For Digital Mass Persuasion, Exposure Matters More Than Persuasiveness

Ben Tappin

London school of Economics and Political Science

b.tappin@lse.ac.uk

This paper is a preprint and has not yet undergone peer review.

The popularity of digital and social media has given rise to widespread concerns about digital mass persuasion, from Cambridge Analytica-style personalised advertising to political and health misinformation, foreign influence operations, and, recently, "deep fakes" and propaganda generated by increasingly sophisticated artificial intelligence. Mobilised by these concerns, a large body of research has studied 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 in this essay I will argue it is usually the less important part. More important is how many people are exposed to the content.

Here's the argument in a nutshell. The impact of a message depends both on how persuasive it is per person and on how many people are exposed to it. While both matter in principle, in practice exposure tends to vary far more across digital media content than persuasiveness does. As we will see, 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.

This argument has the following implication: if you are interested in understanding the mass persuasive impact of digital media content—whether from a research, policy, or practitioner perspective—you should generally pay more attention to content exposure than to its persuasiveness. This implication bears on several areas including AI-driven persuasion, misinformation, communication strategy, the scientific methods of persuasion research, and the question of which groups in society are most advantaged when it comes to mass persuasive impact on digital and social media.

The essay is split into four parts. Part 1 establishes the conceptual foundation of the argument. Part 2 describes the evidence in support. Part 3 discusses the implications and Part 4 considers potential objections and a key boundary condition.

Part 1: Conceptual Foundation

Mass persuasive impact equals persuasiveness multiplied by exposure

For someone to be persuaded by something—an advertisement, podcast, news headline, arguments from an AI chat bot, etc.—they must first be exposed to it. Therefore, we can think about the mass persuasive impact of content as a product of two things: (i) its persuasiveness—how effectively it persuades someone, conditional on them being exposed to it—and (ii) the number of people exposed to it [1].

To illustrate, suppose the persuasiveness of a political video advertisement is such that, for every 100 people exposed to the ad, 1 person is persuaded by it and changes who they vote for. In research nomenclature, this is a persuasiveness effect size of “one percentage point” (1PP). If 5 in every 100 people exposed were persuaded, the persuasiveness would be 5PP. Therefore the persuasive impact of the ad is equal to its persuasiveness multiplied by the number of people exposed to it.

If 10,000 people are exposed to a 5PP ad, this equals a persuasive impact of $0.05×10,000 = 500$ people who changed their vote. If the ad was twice as persuasive, 10PP, the impact would equal $0.1×10,000 = 1000$ changed votes. If the ad was much less persuasive, say 1PP, but many more people were exposed to it, say 1 million, it would nevertheless exert greater impact: $0.01×1,000,000 = 10,000$ changed votes.

(Brief technical aside: This formulation treats exposure as distinct people who fully attend to the content and assumes effects are approximately homogenous across people. This is a simplification; good enough for my purpose here, but the real world clearly involves some repeat/partial exposure and audience heterogeneity.)

In summary, mass persuasive impact depends on both the persuasiveness of the content (conditional on exposure) and the number of people exposed.

Which contributes more to mass persuasive impact—exposure or persuasiveness?

Both exposure and persuasiveness contribute to mass persuasive impact in principle, but they do not necessarily contribute equally in practice.

To illustrate, consider political advertising again. Suppose two ads are equally persuasive: for every 100 people who see them, five change their vote (a persuasiveness of 5PP). If one ad is shown to 1000 people and another to 10,000 people, their impacts will differ tenfold purely because of exposure. The first persuades 50 people; the second persuades 500. In this case, differences in impact are driven entirely by variability in exposure.

Now reverse the scenario. Suppose two ads each reach an audience of 10,000 people but differ in persuasiveness. One persuades 2 out of every 100 viewers (2PP), while the other persuades 10 out of every 100 (10PP). Their impacts subsequently differ by a factor of five: 200 versus 1000 people persuaded. Here, differences in impact are driven entirely by variability in persuasiveness.

These examples are deliberately simplified and unrealistic. In reality, digital media content varies along both dimensions at once. But they help clarify a crucial point: what matters most for mass persuasive impact here is whichever dimension—exposure or persuasiveness—*varies more* across content.

So, in Part 2 I examine the variability in exposure and persuasiveness across content using real data. We can then compare them to see which varies more. For this comparison to make sense, we must define a common range over which to compare them. Thus, I examine variability over the range from “typical” to “high” performing content; where typical content is that with the median level of exposure (or persuasiveness), and high performing is the 90th percentile (where possible) or above.

Part 2: The Evidence

Before looking at the data on real content, we must think about which content is relevant. On the one hand is a broad universe of digital media content: that which is not necessarily designed to be persuasive but could nevertheless persuade. On the other hand is a narrower universe of content: that which is explicitly designed to be persuasive. Focusing on one universe of content versus the other could have implications for whether exposure or persuasiveness varies more across content.

Therefore, in Part 2 I will examine content from both universes, and in Part 4 I consider how the essay's conclusion might change by focusing on one or the other content universe (spoiler: given the available data, my assessment is that the conclusion does not change dramatically but that there are data limitations).

Examining variance in exposure to content

We can categorise exposure broadly into two types: organic exposure and paid exposure. Organic exposure is obtained via user-driven attention, often mediated and amplified by digital media algorithms. In contrast, paid exposure is obtained via exchanging money for exposure, as is the case for many digital advertising campaigns. I will examine each type separately, starting with organic exposure.

It is well known that the distribution of organic exposure across digital media content is highly skewed with a long right-tail: a minority of content receives the majority of people's attention. This suggests large variance in exposure. But how large?

Public estimates of the median and other percentiles of exposure are uncommon but do exist for some digital media content.

McGrady and colleagues [2] innovated a method for estimating the distribution of views (among other features) of all public videos on YouTube, a dominant platform for video content. Their results indicate the median video receives 35 views; the 90th percentile between 1000 and 2000 views (30x to 60x the median); and the 96th percentile approximately 10,000 views (300x median).

These estimates broadly align with those of Munger and colleagues [3] who examined a large convenience sample of *political* YouTube channels (English-speaking only). Their data suggest the median total views across these channels was approximately 100,000 with a 90th percentile of somewhere between 1 million and 10 million views (10x to 100x the median).

Guinaudeau and colleagues [4] also studied a sample of political YouTube channels, focusing instead on the distribution of views *within* (not across) channels. They estimated that, on average, a channel's most viewed video received 40x more views than their median video. These authors additionally studied a comparable sample of videos from TikTok—another widely used video streaming platform—finding even greater variance there: on average, an account's most viewed video was 64x the median.

Beyond video content, the numbers appear of similar magnitude. One of the largest podcast hosting platforms on the Internet, Buzzsprout [5], publishes statistics on new episodes' downloads within their first seven days. Across all episodes, the median is 28 downloads, the 90th percentile is 450 (>15x median) and the 99th percentile is 4350 (>150x median) (note: these numbers were recorded on 2025-11-30).

Xavier [6] recently analysed global web usage patterns, estimating that 3000 websites receive 50% of all online traffic; approximately 250,000 receive a further 30% of traffic; while the remaining 20% goes to the approximate 350 million other websites in existence. Therefore, even excluding the 3000 most visited sites (which are extreme outliers), these numbers suggest that the next most visited 250,000 sites receive at least several orders of magnitude more traffic than the typical site. (To a very rough first approximation, we can calculate the expected proportion of traffic of the 250,000 sites as 0.3÷250,000 = 1.2e-6; and for the 350 million sites as 0.2÷350,000,000 = 5.7e-10; a difference of between 1000x and 10,000x. This comparison is crude and incorrect because it ignores the shape of the distribution within each category of sites, but it nevertheless illustrates the likely scale of the difference.)

For other popular platforms, such as Facebook/Instagram and X (formally Twitter), it is more difficult to find public estimates of the median alongside other percentiles of content exposure—except those that are likely to be highly unrepresentative.

For example, in 2020 Mosleh and Rand [7] identified the Twitter accounts of 816 “elites” which had been fact-checked at least three times by the organisation Politifact, including politicians, businesspeople, media organisations, and so on. The exposure across these accounts was highly variable, as indicated by their number of followers. The median number of followers was approximately 100,000 and the 90th percentile somewhere between 1 million and 10 million (10x to 100x median). However, these accounts are highly unusual and include extreme outliers in terms of follower counts, so the numbers may not generalise well to the broader Twitter ecosystem.

Beyond data such as these, it is easier to find estimates for engagement metrics such as likes and shares—which generally show a similar pattern of high skew with a long right-tail [8]. However, I do not rely on them here because engagement is not the same as exposure and can in fact be a misleading proxy for it [9].

To summarise the evidence on organic exposure so far: various sources suggest that high exposure content plausibly receives between 10x and 100x greater exposure than typical content, perhaps more.

What about *paid* exposure? Since this primarily comes from digital advertising campaigns, information on exposure often remains private and public estimates are rare or unrepresentative. Nevertheless, those that I could find paint a broadly similar picture as for organic exposure.

Gordon and colleagues [10] analysed a set of 663 ad campaigns run on Facebook in 2019/20, chosen to be representative of large scale campaigns run on the platform by US advertisers. The restriction to large scale campaigns in this set means the variance in exposure is likely an underestimate. Even so, the median campaign received 22 million “impressions” (views) against the largest which received between 200 and 500 million (10x to 25x median).

In another study [11], analysing a different set of ad campaigns on Facebook, the median impressions were 10 million while the largest were between 575 and 590 million (60x median).

Since money is exchanged for exposure, another way to examine variance in paid exposure is through variance in the money spent on advertising campaigns (assuming exposure scales strongly with spend, even if not perfectly linearly).

This evidence tells a broadly similar story.

Wernerfelt and colleagues [12] studied 70,000 ad campaigns running on the Meta platform in 2021. They report a median ad spend of $13,000 and a 90th percentile of $195,000 (15x median). Sheingate and colleagues [13] examined digital media spending by US political campaign committees; in the 2020 election cycle (their most recent data), the median spend across the 2500 committees was $31 million, with a 90th percentile of $600 million (20x median).

In summary, for both organic and paid exposure to digital media content, the evidence points toward high exposure content receiving between 10x and 100x greater exposure than typical content.

Now I will examine variance in the *persuasiveness* of content (conditional on exposure), comparing the typical (median) persuasiveness against high (90th percentile or above) persuasiveness.

Examining variance in persuasiveness of content

There are thousands of research studies estimating the persuasiveness of content. Should we calculate the median and other percentiles of persuasiveness across these studies?

No. We want to know the variance in persuasiveness across content; while different studies do test different content, they also differ in many other ways—such as the way in which they measure persuasiveness, the sample of people they expose to the persuasive content, and so on. These “nuisance” factors can affect a study’s estimate of persuasiveness. As a result, even if all *content* was equally persuasive (zero variability in persuasiveness), we could still see variability in persuasiveness across studies. This is a problem.

To avoid being misled in this way, we will examine variance in the persuasiveness of content tested *within the same study*. By design, this holds fixed all of the aforementioned nuisance factors.

There are two additional study design factors to consider. First, the gold standard study design for estimating persuasiveness is a randomised experiment [14], where people are randomly assigned to treatment content or a control group. Accordingly, we will examine randomised experiments only. Second, recall that persuasiveness is *conditional on exposure*. Therefore, we will examine randomised *survey* experiments, because in surveys people's attention is purchased in exchange for money (as opposed to field experiments where it is not and many people will likely ignore the content).

To sum up the type of evidence we require: randomised survey experiments that estimate the persuasiveness of (ideally) a large number of different pieces of content. Fortunately, such evidence exists in the form of survey "megastudies" [15].

I am aware of five such megastudies, testing content related to climate change, electoral advertising, political tolerance, tax compliance, and vaccine scepticism. To spoil the reveal: together these studies indicate that highly persuasive content is at most 10x the persuasiveness of typical content. I will now describe the studies.

Hewitt and colleagues [16] estimated the distribution of persuasiveness across 617 real political ads tested in survey experiments through the 2018 and 2020 US election cycles. Among the ads tested in 2018, the median persuasiveness was approximately 2.5 percentage points (PP) against a 95th percentile of 5.5PP (2-3x the median); among 2020 downballot ads, it was 1.3PP against 2.7PP (2x median); and, among 2020 presidential ads, 0.9PP against 1.8PP (2x median).

An important technical note is that these estimates assume persuasiveness follows a particular shaped distribution across ads. This has an advantage and disadvantage for understanding the variance. The disadvantage is that persuasiveness may in truth follow a different distribution, one with larger variance. In that case, these numbers will underestimate the variance. On the other hand, the advantage is that the estimates account for the statistical noise inherent to experimentation. Without accounting for this noise, the variance in persuasiveness would be overestimated. In general, the overestimation will be larger the fewer people were exposed to each piece of content in the experiment. For the remainder of survey megastudies described below, I rely on the raw persuasiveness estimates that do *not* assume a particular shaped distribution. Therefore the variance is likely overestimated somewhat. I will point out when this overestimation is likely to be particularly severe.

Vlasceanu and colleagues [17] estimated the persuasiveness of 11 different messages encouraging action on climate change in a sample of 60,000 participants from 63 countries. On average across climate belief and policy support outcomes, the median persuasiveness was approximately 0.6PP against a maximum of 2.5PP (4-5x the median).

Voelkel and colleagues [18] tested 25 interventions to reduce antidemocratic attitudes and partisan animosity in a sample of 32,000 Americans. On average across their three outcomes of interest, the median persuasiveness was approximately 1.5PP; the 95th percentile 3PP (2x median); and the maximum 6.3PP (4x median).

Finally, Zickfield and colleagues [19] estimated the effect of 21 different types of “honesty oath” on 21,000 American participants’ behaviour in a tax evasion game, where behaving dishonestly could earn the participants additional money. The median oath reduced dishonesty by approximately 4PP; the 90th percentile by 7PP (1.5-2x the median); and the maximum by 8PP (2x median).

These studies indicate that highly persuasive content is between 2x and 5x more persuasive than typical content. However, they share an important limitation: in all cases, the content was explicitly designed and selected—often by domain experts—to maximise persuasiveness. Therefore, even though experts seemingly struggle to identify highly persuasive content [20], we might still worry that this selection mechanism would exclude some of the least persuasive content out there from making it into the test. If so, this could artificially restrict the variance in persuasiveness. Fortunately, there is a final megastudy we can examine where the content tested *wasn’t* selected to maximise persuasiveness, avoiding this worry.

Allen and colleagues [1] tested a total of 130 Covid vaccine-related headlines on 19,000 Americans' vaccination intentions, where the majority of the headlines were systematically sampled from Facebook to be balanced across topics and the quality of the originating web domain. According to data collected by the authors, approximately 100 of the headlines had a vaccine-sceptical flavour, implying the vaccine was harmful to health (e.g., "Lisa Shaw: presenter's death due to complications of Covid vaccine"). Among these 100 headlines, the median reduction in vaccination intentions they caused among the participants was approximately 0.6PP; the 95th percentile was 4PP (7x median); and the maximum 6PP (10x).

These numbers imply greater variance in persuasiveness than the above studies, likely in part because the content was not selected to be persuasive. However, there is an important caveat. Unlike the other studies, where each piece of content was seen by 1000+ participants, in the vaccine study each piece of content was seen by far fewer participants—perhaps as few as 100 in some cases. As a result, the variance in persuasiveness will be overestimated to a greater extent than in the other studies, for the statistical reason I outlined earlier: noise inherent to experimentation. All things considered, then, this suggests that the figure of 10x median for highly persuasive content should be treated as something of an overestimate.

In summary, the evidence reviewed in this section suggests that highly persuasive content is somewhere between 2x and 10x more persuasive than typical content.

Part 3: Implications

For digital mass persuasion, exposure matters more than persuasiveness

We've seen that the variability in exposure to digital media content is generally larger than the variability in the persuasiveness of content. As a result, typically-persuasive content that receives high exposure will generally be more impactful than highly persuasive content that receives only typical exposure. In this sense, exposure matters more than persuasiveness for mass persuasive impact.

To illustrate this clearly we can put concrete numbers on it.

Suppose the typical exposure is 100 people and the typical persuasiveness is 1PP (these particular numbers are arbitrary and the point below does not depend on them; choose your own numbers if you wish). Using the gaps between high-performing content and typical content—established in the previous section—we can calculate the persuasive impact for different content by multiplying exposure and persuasiveness. Doing so, we see that typically-persuasive content that receives high exposure is more impactful (10-100 people persuaded) than highly persuasive content that receives only typical exposure (2-10 people persuaded) (table below, red text).

Now let us consider some specific implications of this argument.

		Persuasiveness	
		Typical (1PP)	**High** (2x-10x typical)
Exposure	**Typical** (100 people)	1 person persuaded	**2-10 people**
	High (10x-100x typical)	**10-100 people**	20-1000 people

Specific implications

I will discuss several distinct implications. But in each case they boil down to the same idea: if you are interested in understanding the impact of digital mass persuasion—whether from a research, policy, or practitioner perspective—you should generally pay more attention to content exposure than to its persuasiveness.

Misinformation

In the wake of the 2016 election in the United States and the EU referendum in Britain, there was an explosion of concern about the impact of so-called “fake news”, outright falsehoods propagated (in large part) on digital and social media.

Research on the impact of such content proceeded apace, under the banner of misinformation research, with renewed impetus during the COVID-19 pandemic as false claims about the vaccine (etc.) spread via social media. Studies showing that exposure to such content could influence people's beliefs (i.e., that it was persuasive) appeared to validate concerns about online misinformation.

However, converging evidence suggests that people's exposure to outright falsehoods is in fact quite limited relative to other types of content online [21–23]. This is perhaps because falsehoods are often spread by fringe websites that, while unconstrained by journalistic norms against lying, do not enjoy a large public following. As a result, the mass persuasive impact of outright falsehoods may often be limited.

Nevertheless, there are other types of content beyond outright falsehoods that we might also consider misinformation. One example is claims that are technically true yet misleading (for instance: claims that report a true event but omit important context). Such claims are more likely than outright falsehoods to appear in mainstream media—by virtue of their not clearly violating norms against lying—and thereby to receive greater exposure online.

One implication of this essay is that, owing to its greater exposure, technically-true-yet-misleading claims could therefore be much more impactful than outright falsehoods—despite them seeming less egregious to us and even if they were less persuasive than falsehoods.

Research supports this implication.

In their study testing vaccine sceptical headlines (described earlier), Allen and colleagues [1] combined the exposure each headline received on Facebook with its persuasiveness as estimated in their experiment. They found that the most egregiously inaccurate headlines—those explicitly flagged as misinformation by fact checkers—were more persuasive than the unflagged headlines (many of which came from mainstream sources). Yet, the unflagged headlines received much wider exposure. As a result, the less persuasive unflagged headlines had an estimated 46x greater impact on reducing Facebook users' vaccination intentions than did the content flagged as misinformation.

Given the above points, a wider implication of this essay and Allen and colleagues' work is to reorient the field of misinformation research away from a veracity-first lens (is it false?) to an impact-first lens (how persuasive x how many people were exposed?). The most societally impactful misinformation may be not overt falsehoods, but technically-true-yet-misleading content which, while less egregious/persuasive, receives greater exposure online.

While this re-orientation is appealing, it is important to note it could pose deep challenges for the study of misinformation and how society should combat it.

For example, how do we determine whether content is misleading? One definition might be: content that moves people's beliefs away from the truth. But, as the philosopher Daniel Williams has argued, such content is pervasive in human communication. As a result, this definition would expand what content counts as misinformation to an extent that few people would endorse.

Another definition of misleading content might be: claims about reality that omit important context. But how do we decide what's important? The complexity of reality means it is impossible to provide objectively complete context for any claim. This makes it a matter of subjective judgement what context to include, which inevitably invokes people's values (which can differ markedly). Whether these challenges can be sufficiently overcome is unclear and will not be resolved here.

AI-driven persuasion

There is a great deal of concern that advanced artificial intelligence will become superhumanly persuasive and thereby exert powerful influence over public opinion and behaviour. Mobilised by this, a large body of research now exists that exposes people to AI content and measures its persuasiveness [24–30]. This research has provided valuable insights, advancing understanding about the persuasiveness of AI-driven personalisation; its different rhetorical strategies; the persuasive returns to AI model post-training and increasing model size; how (in)accurate its persuasive claims are, and more.

However, an implication of this essay is that none of these "levers" may be the most important for understanding the potential mass persuasive impact of AI. It doesn't matter how persuasive a 10-minute AI conversation is if few people will talk to it that long. The highest value lever could instead be AI's capability to be superhumanly *engaging*, unlocking reliably wider exposure to persuasive content.

For researchers and policymakers interested in understanding the mass persuasive impact of AI, this suggests a re-orientation of the field away from asking “how persuasive is AI?” to “how much exposure can AI generate for persuasive content?”

Methods of persuasion research

The implications of this essay for understanding misinformation and AI persuasion point to a broader implication for the methods of persuasion research.

Researchers have a well established paradigm for studying the persuasiveness of content; the survey experiment, in which people are randomly exposed to treatment or control content, and subsequently have their attitudes and beliefs measured [31]. As we’ve seen, this design suits the study of persuasiveness because it conditions on exposure—typically by paying people money in exchange for their attention. However, this means it is less suited to studying the more important determinant of mass persuasive impact: exposure itself.

Indeed, for this reason, using survey experiments to infer mass persuasive impact can be actively misleading. For example, in Allen and colleagues’ [1] survey experiment on vaccine sceptical headlines, flagged misinformation was more persuasive than unflagged content. But concluding that flagged misinformation is therefore more impactful would have been a mistake. On the contrary, the unflagged content was estimated to have much greater impact owing to its wider exposure.

Related mistakes can be made by confusing estimates of persuasiveness (in survey experiments) with estimates of persuasive impact.

For instance, the average political ad has a persuasiveness of between 1 and 3 percentage points in many survey experiments [16]. Real world elections are frequently decided by such margins. Does this imply that the average political ad is frequently deciding elections? No! The reason why not is that most political ads likely receive very limited exposure, dramatically constraining their impact (ask yourself: when was the last time you watched a political ad from beginning to end?). Even when ads do receive widespread exposure, so too may those from a competing campaign—cancelling impact out on aggregate. It is a mistake to naively generalise estimates of persuasiveness from survey experiments to estimates of persuasive impact.

This implies that researchers interested in persuasive impact need to establish a design for simultaneously studying *exposure* to persuasive content.

One approach could be that of Allen and colleagues [1], who sampled real content with known exposure and could therefore incorporate this exposure data with their estimates of persuasiveness. However, this design does not provide for experimental control over exposure, making it less suited for understanding what causes some content to receive wider exposure than other content.

The political scientist Alex Coppock [31] recently reviewed designs that might solve this problem. He notes the move of some researchers toward using "faux feed" designs, which simulate the real feeds of social media while providing for full experimental control and range of measurements of attention/exposure as well as persuasion [32]. The obvious downside of these designs is that the feed isn't real, which may induce artificial patterns of attention/exposure.

The best of all worlds would of course be experiments on real social media. However, this is often expensive, can raise ethical issues, and may require collaboration with social media companies—and face a range of constraints for that reason.

Clearly no design is perfect. But what is clear is that researchers interested in understanding the persuasive impact of digital media content should systematically incorporate exposure into their studies of persuasiveness.

Communication strategy

This essay argues that the mass persuasive impact of digital content generally depends more on it achieving high (versus typical) exposure than on it being highly (versus typically) persuasive. For communicators who seek to maximise such impact, this suggests they should dedicate more resources—time, energy, money—to achieving high exposure than to crafting highly persuasive content.

There is some recent research that supports this suggestion when it comes to paid exposure. As described earlier, Hewitt and colleagues [16] estimated the variance in persuasiveness of political ads created by real campaigns in the US. They then performed simulations to understand the optimal amount of money a campaign should spend on ad-testing (to find more persuasive ads) versus disseminating those ads to voters (exposure) to maximise impact. Their simulations indicate that a minority of a campaign's budget (between 10% and 20%) should go to ad-testing, while the rest (the majority) should go to dissemination.

When it comes to organic exposure, the picture is a little more complicated.

While achieving high organic exposure may be more important for impact, this says nothing about how much *control* communicators have over achieving it. The factors that determine whether content achieves high organic exposure may be less controllable than the factors that make content highly persuasive, especially because organic exposure is often mediated by opaque algorithms. Therefore, it could be that increasing persuasiveness is in fact more tractable for communicators than increasing organic exposure.

Moreover, even if a communicator "cracks the code" for achieving consistently high exposure, this may only be temporary. On digital and social media there is ferocious competition for a finite pool of attention—meaning that others would soon copy or crack the code, potentially undermining the initial exposure advantage.

In my view, it nevertheless remains plausible that obtaining higher exposure (90th percentile) than typical is at least somewhat tractable for communicators, even if "going viral" (extreme outlier exposure) is not. Indeed, more systematically incorporating the study of exposure into persuasion research (see previous section) could help identify the most important factors for reliably achieving high exposure.

The above points have additional implications for specific debates between practitioners about how to maximise persuasive impact. I will briefly mention two.

In liberal and progressive circles in the US, there has been fierce disagreement over the strategy of *popularism* (not to be confused with populism); the idea that, to win the most votes, candidates for election should talk about issues that are popular, and shy away from those that are less popular.

One way in which proponents of popularism determine which issues are most popular is by using survey experiments to measure the persuasiveness of a candidate talking about different issues. If the candidate talks about reducing the cost of living or protecting Social Security benefits, for example, more participants in the survey might say they would vote for them than if the candidate instead talked about immigration or transgender issues. Therefore, to maximise their vote share (persuasive impact), popularism concludes the candidate should talk about the former set of issues.

As we've seen in this essay, however, saying highly persuasive things is just one part of a message's persuasive impact; achieving high exposure to the message is plausibly the more important part (on digital media at least).

While pocketbook issues like the cost of living or Social Security may be most persuasive conditional on exposure, they could be considered more boring or milquetoast by voters than other—perhaps more culture war-y—issues like immigration or transgender rights. If so, messages on the latter types of issues could receive systematically wider attention and thus exposure on digital media.

If this exposure advantage were large enough, candidates peddling the less persuasive message could in fact achieve greater persuasive impact overall.

Admittedly, this is mostly speculation on my part—I'm not privy to any special data relating persuasiveness to exposure for different types of issue-messaging, and the specific numbers really do matter. But the possibility strikes me as at least somewhat plausible and follows naturally from this essay's argument.

Another live debate surrounds the persuasive impact of non-violent radical protest, such as the throwing of soup over artwork by the group Just Stop Oil.

Protest actions like this are frequently written off as unpersuasive and even counterproductive. However, they can undeniably achieve very wide exposure for a particular message. As a result, even if a message is persuasively weak, it could still achieve much greater impact overall than a well-crafted and compelling message delivered via less disruptive means that consequently receives less exposure. Of course, if the message is not weak but in fact *negatively* persuasive conditional on exposure—that is, it backfires, moving people away from the intended belief—then achieving wide exposure would be very counterproductive for persuasive impact.

Either way, this essay highlights that the non-violent protest action itself can be a powerful vehicle for achieving wide exposure to a message and thus for achieving mass persuasive impact. (James Ozden spells the logic out in more detail in this post.)

The distribution of mass persuasive impact

A final implication of this essay concerns *who* in society is better positioned to generate mass persuasive impact across digital media. If exposure generally matters more for impact than persuasiveness does, then actors whose comparative advantage lies in achieving high exposure—rather than in crafting highly persuasive content—will disproportionately benefit when it comes to mass persuasive impact.

Several groups seem advantaged here.

Wealthy actors, such as corporations, large political campaigns, or other well-resourced organisations can buy large amounts of paid exposure. Insofar as paid exposure scales strongly with spend (albeit not perfectly linearly), these actors could reliably secure high exposure volumes that exceed what most others can obtain, even if their content is not especially persuasive.

Owners of digital and social media platforms occupy another advantageous position: by virtue of controlling ranking systems, visibility policies, and other aspects of platform architecture, in principle they could guarantee high exposure and thus disproportionate persuasive impact even for mediocre persuasive content.

A third advantaged group are actors willing to optimise for exposure by ruthlessly exploiting widespread attentional preferences, perhaps at the expense of other things. For example, research suggests that content which is morally outraged, emotionally negative, or identity-affirming may attract greater attention and exposure than other content [33–35]. Actors willing and able to exploit this could generate greater persuasive impact than their counterparts who are less willing or able. This is not to present a false dichotomy; in principle, content can of course be engaging at the same time as being morally and emotionally balanced and nuanced. But the willingness and ability to ruthlessly exploit widespread attentional preferences is likely advantageous for mass persuasive impact.

In sum, when it comes to mass persuasive impact, this essay suggests digital media rewards advantages in exposure more than advantages in persuasive skill. It is not just that greater exposure equals greater impact—this would be the case even if exposure and persuasiveness mattered similarly for impact. Rather, it is that actors whose comparative advantage lies in exposure are *disproportionately* benefited compared with actors whose comparative advantage lies in persuasion.

As a result, those with money, structural control over content distribution, and those willing and able to ruthlessly optimise content for attention seem disproportionately benefited. In contrast, those whose comparative strength lies more in the building of persuasive arguments—such as subject-matter experts, policy specialists, or long-form journalists/writers—may be systematically disadvantaged in generating mass persuasive impact, even if their messages are more compelling per exposure.

Part 4: Objections And A Boundary Condition

In this final section, I will consider some potential objections to the essay's argument and highlight a key boundary condition. Objections first.

What is the relevant universe of content?

We can think of digital media content as coming from different "universes." On the one hand is a broad universe: content which is not necessarily designed to be persuasive but could nevertheless persuade. On the other hand is a narrower universe: content which is explicitly designed to persuade.

Much of the exposure data I rely on arguably comes from the first universe, while much of the persuasiveness data comes from the second universe. Therefore, one objection could be that I'm comparing the variability of exposure and persuasiveness across different universes of content, making the comparison invalid.

This is a reasonable objection. For example, perhaps the variability in exposure is systematically smaller among content explicitly designed to persuade. Conversely, perhaps the variability in persuasiveness is larger among content not designed to persuade. In these cases, the variability of exposure could shrink relative to persuasiveness; if it shrunk by enough, it could undermine the exposure-first thesis.

To consider the strength of this objection, we can therefore examine data specifically from these cases.

First, let's examine data on the variability in exposure among content designed to persuade. The data on *paid* exposure arguably belongs in this category, since it comes from digital advertising campaigns whose explicit goal is likely to persuade. The data I reviewed in Part 2 suggests high paid exposure is between 10x and 60x that of typical exposure. Most of the data on *organic* exposure I reviewed likely falls outside this category, since it is not necessarily designed to persuade. The organic exposure data from political YouTube channels could be an exception, but we would simply be guessing their goal was to persuade, which could well be wrong. (That data suggests high exposure is between 10x and 100x the typical exposure.)

Now data from the other category: the variability in persuasiveness among content not necessarily designed to persuade. The only data we have there is the study from Allen and colleagues testing vaccine-sceptical headlines, whose estimates suggest highly persuasive content is 7x to 10x more persuasive than typical content.

In summary, given the data available here, it doesn't seem that this objection could overturn the exposure-first thesis. The magnitudes of variability in exposure and persuasiveness seem broadly similar overall as in these specific cases. Nevertheless, it is important to emphasise that we don't have a great deal of data to rely on in these cases—and essentially none that can speak to the variability in organic exposure among content designed to persuade. This is a key limitation of the essay.

What if within-topic variance is smaller than across-topic variance?

The variance in exposure data that I draw upon in Part 2 is mostly *across* topics—for example, all public YouTube videos, which covers videos on cooking, politics, sport, music, health, etc. In contrast, the variance in persuasiveness data is *within* topic—for example, political advertising, *or* climate communication, *or* vaccination misinformation etc.

An objection could be that within-topic variance in exposure is smaller than across-topic variance in exposure. If this were true, it would shrink the importance of exposure over persuasiveness for impact within-topic; if it shrunk by enough, it could undermine the exposure-first thesis. While this objection cannot be ruled out, two pieces of the exposure data speak against it.

On organic exposure first: the study by Guinaudeau and colleagues [4] looked only at political video accounts on YouTube and TikTok. Moreover, their estimates of variance in exposure are *within* account (i.e., they compare high exposure to typical exposure videos from the same account). This gets us closer to within-topic variance. Despite this, their estimates of variance are very similar to the estimates I arrive at overall: high exposure was 40x the typical exposure on YouTube and 64x on TikTok.

On paid exposure too: the data from Sheingate and colleagues [13] I cited concerned political campaign committee spending on digital media within a single US election cycle (2020)—getting us closer to within-topic—and suggests high spending is 20x greater than typical spending. This number is also broadly in keeping with my overall estimates for variance in exposure.

In sum, while it's possible that within-topic variance in exposure is much smaller than across-topic variance in exposure, I don't see clear evidence of this here.

What if variance across content format is larger than within format?

To examine the variance in persuasiveness across content, I drew upon survey megastudies—experiments which tested a large sample of different content within the same study. This was necessary to hold fixed nuisance factors that would have polluted our estimate of the variance otherwise. However, a downside of these megastudies is that the content tested by each tended to be of the same format. For instance, they tested *either* brief passages of text, *or* videos, *or* headlines, and so on. None of them tested content of entirely different formats within the same study.

Therefore, an objection could be that the variance in persuasiveness across content of different formats—such as text versus video versus conversation—is larger than across content of the same format. If so, this could mean that estimates based on within-format data understate the possible variance in persuasiveness.

This is a reasonable objection. Nevertheless, the results of studies that directly compare the persuasiveness of content in different formats suggest it is not a major concern.

Hackenburg and colleagues [30] found that the persuasiveness of a lengthy conversation with AI (e.g., GPT-4.5) was approximately 1.5-2x greater than that of exposure to a static

200-word message written by AI. Argyle and colleagues [24] did not detect any persuasive differences between AI conversation (GPT-4) and static text. Lin and colleagues [29] qualitatively compared the persuasiveness of AI models (e.g., GPT-4o) against that of video ads from other studies, which showed a difference in favour of AI conversation of between 2x and 6x. Chen and colleagues [36] found minimal persuasive differences between video ads and conversation with Anthropic's Claude 3.5 Sonnet. Finally, Wittenberg and colleagues [37] compared 72 political and public health ads in video versus text-only format, finding that video was approximately 2x more persuasive on average.

In summary, while it is plausible that there are moderate or even large differences in persuasiveness across content of alternative formats, these differences generally appear in keeping with the variance in the persuasiveness across content of the same format reviewed in Part 2.

What if absolute levels of exposure are overestimated?

Digital platforms typically count a "view" or "impression" using loose criteria. For instance, a video may register as "viewed" after only a brief fraction of play time. This means platform-reported exposure often overstates the amount of attention people actually pay to content.

This matters a lot for quantifying persuasive impact. If, for example, a video ad's persuasiveness is estimated assuming participants watched the full clip in an experiment, then multiplying that persuasiveness by platform-reported "views" (many of which reflect only partial attention) will inflate estimates of persuasive impact.

Therefore, this is a very important objection for anyone attempting to quantify the mass persuasive impact of digital media content. However, it has limited bearing on the argument of this essay. The thesis depends on *relative* exposure—that some content receives 10x or 100x the exposure of typical content—not on absolute exposure. Whatever the precise threshold platforms use for counting a "view," it presumably applies consistently across content. As a result, even if the absolute levels are inflated, the ratios between high-exposure and typical-exposure content will remain broadly intact. For this reason, it does not undermine the core claim: that differences in exposure are generally larger than differences in persuasiveness, and thus disproportionately matter for mass persuasive impact.

What if exposure and persuasiveness are negatively correlated?

It is possible that exposure and persuasiveness are negatively correlated: content that achieves high exposure is systematically less persuasive than content that receives more limited exposure.

There are good reasons to take this possibility seriously. Prominent theories of persuasion emphasise that attitude change is strongest when people engage deeply and systematically with a message, rather than relying on shallow heuristics or cues [38]. Consistent with this, some of the largest persuasive effects observed in experiments arise from lengthy, interactive conversations—whether with other people or with AI systems—which demand sustained attention and engagement [26,30,39].

By contrast, much high exposure digital content, such as short videos, ads, or headlines, is designed for rapid consumption and minimal cognitive effort. Because people are often effort-averse (especially in digital environments optimised for scrolling and multitasking) content that spreads widely may be that which demands shallower engagement and therefore is less persuasive. If so, exposure and persuasiveness would be negatively correlated across content. This could potentially undermine the claim that exposure matters more for mass persuasive impact.

The practical force of this objection depends on the *strength* of the negative relationship. The core comparison in this essay concerns the gap between content with typical and high exposure versus the gap between typical and high persuasiveness. As reviewed in Part 2, for exposure this gap is likely a factor of between 10x and 100x—for simplicity let's say 50x—whereas for persuasiveness it is between 2x and 10x—for simplicity, 5x. Therefore, high exposure content would need to be systematically less persuasive by roughly one order of magnitude to close this difference (50x versus 5x) and overturn the exposure-first thesis.

The evidence does not suggest a persuasion penalty this strong.

For example, as described earlier, the difference in persuasiveness between content of alternative formats—particularly those most likely to receive different exposures owing to their different demands on attention and engagement (e.g., conversation versus video versus text)—is generally less than 10x.

The importance of “mass” in mass persuasive impact

Finally, a key boundary condition of this essay is that it concerns *mass* persuasive impact, in which impact is primarily determined by the numbers of people persuaded. But persuasive impact can be conceived of in other ways—and importantly can be large even when exposure is tiny.

For example, if the target audience of persuasion is small but highly leveraged—political leaders, senior civil servants, judges, regulators, journalists, corporate executives etc.—then persuading a few individuals could produce outsized downstream impact. In these high-leverage contexts, exposure is small by definition, and the persuasiveness of content can matter far more for impact.

A concrete example comes from current concerns about AI-driven persuasion. Many AI persuasion fears focus on its capability to influence public beliefs and behaviour at scale; in that case, exposure remains central to its persuasive impact. But a different class of risk involves targeted persuasion of key decision makers. For instance, an AI system that is able to win trust and influence the judgement of a small set of politicians, business executives, or military personnel. Here, the number exposed may be small by definition, yet the downstream impact could nevertheless be large.

This is best understood as “elite” persuasion rather than mass persuasion, and does not contradict the exposure-first thesis of this essay; it simply falls outside its intended scope.

Acknowledgements

I am grateful to Jennifer Allen and Luke Hewitt for feedback on an earlier version of this paper. This does not imply they endorse everything written, nor are they responsible for any errors.

References

1. Allen, J., Watts, D. J. & Rand, D. G. Quantifying the impact of misinformation and vaccine-skeptical content on Facebook. *Science* **384**, eadk3451 (2024).

2. McGrady, R., Zheng, K., Curran, R., Baumgartner, J. & Zuckerman, E. Dialing for Videos: A Random Sample of YouTube. *J. Quant. Descr. Digit. Media* **3**, (2023).

3. Munger, K., Hindman, M., Yalcin, O., Phillips, J. & Bisbee, J. Pressing Play on Politics: Quantitative Description of YouTube. *J. Quant. Descr. Digit. Media* **5**, (2025).

4. Guinaudeau, B., Munger, K. & Votta, F. Fifteen Seconds of Fame: TikTok and the Supply Side of Social Video. *Comput. Commun. Res.* **4**, 463–485 (2022).

5. Buzzsprout. Podcasting Industry Stats. https://www.buzzsprout.com/stats.

6. Xavier, H. S. The Web unpacked: a quantitative analysis of global Web usage. in *Proceedings of the 20th International Conference on Web Information Systems and Technologies* 183–190 (2024). doi:10.5220/0012905900003825.

7. Mosleh, M. & Rand, D. G. Measuring exposure to misinformation from political elites on Twitter. *Nat. Commun.* **13**, 7144 (2022).

8. Narayanan. Understanding Social Media Recommendation Algorithms. *Knight First Amendment Institute* http://knightcolumbia.org/content/understanding-social-media-recommendation-algorithms.

9. Barrie, C., Atalay, A. & ElKattan, A. Information diets are more diverse in attention than in engagement. Preprint at https://doi.org/10.31219/osf.io/kvbzd_v1 (2025).

10. Gordon, B. R., Moakler, R. & Zettelmeyer, F. Close Enough? A Large-Scale Exploration of Non-Experimental Approaches to Advertising Measurement. *Mark. Sci.* **42**, 768–793 (2023).

11. Gordon, B. R., Zettelmeyer, F., Bhargava, N. & Chapsky, D. A Comparison of Approaches to Advertising Measurement: Evidence from Big Field Experiments at Facebook. *Mark. Sci.* **38**, 193–225 (2019).

12. Wernerfelt, N., Tuchman, A., Shapiro, B. T. & Moakler, R. Estimating the Value of Offsite Tracking Data to Advertisers: Evidence from Meta. *Mark. Sci.* **44**, 268–286 (2025).

13. Sheingate, A., Scharf, J. & Delahanty, C. Digital Advertising in U.S. Federal Elections, 2004-2020. *J. Quant. Descr. Digit. Media* **2**, (2022).

14. Coppock, A. *Persuasion in Parallel*. (University of Chicago Press, Chicago, 2022).

15. Duckworth, A. L. & Milkman, K. L. A guide to megastudies. *PNAS Nexus* **1**, pgac214 (2022).

16. Hewitt, L. *et al.* How Experiments Help Campaigns Persuade Voters: Evidence from a Large Archive of Campaigns' Own Experiments. *Am. Polit. Sci. Rev.* 1–19 (2024) doi:10.1017/S0003055423001387.

17. Vlasceanu, M. *et al.* Addressing climate change with behavioral science: A global intervention tournament in 63 countries. *Sci. Adv.* **10**, eadj5778 (2024).

18. Voelkel, J. G. *et al.* Megastudy testing 25 treatments to reduce antidemocratic attitudes and partisan animosity. *Science* **386**, eadh4764 (2024).

19. Zickfeld, J. H. *et al.* Effectiveness of ex ante honesty oaths in reducing dishonesty depends on content. *Nat. Hum. Behav.* **9**, 169–187 (2025).

20. Broockman, D. E., Kalla, J. L., Caballero, C. & Easton, M. Political practitioners poorly predict which messages persuade the public. *Proc. Natl. Acad. Sci.* **121**, e2400076121 (2024).

21. Grinberg, N., Joseph, K., Friedland, L., Swire-Thompson, B. & Lazer, D. Fake news on Twitter during the 2016 U.S. presidential election. *Science* **363**, 374–378 (2019).

22. Allen, J., Howland, B., Mobius, M., Rothschild, D. & Watts, D. J. Evaluating the fake news problem at the scale of the information ecosystem. *Sci. Adv.* **6**, eaay3539 (2020).

23. Altay, S., Nielsen, R. K. & Fletcher, R. Quantifying the "infodemic": People turned to trustworthy news outlets during the 2020 coronavirus pandemic. *J. Quant. Descr. Digit. Media* **2**, (2022).

24. Argyle, L. P. *et al.* Testing theories of political persuasion using AI. *Proc. Natl. Acad. Sci.* **122**, e2412815122 (2025).

25. Hackenburg, K. *et al.* Scaling language model size yields diminishing returns for single-message political persuasion. *Proc. Natl. Acad. Sci.* **122**, e2413443122 (2025).

26. Costello, T. H., Pennycook, G. & Rand, D. G. Durably reducing conspiracy beliefs through dialogues with AI. *Science* **385**, eadq1814 (2024).

27. Bai, H., Voelkel, J. G., Muldowney, S., Eichstaedt, J. C. & Willer, R. LLM-generated messages can persuade humans on policy issues. *Nat. Commun.* **16**, 6037 (2025).

28. Hackenburg, K. & Margetts, H. Evaluating the persuasive influence of political microtargeting with large language models. *Proc. Natl. Acad. Sci.* **121**, e2403116121 (2024).

29. Lin, H. *et al.* Persuading voters using human–artificial intelligence dialogues. *Nature* **648**, 394–401 (2025).

30. Hackenburg, K. *et al.* The levers of political persuasion with conversational artificial intelligence. *Science* **390**, eaea3884 (2025).

31. Coppock, A. *Notes on Attention*. https://alexandercoppock.com/attention.pdf (2025).

32. Epstein, Z., Lin, H., Pennycook, G. & Rand, D. Quantifying attention via dwell time and engagement in a social media browsing environment. Preprint at https://doi.org/10.48550/arXiv.2209.10464 (2022).

33. Brady, W. J., Gantman, A. P. & Van Bavel, J. J. Attentional capture helps explain why moral and emotional content go viral. *J. Exp. Psychol. Gen.* **149**, 746–756 (2020).

34. Brady, W. J., Wills, J. A., Jost, J. T., Tucker, J. A. & Van Bavel, J. J. Emotion shapes the diffusion of moralized content in social networks. *Proc. Natl. Acad. Sci.* **114**, 7313–7318 (2017).

35. Williams, D. The marketplace of rationalizations. *Econ. Philos.* 1–25 (2022) doi:10.1017/S0266267121000389.

36. Chen, Z. *et al.* A Framework to Assess the Persuasion Risks Large Language Model Chatbots Pose to Democratic Societies. Preprint at https://doi.org/10.48550/arXiv.2505.00036 (2025).

37. Wittenberg, C., Tappin, B. M., Berinsky, A. J. & Rand, D. G. The (minimal) persuasive advantage of political video over text. *Proc. Natl. Acad. Sci.* **118**, (2021).

38. Petty, R. E. & Cacioppo, J. T. The Elaboration Likelihood Model of Persuasion. in *Communication and Persuasion: Central and Peripheral Routes to Attitude Change* (eds Petty, R. E. & Cacioppo, J. T.) 1–24 (Springer, New York, NY, 1986). doi:10.1007/978-1-4612-4964-1_1.

39.Broockman, D. & Kalla, J. Durably reducing transphobia: A field experiment on door-to-door canvassing. *Science* **352**, 220–224 (2016).