https://doi.org/10.1038/s41562-026-02644-9

Exposure matters more than persuasiveness for digital mass persuasion

Ben M. Tappin

Check for updates

The persuasive impact of digital media content depends both on its persuasiveness per person and on how many people are exposed to it, and there is greater variability in exposure than persuasiveness across content. Together, these facts suggest exposure matters more than persuasiveness for digital mass persuasion.

The popularity of digital and social media has given rise to widespread concerns about digital mass persuasion, from microtargeted personalized advertising to political and health misinformation, foreign influence operations and, most recently, 'deep fakes' and propaganda generated by increasingly sophisticated artificial intelligence (AI). When studying these concerns, researchers typically assess the persuasive impact of such content by exposing people to it and measuring the extent to which it persuades them; how much it changes the way they think, feel or act.

However, the persuasive impact of such content depends only in part on how persuasive it is, and it is usually the less important part. More important is how many people are exposed to the content. Why? In short, the impact of a message depends both on its persuasiveness per person and on how many people are exposed to it. Although exposure and persuasiveness both matter for impact, in practice exposure tends to vary more across digital media content than persuasiveness does: high-exposure content reaches tens or hundreds of times more people than typical content, whereas highly persuasive messages are not tens or hundreds of times more persuasive than typical ones. As a result, typically persuasive content that receives high exposure is generally more impactful than highly persuasive content that receives only typical exposure. In this sense, exposure matters more than persuasiveness for digital mass persuasion.

The upshot is that if we are interested in understanding the mass persuasive impact of digital media content — whether from a research, policy or practitioner perspective — we should generally pay more attention to content exposure than to its persuasiveness. This argument bears directly on several influential and highly active areas of research, including work on misinformation and AI-driven persuasion, and has important practical implications for how scholars study mass persuasion and influence on digital and social media. I describe these implications after laying out the argument.

Measuring mass persuasive impact

For someone to be persuaded by something — an advertisement, news headline, AI chatbot and so on — they must first be exposed to it. Therefore, we can think about the mass persuasive impact of content as the product of two things: (1) its persuasiveness, defined as how effectively it persuades someone conditional on their being exposed to it, and (2) the number of people exposed to it[1].

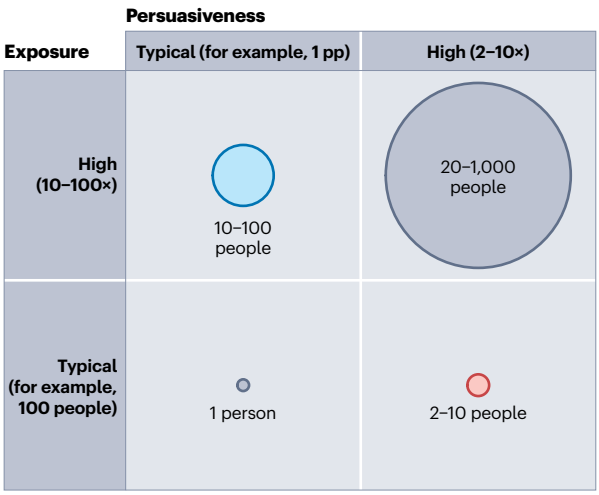

Fig. 1 | Persuasive impact as a function of exposure and persuasiveness. Hypothetical numbers of people persuaded in expectation, assuming typical exposure is 100 people and typical persuasiveness is 1 pp (these particular numbers are illustrative only).

To illustrate, suppose a political advertisement persuades 5 of every 100 viewers who see it to change their vote — a persuasiveness of five percentage points (pp). If 10,000 people are exposed to this advertisement, its expected persuasive impact is $0.05 \times 10{,}000 = 500$ changed votes. If the advertisement were twice as persuasive (10 pp), its expected impact would be $0.10 \times 10{,}000 = 1{,}000$ changed votes; however, if the advertisement were much less persuasive (1 pp) but nevertheless reached a much larger number of people — say, one million — its expected impact would be 10,000 changed votes.

These examples show that both persuasiveness and exposure matter for impact in principle. But does one matter more than the other in practice? To answer this question, we can compare the variances of persuasiveness and exposure across content. For instance, suppose that variance is greater in persuasiveness than exposure: highly persuasive content is 10× as convincing as typical content, whereas high-exposure content reaches only 5× as many people as typical-exposure content. In this scenario, persuasiveness would generally matter more for impact than exposure, whereas if variance were greater in exposure than persuasiveness the reverse would be true.

Exposure has higher variance

We need not limit ourselves to hypothetical scenarios. Studies have empirically estimated variance in exposure and persuasiveness across

BOX 1

Avenues for future work

The present thesis suggests concrete avenues for future work, several of which I outline below.

Refining the thesis
Not all digital media content is alike. The distribution of persuasiveness and exposure may be notably different for video content versus social media posts versus podcasts, for instance. These distributions may further depend on whether the content was explicitly designed to persuade versus something else (for example, entertain or inform). Collecting exposure and persuasiveness data systematically for content in different domains would enable more direct comparison; exposure may be revealed to dominate persuasiveness in some domains more than others, and in others still perhaps not at all.

Validating exposure-integrated designs
Understanding the mass persuasive impact of digital media content requires us to better understand exposure to content, not just its persuasiveness. 'Faux feed' designs are emerging as a method for doing so by enabling close experimental control and measurement over both exposure and persuasiveness of content[15]. However, they are artificial and their exposure results may generalize poorly to the real world. To validate these designs, future work could collect content with a range of exposures from real digital media, and test whether the attention given to that content in faux feed experiments replicates the real-world pattern. If not, the design could be iterated upon until validation improves.

Studying AI-designed content
The capabilities of AI are advancing apace, including in content creation (for example, video or audio). If mass persuasive impact depends more on exposure than persuasiveness, future work could compare the attention-grabbing performance of AI-designed content against that of human-designed content collected from the real world — perhaps in the validated faux feed designs described above. Whether and to what extent AI beats human content along this dimension has implications for AI-driven mass persuasive impact.

Estimating subgroup heterogeneity in exposure
A consistent finding in persuasion research is that, conditional on exposure, the persuasiveness of content does not differ dramatically from person to person — that is, subgroup heterogeneity is small[15]. However, this may not be the case for attention given to content, which plausibly depends a lot on the person. Exposure could thus differ dramatically across people and, as a result, persuasive impact could be highly heterogenous across subgroups. Studies that incorporate both exposure and persuasiveness can answer this question and may refine the current consensus of small subgroup heterogeneity.

content, which allows for comparison. To make the comparison fair, we must use the same scale – say, establishing the gap between the median (typical) and 90th+ percentile (high-performing) content for each of exposure and persuasiveness, and then comparing the size of these gaps.

Starting with exposure, a recent study[2] estimated the distribution of views across all public videos on YouTube, and found that the median video received 35 views; the 90th percentile between 1,000 and 2,000 views (30–60× median); and the 96th percentile approximately 10,000 (300×). Other studies that examined YouTube and TikTok accounts about politics report similar magnitudes, with accounts in the 90th percentile receiving 10–100× the median total views[3], and accounts' most-viewed videos receiving 40× (YouTube) and 64× (TikTok) their median video[4]. Broadly similar magnitudes also characterize other types of content, such as podcast downloads, social media account followers and website traffic[5].

This evidence describes organic exposure (user-driven attention), yet for paid exposure – such as digital marketing campaigns – the picture appears comparable. For example, studies of advertising campaigns on Facebook find the number of impressions (views) received by the largest campaigns to be 10–60× the median campaign[6,7]. Studies suggest the amount of money spent on campaigns also follows a similar distribution[5] – further evidence, insofar as exposure scales approximately linearly with spend.

In summary, for both organic and paid exposure, there is evidence to suggest high-performing content receives somewhere between 10× and 100× greater exposure than typical content.

Moving on, we can estimate the variance in persuasiveness across content by consulting 'megastudies' – specifically, survey experiments in which researchers test the persuasiveness of different content, holding fixed other features of the study design.

The results of three recent megastudies – which tested 11 messages encouraging action on climate change[8], 25 interventions to reduce US partisan animosity[9] and 21 honesty oaths in a tax compliance game[10] – suggest that highly persuasive content is approximately 2–5× more persuasive than the typical content tested. Another study that estimated the distribution of persuasiveness across 617 real political advertisements tested through the 2018 and 2020 US election cycles reported the 95th percentile advertisement to be 2–3× more persuasive than the median ad[11]. The content in these studies was preselected to maximize its persuasiveness, which plausibly compresses the variance somewhat. A recent study mitigates this concern; it tested approximately 100 vaccine-sceptical headlines that were not preselected for persuasiveness[1]. Among these headlines, the 95th percentile was approximately 7× more persuasive than the median, and the maximum 10× the median.

In summary, this evidence suggests that highly persuasive content is somewhere between 2× and 10× more persuasive than typical content. An important caveat is that these estimates do not come from a systematic review of all relevant studies; future research is needed to systematically characterize the state of existing knowledge on the variance in exposure and persuasiveness across digital media content.

Implications

The variance in exposure to digital media content is plausibly larger than the variance in persuasiveness. As a result, typically persuasive content that receives high exposure will generally be more impactful on the population level than highly persuasive content that receives only typical exposure. In this sense, exposure matters more than persuasiveness for mass persuasive impact (Fig. 1).

This argument has both substantive and methodological implications for highly influential and active areas of research, which I discuss below.

AI-driven persuasion. There is widespread concern that advanced AI could become superhumanly persuasive and thereby exert powerful influence over public opinion and behaviour. Mobilized by this, a rapidly growing body of research exposes people to AI content and measures its persuasiveness[12,13]. This work has yielded valuable insights about the persuasiveness of AI-driven personalization, its rhetorical strategies, the effect of AI model size and training, and how it compares to human persuasion.

However, the argument of this paper suggests these insights may not be the most important for understanding AI's potential mass persuasive impact. It does not matter how persuasive a ten-minute AI conversation is if few people will talk to it that long. A potentially higher-value lever for AI mass influence could be its capability to generate reliably high exposure to persuasive content — perhaps through the formation of parasocial relationships that sustain repeated user engagement, and through the rapidly growing adoption of AI chatbots as everyday interfaces for information seeking. This suggests reorienting the field away from asking 'how persuasive is AI?' towards 'how much exposure can AI generate for persuasive content?'

Misinformation. Over the past decade, an explosion of concern about 'fake news' has driven a large body of research that examines the persuasiveness of falsehoods propagated on digital media, with renewed impetus since the COVID-19 pandemic and recent advances in AI. Studies that show that such content can influence people's beliefs has worked to validate concerns about online misinformation. However, evidence suggests people's exposure to outright falsehoods is in fact quite limited relative to other content online[1,14]. This is plausibly because falsehoods are often spread by fringe websites that, although unconstrained by journalistic norms against lying, do not enjoy a large public following. By contrast, content that is technically true yet misleading — for instance, claims that report a true event but omit important context — may often be more likely to appear in mainstream media and thereby receive greater exposure[1].

This paper's argument implies that such technically true, yet misleading, content could be disproportionately more impactful than outright falsehoods, despite it seeming less egregious and even if it were less persuasive per exposure. The study of vaccine-sceptical headlines described above provides direct evidence for this. Headlines explicitly flagged as misinformation by fact-checkers were more persuasive per exposure than unflagged headlines (many of which came from mainstream sources). Yet the unflagged headlines received much wider exposure. As a result, the less persuasive unflagged content had an estimated 46× greater overall impact on reducing users' vaccination intentions[1].

This suggests a reorientation of misinformation research away from a focus on the persuasiveness of outright falsehoods, and towards greater focus on the exposure profile of content — especially that which is technically true yet misleading.

Methods of persuasion research. The dominant paradigm for studying persuasion is the survey experiment, in which people are randomly exposed to treatment or control content and subsequently have their attitudes measured. This design conditions on exposure (typically by paying for people's attention), which makes it well suited for estimating persuasiveness, but poorly suited for studying the more important determinant of mass persuasive impact outlined here: exposure itself. Indeed, as we have just seen, mistaking persuasiveness for persuasive impact can be actively misleading.

Researchers interested in persuasive impact therefore need designs that simultaneously incorporate exposure. Promising options include sampling real content with known exposure and combining it with experimental persuasiveness estimates[1]; simulated social-media-feed experiments[15]; or field experiments on real digital media platforms, although these can be expensive and raise ethical issues. No design is perfect, but what is clear is that the study of mass persuasive impact should more systematically incorporate exposure.

Conclusion

Digital mass persuasion is of considerable interest to scholars, policymakers and the public. Here I have outlined the argument that exposure generally matters more than persuasiveness for the impact of digital mass persuasion, and should thus receive greater attention. Future work can further explore and test the implications of this thesis (Box 1).

Ben M. Tappin
Department of Psychological and Behavioural Science, London School of Economics and Political Science, London, UK.
e-mail: b.tappin@lse.ac.uk

References

1. Allen, J., Watts, D. J. & Rand, D. G. *Science* **384**, eadk3451 (2024).
2. McGrady, R., Zheng, K., Curran, R., Baumgartner, J. & Zuckerman, E. *J. Quant. Descr. Digit. Media* https://doi.org/10.51685/jqd.2023.022 (2023).
3. Munger, K., Hindman, M., Yalcin, O., Phillips, J. & Bisbee, J. *J. Quant. Descr. Digit. Media* https://doi.org/10.51685/jqd.2025.006 (2025).
4. Guinaudeau, B., Munger, K. & Votta, F. *Comput. Commun. Res.* **4**, 463–485 (2022).
5. Tappin, B. M. Exposure matters more than persuasiveness for digital mass persuasion - supplement. *Zenodo* https://doi.org/10.5281/zenodo.23686752 (2026).
6. Gordon, B. R., Zettelmeyer, F., Bhargava, N. & Chapsky, D. *Mark. Sci.* **38**, 193–225 (2019).
7. Gordon, B. R., Moakler, R. & Zettelmeyer, F. *Mark. Sci.* **42**, 768–793 (2023).
8. Vlasceanu, M. et al. *Sci. Adv.* **10**, eadj5778 (2024).
9. Voelkel, J. G. et al. *Science* **386**, eadh4764 (2024).
10. Zickfeld, J. H. et al. *Nat. Hum. Behav.* **9**, 169–187 (2025).
11. Hewitt, L. et al. *Am. Polit. Sci. Rev.* **118**, 2021–2039 https://doi.org/10.1017/S0003055423001387 (2024).
12. Summerfield, C. et al. *Nat. Hum. Behav.* **9**, 2420–2430 https://doi.org/10.1038/s41562-025-02309-z (2025).
13. Hackenburg, K. et al. *Science* **390**, eaea3884 (2025).
14. Allen, J., Howland, B., Mobius, M., Rothschild, D. & Watts, D. J. *Sci. Adv.* **6**, eaay3539 (2020).
15. Coppock, A. Notes on attention. *alexandercoppock.com* https://alexandercoppock.com/attention.pdf (2025).

Acknowledgements

I am grateful to J. Allen and L. Hewitt for feedback on an earlier version.

Competing interests

The author declares no competing interests.