



Title	Can the Wikipedia moderation model rescue the social marketplace of ideas?
Authors(s)	Yasseri, Taha, Menczer, Filippo
Publication date	2023-09
Publication information	Yasseri, Taha, and Filippo Menczer. "Can the Wikipedia Moderation Model Rescue the Social Marketplace of Ideas?," September 2023. https://doi.org/10.1145/3578645 .
Item record/more information	http://hdl.handle.net/10197/12712
Publisher's statement	© the Authors 2023. This is the author's version of the work. It is posted here by permission of ACM for your personal use. Not for redistribution. The definitive version was published in Computing Profession {66, 9, (2023)} http://doi.acm.org/10.1145/3578645
Publisher's version (DOI)	10.1145/3578645

Downloaded 2026-05-01 23:47:42

The UCD community has made this article openly available. Please share how this access benefits you. Your story matters! (@ucd_oa)



© Some rights reserved. For more information

Can the Wikipedia moderation model rescue the social marketplace of ideas?

Taha Yasseri^{1,2,3,4*} and Filippo Menczer^{5,6}

¹School of Sociology, University College Dublin, Dublin, Ireland.

²Geary Institute for Public Policy, University College Dublin, Dublin, Ireland.

³Oxford Internet Institute, University of Oxford, Oxford, UK.

⁴Alan Turing Institute for Data Science and AI, London, UK.

⁵Observatory on Social Media, Indiana University, Bloomington, USA.

⁶Luddy School of Informatics, Computing, and Engineering, Indiana University, Bloomington, USA.

* Corresponding author: taha.yasseri@ucd.ie

Abstract

Facebook announced a *community review program* in December 2019 and Twitter launched a *community-based platform* to address misinformation, called Birdwatch, in January 2021. We provide an overview of the potential affordances of such community based approaches to content moderation based on past research. While our analysis generally supports a community-based approach to content moderation, it also warns against potential pitfalls, particularly when the implementation of the new infrastructures does not promote diversity. We call for more multidisciplinary research utilizing methods from complex systems studies, behavioural sociology, and computational social science to advance the research on crowd-based content moderation.

Keywords: Social Media, Echo Chambers, Content Moderation, Crowd-sourcing, Homophily

In the late 1980s, Tim Barners-Lee invented the World Wide Web to facilitate information sharing and collaboration among researchers. Rapid and widespread adoption by academic and commercial entities followed in the 1990s. After the bursting of the dot-com bubble at the dawn of the new millennium, several social web platforms deeply transformed the information ecosystem within a few years. Wikis, blogs, social tagging and sharing sites, online social networks, and microblogs all significantly lowered the costs of producing and sharing content. This transition was hailed as a harbinger of more democratic participation in the information and knowledge economy — the realization of the “marketplace of ideas.” No project was a better icon of such utopian dreams than Wikipedia: the revolutionary vision of an encyclopedia that can be read, and more importantly, edited by anyone, surprised media scholars and practitioners. The potential of the largest online collaborative project in history was soon realized, and the quality of Wikipedia articles reached a level comparable to professionally edited encyclopedias (6) despite its reported limitations and biases. Facebook and Twitter appeared a few years after Wikipedia. Ten years into the millennium, democratic revolutions were credited to these social media platforms.

It did not take long for this sweet story to take a bitter twist. Hate speech, misinformation, polarization, manipulation, extremism, and data privacy scandals became commonplace online. January 2021 saw Wikipedia celebrate its 20th birthday, enduring the test of time as a reputable source of information. During the same month, the spread of disinformation and conspiracy theories on social media led to one of the darkest moments in the history of the American democracy, when the US Capitol was attacked on January 6.

How is it possible that the marketplace of ideas works on Wikipedia but not on social media? Of two social web technologies, one stands as the most successful example of collaborative and healthy information sharing while the other is blamed for epistemic chaos. Let us explore this paradox by focusing on network effects and content moderation strategies. To be sure, there are many other important differences between Wikipedia and social media platforms, including design, business models, and user motivations and characteristics. However,

a review of past research on Wikipedia and social media points to Wikipedia content moderation as a model for rescuing the social media marketplace of ideas, provided there is a serious intention by the commercial platform owners to promote a healthier information environment.

Why Wikipedia Works

Wikipedia is not the only community-based information sharing initiative, but arguably is the most successful and for sure the most researched one. Despite its generally acknowledged reliability, any experienced Wikipedia editor can tell stories of “edit wars” in which they have found themselves. Wikipedia edit wars have been studied in detail, showing that in most cases they lead to consensus after enough time (15). Agent-based simulations have helped reveal the secret of consensus-reaching mechanisms in Wikipedia. Articles, as shared products of editorial activities, facilitate indirect interactions between editors with opposing opinions, who would not engage in any positive interactions otherwise (14). These indirect interactions provide both mechanisms and incentives to bring editors closer to one another and reach consensus. The same models, validated and calibrated against empirical data from Wikipedia editorial activities, were used to examine the effects of banning editors with extreme opinions. Counterintuitively, the exclusion of extremists hinders the consensus reaching process (9). Moreover, a moderate amount of controversy is correlated with higher article quality (10) and higher completeness and neutrality ratings. Wikipedia is not free of bias and abuse, e.g., inauthentic and coordinated accounts to manipulate content. Over time, the platform has dealt with these threats by evolving a complex set of community rules and norms along with a hierarchy of user and administrator roles, privileges, and tools (including benign bots).

Why Social Media Do Not Work

As opposed to Wikipedia, where articles are meant for a diverse community, the content created and shared on social media is targeted to one’s social connections. Studies have reported that social networks on Twitter (4) and Facebook (12) tend to be structured around ideologically homogeneous echo-chambers (Fig. 1). Homophily, the tendency to connect with like-minded individuals, leads to homogeneous communities both online and offline. However, social media mechanisms make it very easy to create as well as to cut off a social tie. In fact, platforms encourage us to “block” users whose content we dislike. Mathematical analysis (1) and agent-based simulations (11) show that such costless unfollowing accelerates the formation of segregated echo chambers. Influence and social pressure ensure strong homogeneity of opinions within an echo chamber. One can generate engagement and gain popularity more easily by addressing their homogeneous audience than by appealing to common ground across diverse communities.




Figure 1. Echo-chamber structure in the friend/follower network of a sample of news-sharing Twitter users. Node colour is based on partisanship, inferred from the political alignment of shared news sources. Node size is proportional to the fraction of shared links to low-credibility sources. (Source: doi.org/10.37016/mr-2020-55).

Another important ingredient of social media is algorithmic bias. News feed ranking algorithms favour content that is likely to trigger engagement. Popularity also plays a role on Wikipedia, where content tends to be generated around topics that attract attention (2). But on social media, engagement is directly amplified by the platforms. This leads to greater exposure of low-quality content (3) as well as ideologically-aligned content that confirms or reinforces existing beliefs. Algorithmic bias also increases vulnerability to manipulation by inauthentic or coordinated accounts, compared to Wikipedia. The incentives created by all of these social media characteristics — targeting of homogeneous audiences as well as amplification of low-quality, confirmatory, and manipulated content — go against reaching consensus outside of one’s group. Worse, hate speech and verbal abuse are natural byproducts of this highly polarized information ecosystem and spill over to real-world violence (5).

A Community Based Solution

Ronald Reagan once said: “Peace is not absence of conflict, it is the ability to handle conflict by peaceful means.” The current social media approach of hiding conflict through self-selection, unfollowing, removal of content, and account bans neither prevents nor mitigates conflict. Some researchers suggested exposing users to counter-attitudinal content, a positive algorithmic bias designed to break the bubbles. However, experiments show that partisan users become more entrenched in their beliefs once they are exposed to opposing views (7). Moreover, mathematical models suggest that no matter how much positive algorithmic bias is applied, as long as unfollowing is costless, echo chambers will inevitably form (1). What has been shown to be effective, in contrast, is sharing personal experience (7) and —if we have learned one thing from the Wikipedia experience— collaborative interaction. Such collaboration in the context of social media could be aimed at tackling misinformation and community policy violations. Recent experiments suggest that crowdsourced layperson judgments can be effective at identifying misinformation (8). Such a community approach could scale up fact-checking and moderation practices while mitigating both misinformation and polarization.

One might argue that the population of Wikipedia editors is more homogenous than that of social media. While this may be true, a polarized crowd can be even more effective in producing high quality content compared with a homogenous team (13). Just like the population of Wikipedia editors is much narrower than that of its consumers in both size and diversity, the community of volunteer moderators will be small and self-selecting compared to social media users. Robustness to coordinated manipulation will therefore require mechanisms to enforce diverse and representative participants. But Wikipedia teaches us that it is possible for community rules and norms to overcome these challenges.

Conclusion

Solutions to the epistemic chaos brought about by social media platforms will have to involve multiple dimensions and stakeholders, including digital media literacy and government regulation. Past research on Wikipedia and social media platforms has shown that content moderation and network effects are an important part of the equation. Collaboration, civilized challenge, and mediated controversy are keys to the success of Wikipedia, while platform-amplified engagement focused on homogeneous audiences contributes to clash and abuse — online and offline. Wikipedia offers a community model for content moderation. Facebook announced a community review program (<https://about.fb.com/news/2019/12/helping-fact-checkers>) in December 2019 and Twitter launched a community platform to address misinformation, called Birdwatch (https://blog.twitter.com/en_us/topics/product/2021/introducing-birdwatch-a-community-based-approach-to-misinformation.html), in January 2021. We do not yet know how these systems will work or whether they will succeed. But the research presented above supports such initiatives as perhaps the most viable option at hand to rescue the social media marketplace of ideas.

Acknowledgments

We are grateful to David G. Rand, Janos Kertesz, Ksenia Musaelyan, and Dimitar Nikolov for feedback and comments on the manuscript. Funding: TY is partially supported by Google through “The Offensive Internet? Examining Cultures of Hate and Prejudice Online” project. FM is supported in part by John S. and James L. Knight Foundation, and Craig Newmark Philanthropies.

Competing Interest Statement: The opinions and recommendations expressed here are those of the authors and do not necessarily reflect the views of these funders. The authors declare no competing interests.

References

1. Blex, C. and Yasseri, T. 2020. Positive algorithmic bias cannot stop fragmentation in homophilic networks. *The Journal of Mathematical Sociology*. (Sep. 2020), 1–18. DOI:<https://doi.org/10.1080/0022250x.2020.1818078>.
2. Ciampaglia, G.L. et al. 2015. The production of information in the attention economy. *Scientific Reports*. 5, 1 (May 2015). DOI:<https://doi.org/10.1038/srep09452>.
3. Ciampaglia, G.L. et al. 2018. How algorithmic popularity bias hinders or promotes quality. *Scientific Reports*. 8, 1 (Oct. 2018). DOI:<https://doi.org/10.1038/s41598-018-34203-2>.
4. Conover, M.D. et al. 2012. Partisan asymmetries in online political activity. *EPJ Data Science*. 1, 1 (Jun. 2012). DOI:<https://doi.org/10.1140/epjds6>.
5. Gallacher, J.D. et al. 2021. Online Engagement Between Opposing Political Protest Groups via Social Media is Linked to Physical Violence of Offline Encounters. *Social Media + Society*. 7, 1 (Jan. 2021), 205630512098444. DOI:<https://doi.org/10.1177/2056305120984445>.
6. Giles, J. 2005. Internet encyclopaedias go head to head. *Nature*. 438, 7070 (Dec. 2005), 900–901. DOI:<https://doi.org/10.1038/438900a>.
7. Kubin, E. et al. 2021. Personal experiences bridge moral and political divides better than facts. *Proceedings of the National Academy of Sciences*. 118, 6 (Jan. 2021), e2008389118. DOI:<https://doi.org/10.1073/pnas.2008389118>.
8. Pennycook, G. and Rand, D.G. 2019. Fighting misinformation on social media using crowdsourced judgments of news source quality. *Proceedings of the National Academy of Sciences*. 116, 7 (Jan. 2019), 2521–2526. DOI:<https://doi.org/10.1073/pnas.1806781116>.
9. Rudas, C. et al. 2017. Understanding and coping with extremism in an online collaborative environment: A data-driven modeling. *PLOS ONE*. 12, 3 (Mar. 2017), e0173561. DOI:<https://doi.org/10.1371/journal.pone.0173561>.
10. Ruprecht, T. et al. 2020. Relating Wikipedia article quality to edit behavior and link structure. *Applied Network Science*. 5, 1 (Sep. 2020). DOI:<https://doi.org/10.1007/s41109-020-00305-y>.
11. Sasahara, K. et al. 2020. Social influence and unfollowing accelerate the emergence of echo chambers. *Journal of Computational Social Science*. 4, 1 (Sep. 2020), 381–402. DOI:<https://doi.org/10.1007/s42001-020-00084-7>.
12. Schmidt, A.L. et al. 2017. Anatomy of news consumption on Facebook. *Proceedings of the National Academy of Sciences*. 114, 12 (Mar. 2017), 3035–3039. DOI:<https://doi.org/10.1073/pnas.1617052114>.
13. Shi, F. et al. 2019. The wisdom of polarized crowds. *Nature Human Behaviour*. 3, 4 (Mar. 2019), 329–336. DOI:<https://doi.org/10.1038/s41562-019-0541-6>.
14. Török, J. et al. 2013. Opinions, Conflicts, and Consensus: Modeling Social Dynamics in a Collaborative Environment. *Physical Review Letters*. 110, 8 (Feb. 2013). DOI:<https://doi.org/10.1103/physrevlett.110.088701>.
15. Yasseri, T. et al. 2012. Dynamics of Conflicts in Wikipedia. *PLoS ONE*. 7, 6 (Jun. 2012), e38869. DOI:<https://doi.org/10.1371/journal.pone.0038869>.