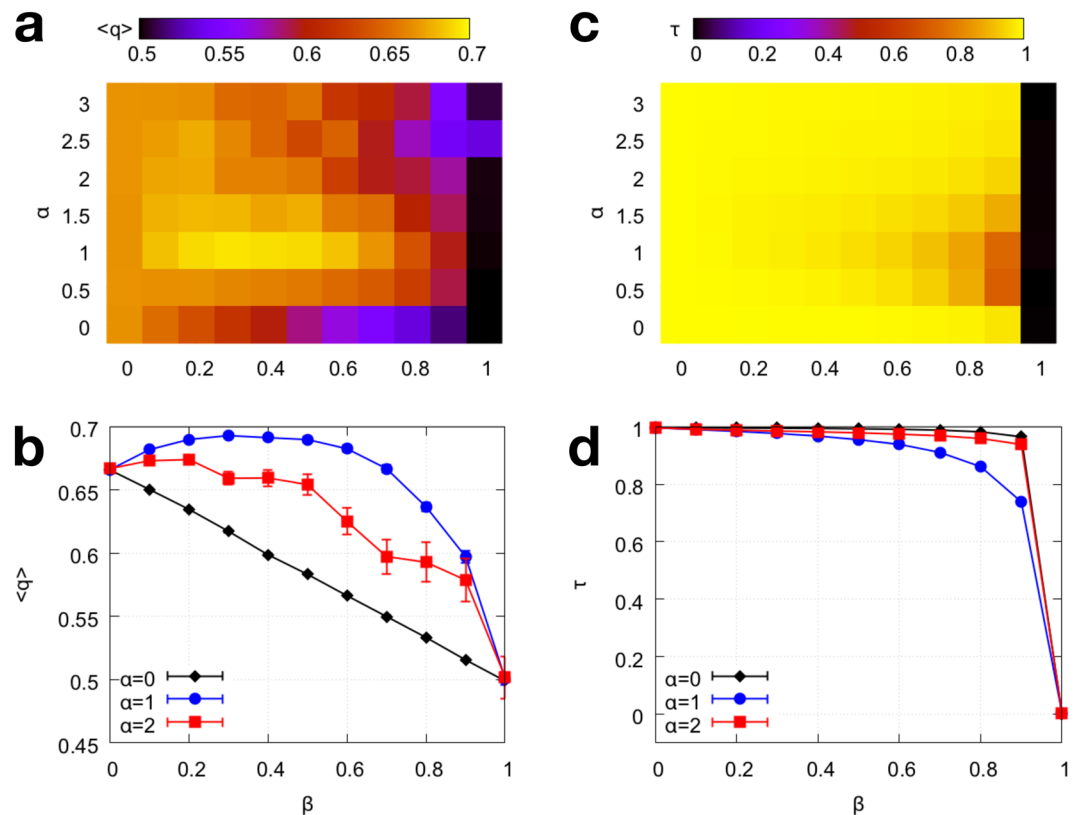
$$P_k^{qual} \approx \frac{k/N}{\sum_{j=1}^N j/N} = \frac{2k}{N(N + 1)} \tag{4}$$

For the second term of the probability, we assume that the rank will converge to a value that depends on  $\beta$ . In the extreme  $\beta = 0$ , the rank of the item with quality  $k/N$  should be narrowly peaked around  $r_k = N - k + 1$ , yielding the top rank when  $k = N$  (maximum quality) and the bottom rank when  $k = 1$  (lowest quality). In the extreme  $\beta = 1$ , quality plays no role, and therefore all ranks are equally probable, no matter the value of the item quality. For intermediate  $\beta$  we interpolate between these extremes, assuming that  $r_k$  is uniformly distributed between two limits  $r_k^{min} = 1 + (N - k)(1 - \beta^\gamma)$  and  $r_k^{max} = N - (k - 1)(1 - \beta^\gamma)$ . The heuristic parameter  $\gamma$  captures how effectively the perfect information produced (for ranking) by quality-based choices is preserved as  $\beta$  increases. From these assumptions and Eq. 1 we obtain

$$P_k^{pop} \approx \frac{1}{r_k^{max} - r_k^{min}} \sum_{r_k=r_k^{min}}^{r_k^{max}} \frac{r_k^{-\alpha}}{\sum_{j=1}^N j^{-\alpha}} \tag{5}$$

We plugged the approximations from Eqs 4 and 5 into Eq. 3 and solved numerically using  $\gamma = 4$  to obtain the equilibrium prediction shown in Fig. 1. The analysis suggests a non-trivial behavior of the system, with a maximum in average quality  $\bar{q}$  for intermediate values of  $\beta$  when  $\alpha$  is not too small and another maximum for high  $\beta$  and  $\alpha$ .

Given the strong assumptions in our derivation and the consequent uncertainty of the solution, let us turn to simulations for a more reliable analysis of the model’s behavior. We vary  $\beta$  systematically in  $[0, 1]$  and consider




**Figure 2.** Effects of popularity bias on average quality and faithfulness from model simulations. **(a)** Heatmap of average quality  $\bar{q}$  as a function of  $\alpha$  and  $\beta$ , showing that popularity bias can either hinder or promote average quality depending on exploration cost. **(b)** The location of the maximum  $\bar{q}$  as a function of  $\beta$  depends on  $\alpha$ , here shown for  $\alpha=0, 1, 2$ . When  $\alpha=1$  an intermediate amount of popularity bias is optimal. **(c)** Faithfulness  $\tau$  of the algorithm as a function of  $\alpha$  and  $\beta$ . **(d)**  $\tau$  as a function of  $\beta$  for the same three values of  $\alpha$ . Standard errors are shown in panels **(b,d)** and are smaller than the markers if not visible.

different values of  $\alpha$  between 0 and 3. We simulate 100 realizations for each parameter configuration. In each realization we perform  $T=10^8$  selections among  $N=1000$  items using Eqs 1 and 2 and store the final popularity values.

The dependence of the average quality  $\bar{q}$  on the popularity bias  $\beta$  and exploration cost  $\alpha$  is shown in Fig. 2(a,b). We observe that the predictions for  $\beta=0$  ( $\bar{q}=2/3$ ) and  $\beta=1$  ( $\bar{q}=0.5$ ) were correct. For intermediate popularity bias, the derived solution predicted that the optimal amount of popularity bias depends non-trivially on the exploration cost; however the simulation results provide a more accurate picture. When attention is abundant and the exploration cost is small ( $\alpha < 1$ ), popularity bias only hinders quality; the best average quality is obtained for  $\beta=0$ . In the opposite extreme of very limited attention, when popularity-based choices are strongly focused on the top-ranked items ( $\alpha=3$ ), the trend is more noisy but popularity bias tends to hinder quality; the optimal value of  $\bar{q}$  is again attained for  $\beta=0$ . The most interesting behavior is observed at a critical regime of attention, when the exploration cost is around  $\alpha=1$ . Here, the optimal value of  $\bar{q}$  is attained for an intermediate value of popularity bias  $\hat{\beta} = \arg \max_{\beta} \bar{q}(\beta) \approx 0.3$ . When  $\beta$  is smaller, the system may not be taking advantage of quality signals crowdsourced from other users. When  $\beta$  is higher, the system may be amplifying random initial fluctuations in popularity.

In Fig. 2(c,d) we show the behavior of faithfulness  $\tau$  as a function of  $\alpha$  and  $\beta$ . We observe that  $\tau=1$  for  $\beta=0$ , as predicted. Popularity bias always hinders the algorithm's faithfulness, however the effect is small as long as  $\beta$  is not too large. This suggests that in the regime where popularity bias improves quality on average, there is a small price to be paid in terms of over-represented low-quality items and under-represented higher-quality items. In general, the algorithm can retain faithfulness in the presence of moderate popularity bias. Near  $\alpha=1$  we observe a degradation in faithfulness as  $\beta$  grows larger, as low-quality items are wrongly picked up and become popular. This explains the degradation in average quality. But in general,  $\tau$  remains relatively high over a wide range of popularity bias values. We observe a sharp transition to the predicted value  $\tau=0$  at  $\beta=1$ , when quality no longer plays a role and popularity merely amplifies random fluctuations.

For a given value of  $\beta$ , if  $\alpha$  is low, the popularity bias hinders quality because it fails to enhance the signal provided by the quality-based choices, which are supported by exploration. To understand why quality is also hindered by the popularity bias when  $\alpha$  takes higher values, consider the evolution of the average quality in simulations of the model for different values of  $\alpha$ , shown in Fig. 3. By focusing only on the top ranked items ( $\alpha=2$ ), the system converges prematurely to a sub-optimal ranking, producing lower quality on average. In other words, with



**Figure 3.** Temporal evolution of average quality. Average quality  $\bar{q}$  is traced over time for  $\beta = 0.6$  and different values of exploration cost. Error bars represent standard errors across runs. Compared to the optimal case  $\alpha = 1$ , with more exploration ( $\alpha = 0$ ) the popularity bias just adds noise, and with less exploration ( $\alpha = 2$ ) it causes the system to converge prematurely to sub-optimal quality.

insufficient exploration the popularity bias risks enhancing initial noise rather than the quality-based signal. With more exploration ( $\alpha = 1$ ),  $\bar{q}$  continues to grow. The premature convergence to sub-optimal ranking caused by excessive popularity bias is also reflected in the increased variance of the average quality across runs of the model (larger error bars). This is consistent with the increase in variance of outcomes observed in other studies<sup>12,22</sup>.

## Discussion

Cultural markets like social media and the music and fashion industry account for multi-billion dollar businesses with worldwide social and economic impact<sup>1</sup>. Success in these markets may strongly depend on structural or subjective features, like competition for limited attention<sup>2,35</sup>. The inherent quality of cultural products is often difficult to establish, therefore relying on measurable quantitative features like the popularity of an item is hugely advantageous in terms of cognitive processing and scalability.

Yet, previous literature has shown that recommending already popular choices can be detrimental to the predictability and overall quality of a cultural market<sup>12</sup>. This left open the question of whether there exist situations in which a bit of popularity bias can help high-quality items bubble up in a cultural market.

In this paper we answered this question using an extremely simplified abstraction of cultural market, in which items are endowed with inherent quality. Our results show that optimizing the average quality of consumed items requires a careful tuning of quality- and popularity-based choices that depends on the focus on the most popular items. Popularity bias hinders average quality when users are capable of exploring many items, as well as when they only consider very few top items due to scarce attention. Critically, we found an intermediate regime of mild exploration cost in which some popularity bias is good, but too much is bad.

The model could be extended in many directions, for example assuming a population of networked agents with heterogeneous parameters. However, our approach leads to very general findings about the effects of popularity bias. While we confirmed that such a bias can distort assessments of quality, the scenario emerging from our analysis is richer than suggested by prior literature. First, it is possible to maintain a good correspondence between popularity and quality rankings of consumed items even when our reliance on popularity for our choices is relatively high. Second, one can leverage the wisdom of the crowd in the presence of limited attention, or let users make their own decisions when they are able to explore many items.

From a normative perspective, our results provide a recipe for improving the quality of content in techno-social cultural markets driven by engagement metrics, such as social media platforms. It is possible in these systems to estimate the exponent  $\alpha$  empirically, by measuring the probability that a user engages with an item as a function of the item's position in the feed. Given a statistical characterization (e.g., average or distribution) of the exploration cost, the bias  $\beta$  of the ranking algorithm could be tuned to maximize expected average quality.

These findings are important because in our information-flooded world we increasingly rely on algorithms to help us make consumption choices. Platforms such as search engines, shopping sites, and mobile news feeds save us time but also bias our choices. Their algorithms are affected by and in turn affect the popularity of products and information, and ultimately drive what we collectively consume in ways that we do not entirely comprehend. It has been argued, for example, that the engagement bias of social media ranking algorithms is partly responsible for the spread of low-quality content over high-quality material<sup>39</sup>. Evaluating such a claim is challenging, but the present results may lead to a better understanding of algorithmic bias.

## References

1. Park, J., Ciampaglia, G. L. & Ferrara, E. Style in the age of Instagram: predicting success within the fashion industry using social media. In *proc. of 19th conf. on computer-supported cooperative work & social computing (CSCW)*, pages 64–73 (2016).
2. Weng, L., Flammini, A., Vespignani, A. & Menczer, F. Competition among memes in a world with limited attention. *Scientific reports* 2(335), e335 (2012).

3. Weng, L., Menczer, F. & Ahn, Y.-Y. Virality prediction and community structure in social networks. *Scientific reports* **3**(2522), e2522 (2013).
4. Nikolov, D., Lalmas, M., Flammini, A. & Menczer, F. Quantifying biases in online information exposure. *Journal of the American Society for Information Science and Technology*, In press; preprint arXiv:1807.06958 (2018).
5. Surowiecki, J. *The Wisdom of Crowds*. (Anchor, 2005).
6. Bakshy, E., Hofman, J. M., Mason, W. A. & Duncan, J. W. Everyone's an influencer: quantifying influence on Twitter. In *proc. of 4<sup>th</sup> int. conf. on Web search and data mining*, pages 65–74 (2011).
7. Christakis, N. A. & Fowler, J. H. *Connected: The surprising power of our social networks and how they shape our lives*. Little, Brown (2009).
8. Krumme, C., Cebrian, M., Pickard, G. & Pentland, S. Quantifying social influence in an online cultural market. *PLOS ONE* **7**(5), e33785 (2012).
9. Kristina, L. & Hogg, T. Leveraging position bias to improve peer recommendation. *PLOS ONE* **9**(6), e98914 (2014).
10. Lorenz, J., Rauhut, H., Schweitzer, F. & Helbing, D. How social influence can undermine the wisdom of crowd effect. *Proceedings of the national academy of sciences* **108**(22), 9020–9025 (2011).
11. Muchnik, L., Aral, S. & Taylor, S. J. Social influence bias: A randomized experiment. *Science* **341**(6146), 647–651 (2013).
12. Salganik, M. J., Dodds, P. S. & Watts, D. J. Experimental study of inequality and unpredictability in an artificial cultural market. *Science* **311**(5762), 854–856 (2006).
13. Banerjee, A. V. A simple model of herd behavior. *The Quarterly Journal of Economics* **107**(3), 797–817 (1992).
14. Ferrara, E., Varol, O., Davis, C., Menczer, F. & Flammini, A. The rise of social bots. *Communication of the ACM* **59**(7), 96–104 (2016).
15. Ratkiewicz, J. *et al.* Detecting and tracking political abuse in social media. In *proc. of 5th int. AAAI conf. on weblogs and social media (ICWSM)*, pages 297–304 (2011).
16. Fortunato, S., Flammini, A., Menczer, F. & Vespignani, A. Topical interests and the mitigation of search engine bias. *Proceedings of the national academy of sciences* **103**(34), 12684–12689 (2006).
17. Gleeson, J. P., Ward, J. A., O'Sullivan, K. P. & William, T. L. Competition-induced criticality in a model of meme popularity. *Physical review letters* **112**(4), 048701 (2014).
18. Sunstein, C. R. *Republic.com 2.0*. Princeton university press (2009).
19. Pariser, E. The filter bubble: How the new personalized Web is changing what we read and how we think. *Penguin* (2011).
20. Adler, M. Stardom and talent. *American economic review* **75**(1), 208–12 (1985).
21. Rendell, L. *et al.* Why copy others? insights from the social learning strategies tournament. *Science* **328**(5975), 208–213 (2010).
22. Hogg, T. & Lerman, K. Disentangling the effects of social signals. *Human computation* **2**(2), 189–208 (2015).
23. Hodas, N. O. & Lerman, K. How limited visibility and divided attention constrain social contagion. In *proc. of ASE/IEEE int. conf. on social computing* (2012).
24. Gilbert, E. Widespread underprovision on Reddit. In *proc. of the 2013 conference on computer supported cooperative work (CSCW)*, pages 803–808 (2013).
25. Salganik, M. J. & Watts, D. J. Leading the herd astray: An experimental study of self-fulfilling prophecies in an artificial cultural market. *Social psychology quarterly* **71**(4), 338–355 (2008).
26. Salganik, M. J. & Watts, D. J. Web-based experiments for the study of collective social dynamics in cultural markets. *Topics in cognitive science* **1**(3), 439–468 (2009).
27. Van Hentenryck, P., Abeliuk, A., Berbeglia, F., Maldonado, F. & Berbeglia, G. Aligning popularity and quality in online cultural markets. In *proc. of 10th int. AAAI conf. on Web and social media (ICWSM)*, pages 398–407 (2016).
28. Abeliuk, A., Berbeglia, G., Van Hentenryck, P., Hogg, T. & Lerman, K. Taming the unpredictability of cultural markets with social influence. In *Proc. of 26th Int. Conf. on World Wide Web (WWW)*, pages 745–754 (2017).
29. Abeliuk, A., Berbeglia, G., Cebrian, M. & Van Hentenryck, P. The benefits of social influence in optimized cultural markets. *PLOS ONE*, **10**(4), 1–20, 04 (2015).
30. Becker, J., Brackbill, D. & Centola, D. Network dynamics of social influence in the wisdom of crowds. *Proceedings of the national academy of sciences* **114**(26), E5070–E5076 (2017).
31. Bianconi, G. & Barabási, A.-L. Competition and multiscaling in evolving networks. *EPL (Europhysics Letters)* **54**(4), 436 (2001).
32. Barabási, A.-L. & Albert, R. Emergence of scaling in random networks. *Science* **286**(5439), 509–512 (1999).
33. Dawkins, R. *The selfish gene*. Oxford university press (1989).
34. Fortunato, S., Flammini, A. & Menczer, F. Scale-free network growth by ranking. *Physical review letters* **96**(21), 218701 (2006).
35. Xiaoyan, Q. *et al.* Limited individual attention and online virality of low-quality information. Preprint arXiv:1701.02694 (2017).
36. Estes, W. K. The problem of inference from curves based on group data. *Psychological bulletin* **53**(2), 134–140 (1956).
37. Mahmoud, H. *Polya Urn Models*. Chapman & Hall/CRC (2008).
38. Kendall, M. G. & Gibbons, J. D. *Rank correlation methods*. E. Arnold (1990).
39. Bell, E. J., Owen, T., Brown, P. D., Hauka, C. & Rashidian, N. The platform press: how Silicon Valley reengineered journalism. Technical report, Tow Center for Digital Journalism (2017).

## Acknowledgements

The authors would like to thank Diego F.M. Oliveira and Kristina Lerman for helpful feedback. This material is based upon work supported in part by the James S. McDonnell Foundation (grant 220020274), the National Science Foundation (award CCF-1101743), and the Defense Advanced Research Projects Agency (DARPA) and the Army Research Office (ARO) under Contract No. W911NF-17-C-0094. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the funders.

## Author Contributions

A.N., G.C., F.M. and A.F. designed the research; A.N. and F.M. performed the analysis; F.M., G.C. and A.F. wrote the paper.

## Additional Information

**Competing Interests:** The authors declare no competing interests.

**Publisher's note:** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2018