



The Oxygen of Amplification

Better Practices for Reporting on Extremists, Antagonists, and Manipulators Online

By Whitney Phillips

Data&Society

EXECUTIVE SUMMARY

MAPPING THE MEDIA ECOSYSTEM

We live in a time where new forms of power are emerging, where social and digital media are being leveraged to reconfigure the information landscape. This new domain requires journalists to take what they know about abuses of power and media manipulation in traditional information ecosystems and apply that knowledge to networked actors, such as white nationalist networks online. These actors create new journalistic stumbling blocks that transcend attempts to manipulate reporters solely to spin a beneficial narrative – which reporters are trained to decode – and instead represent a larger effort focused on spreading hateful ideology and other false and misleading narratives, with news coverage itself harnessed to fuel hate, confusion, and discord.

The choices reporters and editors make about what to cover and how to cover it play a key part in regulating the amount of oxygen supplied to the falsehoods, antagonisms, and manipulations that threaten to overrun the contemporary media ecosystem—and, simultaneously, threaten to undermine democratic discourse more broadly. This context demands that journalists and the newsrooms that support them examine with greater scrutiny how these actors and movements endeavor to subvert journalism norms, practices, and objectives. More importantly, journalists, editors, and publishers must determine how the journalistic rule set must be strengthened and fortified against this newest form of journalistic manipulation—in some cases through the rigorous upholding of long-standing journalistic principles, and in others, by recognizing which practices and structural limitations make reporters particularly vulnerable to manipulation.

With a particular focus on coverage of internet trolls, conspiracy theories, and networks of white nationalists during and after the 2016 US presidential election, this report explores these issues through the perspectives of those who must navigate this territory every day: the journalists themselves. The report's three parts incorporate interviews with 50 individuals with intimate knowledge of the contemporary news media. Fifty-six percent of these respondents are women, 30% are people of color, and 26% are natural-born citizens of countries outside the United States, with additional insights gleaned from the scores of the more informal discussions the author – a frequent expert commentator on stories about internet trolling – has had with reporters since 2010. While each part may be read on its own, each informs and is informed by the others.

PART ONE

“**In Their Own Words: Trolling, Meme Culture, and Journalists’ Reflections on the 2016 US Presidential Election,**” provides a historical overview of the relationship between the news media and far-right manipulators during the 2016 US presidential election.

PART TWO

“**‘At a Certain Point You Have to Realize That You’re Promoting Them’: The Ambivalence of Journalistic Amplification,**” identifies the intended and unintended consequences of reporting on bigoted, damaging, or otherwise problematic information and the structural limitations of journalism (economic, labor, and cultural) that exacerbate these tensions; and

PART THREE

“**The Forest and the Trees: Proposed Editorial Strategies,**” recommends practices on establishing newsworthiness; handling objectively false information; covering specific harassment campaigns or manipulators, bigots, and abusers; and reporting on the internet that are particularly critical in an era of disinformation.

Author: Whitney Phillips; PhD 2012, English with an emphasis on folklore, University of Oregon

This report is funded by Craig Newmark Philanthropies and News Integrity Initiative, with additional programmatic and general support from other Data & Society funders.

For more information, please visit <https://datasociety.net/about/#funding>.

PART ONE IN THEIR OWN WORDS:

TROLLING, MEME CULTURE, AND JOURNALISTS' REFLECTIONS ON THE 2016 US PRESIDENTIAL ELECTION

Analyzing the relationship between journalism and the amplification of harmful, polluted, or false information before, during, and after the election requires understanding the influence of earlier online subcultures on the journalists involved. In particular, the subculture that organized around 4chan during the previous decade had a direct impact on the ways many journalists conceptualized and reported on the emerging alt-right narrative. This impact hinged on the changing use of the term “troll” across that decade, as well as an under-examined division among journalists who were troll-trained (through previous exposure to and familiarity with the subculture) and those who were not troll-trained, or who simply rejected the category of trolling.

THE AMBIGUITY OF “TROLLING”

Currently, the term “trolling” is used to describe an enormous range of behaviors online. Depending on who is speaking, “trolling” can subsume acts as simple as disagreeing with someone on social media, as contradictory as feminist activism and violent attacks against feminists, and just about everything in between, rendering the term so slippery it has become almost meaningless.¹ Despite the nebulousness of the “troll” framing, many within the news media, on social media, and even in some academic circles have credited some combination of “trolls” with shifting the norms of acceptable public discourse and contributing directly to Trump’s electoral victory in 2016. This narrative is problematic because it lumps too many individuals and actions into an imprecise category, in turn providing violent bigots, antagonists, and manipulators a built-in defense of plausible deniability, summarized by the justification “I was just trolling.” The narrative is also problematic because it obscures the history of subcultural trolling and its direct influence on many of the journalists who covered 2016 US politics.

THE TROLL-TRAINED VERSUS NOT TROLL-TRAINED

DISTINCTION

Younger respondents (in the 28–32-year-old range) frequently prefaced discussions of alt-right memetic warfare with unprompted discussions of their own teenaged experiences with subcultural trolling and/or 4chan. The precise orientation to trolling culture of these troll-trained reporters varied between those who had been subcultural trolls, those who were troll adjacent, and those who were steeped in the norms of meme culture.² For those whose careers required them, daily, to plunge the internet depths, the widespread sharing of antagonistic memes, racist jokes, and general “shitposting” during 2016 was entirely par for the internet course. These journalists’ news reports often focused on how “funny and bizarre” this content was, often with the assumption that much of it was ironic.

Running just a few steps behind these (typically) younger troll-trained reporters were more traditional, (typically) older reporters – as well as reporters whose bodies numbered among those being targeted by far-right antagonists’ violent bigotries – inclined to approach trollish materials with much more credulity. From their vantage point, there was nothing funny, bizarre, or ironic about any of it, so they tended to amplify what the “trolls” said as fact, with more forceful urgency. Some within this group actively rejected the “troll” frame, taking, instead, a hard line against any element of irony in their reporting. However, as the far-right antagonists in question actively employed “classic” trolling strategies, those who were not-troll trained and those who rejected trolling strategies were particularly vulnerable to their subsequent onslaught of targeted manipulations. The result of the interplay between troll-trained and troll-untrained (or troll-rejecting) reporters was to further muddle the term “troll,” to filter violent bigotries into mainstream discourse, and to catalyze the visibility of alt-right manipulators. As the stories themselves and social media reactions to these stories ricocheted across and between online collectives, what was meant as trolling was reported (and reacted to) seriously, and what was meant seriously was reported (and reacted to) as trolling—all while those on the far-right fringes laughed and clapped.

The nonstop coverage devoted to alt-right antagonists operating under the banner of trolling illustrates the fundamental ambivalence of amplification. However critically it might have been framed, however necessary it may have been to expose, coverage of these extremists and manipulators gifted bad actors a level of visibility and legitimacy that even they could scarcely believe, as nationalist and supremacist ideology metastasized from culturally peripheral to culturally principal in just a few short months.

PART TWO

"AT A CERTAIN POINT YOU HAVE TO REALIZE THAT YOU'RE PROMOTING THEM":

THE AMBIVALENCE OF JOURNALISTIC AMPLIFICATION

Journalists, particularly those assigned to politics and technology beats, were presented with a unique challenge before, during, and after the 2016 US presidential election. The bigoted, dehumanizing, and manipulative messages emanating from extremist corners of the internet were impossible, and maybe even unethical, to ignore. At the same time, news coverage of those messages helped make the messages, and their messengers, much more visible than they would have been otherwise, even when the reporting took an explicitly critical stance. Similar tensions marked stories that didn't have, or at least didn't seem to have, an explicit political agenda, including online harassment campaigns, social media hoaxes following mass shootings and other tragedies, and the plethora of misleading narratives circulating social media.

Speaking to this tension pre-and post-election (and particularly looking forward to the 2018 midterms), all of the reporters interviewed for this project acknowledged, and most expressed deep concern, about the impact of publicizing polluted or potentially damaging information. Just as many reporters expressed deep concern about the impact of *not* publicizing such information. As a result, responses to the question "to amplify or not to amplify," or alternatively put, "to report or not to report" often contained a baked-in ambivalence that simultaneously reflected long-standing tensions within journalism *and* seismic shifts within the information ecosystem ushered in by social media; as soon as the reporter finished listing the dangers of amplification, they would then explain the dangers of doing nothing.

TO AMPLIFY OR NOT TO AMPLIFY

Amplification of harmful, polluted, or false information:

- Increases the likelihood, and raises the stakes, of harassment
- Increases the likelihood that similar disinformation and harassment tactics will be used in the future
- Makes particular stories, communities, and bad actors bigger – more visible, more influential – than they would have been otherwise

STRUCTURAL LIMITATIONS

- Makes it very difficult, if not impossible, not to benefit those looking to manipulate journalists
- Risks normalizing and desensitizing people to harmful views
- Risks lending credence to false narratives
- Relinquishes control of the narrative to bad actors and bad information
- Privileges certain kinds of personalities and information
- Flattens more complicated and contested conversations

On the other hand, not covering stories with false, dehumanizing, or manipulative elements can be just as problematic. Not amplifying harmful, polluted, or false information:

- Allows for the possibility that worse information will take its place
- Means that another reporter will get there first, and cover the story poorly because of a lack of understanding of the story's subject and/or the overall information landscape
- Risks missing an opportunity to educate the public
- Risks reducing specific instances of abuse, harm, or manipulation to abstract concepts, rather than individual lived experiences
- Allows poisonous ideology to flourish, and cedes cultural territory to bigots and manipulators
- Can inadvertently contribute to the process of radicalization
- Doesn't mean that the issue, whatever it is, will go away

STRUCTURAL LIMITATIONS

Journalists' decision to amplify a story or not isn't solely a matter of personal ethics or experience. Rather, these choices reflect a variety of political-economic, techno-cultural, and ideological forces that reporters and editors say direct, or at least strongly influence, the kinds of editorial choices they are able to make. This report identifies four categories of broader cultural forces, as well as a number of long-standing tensions within the profession of journalism, that simultaneously catalyze the spread of bad information and stymie opportunities to contextualize, decode, and respond ethically to networked mis- and disinformation campaigns online.

The Tyranny of Analytics

In the social media age, the measurability and commoditization of content, in the form of traffic, clicks, and likes, has tethered editorial strategy to analytics like never before. The emphasis on quantifiable metrics stacks the news cycle with stories most likely to generate the highest level of engagement possible, across as many platforms as possible. Things traveling too far, too fast, with too much emotional urgency, is exactly the point, but these are also the conditions that can create harm.

STRUCTURAL LIMITATIONS

The Information Imperative

Journalism is guided by the basic tenet to publish, and therefore to spread, newsworthy information. Stories deemed relevant to the public interest are thus marked by what can be described as an information imperative: the norms of journalism dictate that these stories must be amplified. While the information imperative serves a critical democratic function, it can also be harnessed as a tool of manipulation, a point exacerbated by the ubiquity of social media. According to respondents, two primary factors complicating the information imperative, particularly in digital environments, are the prevalence of “iterative reporting” and the frequent inclusion of false equivalencies in news reports, particularly in the US. Iterative reporting is the expectation that journalists should report on what other journalists are already covering. The inclusion of false equivalencies in news reports represents the journalistic norm of reporting on both sides of a story (described by several reporters as “both sides-ism”) on steroids, as positions that are false, manipulative, dehumanizing, and in many cases not worth reporting at all, are given an equal platform to positions that are factually true, relevant to the public interest, and unquestionably newsworthy.

Labor Issues

Reporters – especially interns, reporters early in their careers, and freelancers – are often required to meet excessive word, story, and/or traffic quotas, too easily resulting in rushed and imprecise reporting. Excessive writing demands, and the liabilities they can create for reporters, are not new to the institution of journalism; reporters and editors alike have long needed to navigate this perennial journalism challenge. However, the ever-increasing demands of producing content online further exacerbates the problem. In addition, the push to make reporters as visible as possible, which connects to the push to make their reporting as lucrative as possible, serves as a greased wheel for harassment. Many reporters have inadequate protections from the onslaught of harassment they face on social media, a particular problem for female journalists, journalists of color, and queer journalists. This harassment often prompts news coverage of the harassment, in turn spreading the abusers’ messages, normalizing abuse as part of the job, and incentivizing future attacks.

The Homogeneity of (Imagined) Audiences and Hegemony of Newsrooms

Concerns about “bad information” entering the media ecosystem hinge as much on *who* is doing the reporting as on *what* is being reported. Who is reacting to that reporting is also significant. Many journalists observe that mainstream news is still produced for an audience presumed to be majority white, who are interested primarily in dramatic, emotionally reactive content. Many journalists also observe that the overwhelming whiteness of mainstream newsrooms has a direct impact not just on how these stories are told, but what stories are deemed worthy of telling in the first place.

PART THREE

THE FOREST AND THE TREES:

PROPOSED EDITORIAL STRATEGIES

Part Three departs from Part One and Two's focus on the forest of journalism and instead homes in on the trees: specific editorial best practices designed to minimize narrative hijacking by bad-faith actors, and to maximize reporters' ability to communicate critical truths. These sets of recommendations identify best practices for assessing newsworthiness; reporting on objectively false information; covering targeted manipulation campaigns and specific manipulators; and more broadly, reporting on the internet.

Many of the recommendations made here echo the core tenets of good journalism, which have guided reporters, editors, and publishers in their efforts to grapple with issues of newsworthiness, untruth, and manipulation by state and business actors since the start of the profession. These recommendations – shaped by interviewees themselves – build on these tenets to reflect the challenges specific to social media and networked manipulation campaigns. While Part Three is geared most directly to working journalists, these suggestions about how best to respond to problematic information are also applicable to everyday users of social media, who serve as critical links in the overall amplification chain.

TIPS FOR ESTABLISHING NEWSWORTHINESS

Journalists must always assess the newsworthiness of information, especially when it concerns potential manipulations. There are three broad criteria to assess:

- *Tipping Point* – has the story extended beyond the interests of the community being discussed (Moschella and Watts 2017)? In the case of online memetic content, this question would direct reporters to consider whether a particular meme has been broadly shared by anyone outside the core group of participants.
- *Social Benefit* – will the story have a positive social benefit, open up a new conversation, or add weight or exemplars to an existing conversation?
- *Potential Harms* – will the story produce harm (embarrassment, retraumatization, professional damage), or could an audience use the story to cause harm (attacking sources, imitating crimes)?

TIPS FOR REPORTING ON OBJECTIVELY FALSE INFORMATION

As in all contexts, journalists must take special care when reporting on objectively false information, whether satire, hoaxes, conspiracy theories, or political manipulations.

One research participant offered the following criteria:

- Determine if the story reaches the tipping point
- Determine if there would be a public health takeaway (i.e., something worth learning) from the debunking; for example, explanations that identify and analyze manipulators' rhetorical strategies, including their use of humor
- Determine if there is a political or social action point (i.e., something worth doing) related to the falsehood itself; for example, editorials that provide media literacy strategies for recognizing and resisting networked manipulation campaigns
- Determine if the risk of entrenching/rewarding the falsehood in some stories is worth dislodging the falsehood in others.

If the answer to each of these questions is no, then the story isn't worth reporting at that time. If a story ultimately passes the tipping point and does become appropriate to report, reporters should be transparent about the origins and context of the information. Whenever possible, experts in the particular subject area should be recruited to write or consult, to ensure the clearest and most informed refutations possible.

TIPS FOR REPORTING ON SPECIFIC HARRASSMENT

CAMPAIGNS OR OTHER COORDINATED

MANIPULATION ATTACKS

Journalists must take special care when covering the coordinated campaigns of polluted information. The following journalism best practices were suggested by research participants:

- Treat **violent antagonisms as inherently contagious**, akin to coverage of suicide,³ mass shootings,⁴ and terrorism,⁵ all of which are known to inspire and even provide behavioral blueprints for future copycat attacks. Similarly, **wall-to-wall coverage of online harassment and manipulation incentivizes future attacks** by signaling that such behaviors will result in the desired outcome – non-stop attention for the attackers.
- When stories focus on the targets of online harassment, reporters should be careful **not to minimize their subjects' experiences by suggesting that digitally mediated harm is less serious or less real** than physical harm.

TIPS FOR REPORTING ON OBJECTIVELY FALSE INFORMATION

- Reporters should reflect on how stories profile primary and secondary victims, and what information that reporting publicizes about them. Reporters should **minimize the inclusion of unnecessary identifying information**, talk to the victim – and ideally their friends and family – about these details, and see what information they are comfortable sharing publicly.
- To the extent possible, **stories should specify the number of participants in a particular online attack/campaign**, rather than using vague mass nouns (i.e., trolls did this, the alt-right did that). Important contextualizing information includes the apparent number of online participants (based on observational data), why the reporter believes this count is accurate, and any unknown variables that might impact the readers' understanding of the story. When describing media manipulation campaigns of any kind, stories and their headlines should employ the most precise language possible.
- Given how nebulous the term has become, and how easily it is used to cloak hateful behaviors, **“troll” should be used sparingly in stories and headlines**, if at all.
- Publications should **avoid publishing listicles that round up the worst examples of racist or misogynist expression without significantly commenting** on that expression; while they may seek to call attention to abuse in order to condemn it, these articles inadvertently provide abusers, and their abusive messages, a much larger platform.
- Stories should avoid framings that fall back on “both sides-ism,” in which the attack is described, followed by an overview of what both the attackers and the attacked have to say about it; such a framing elevates the perpetrators to an equal platform as those being harmed. Particularly in cases where the attacks are motivated by racial or gender animus, **reporters should foreground the perspectives of those harmed**, and not give aggressors an opportunity to justify, spin, or further normalize their hateful behavior.
- Reporters and their **editorial teams should exercise an abundance of caution when reprinting memetic images** used during particular attacks, especially when the images are dehumanizing and bigoted. When sharing an image is deemed necessary, editorial teams, along with members of the communications team, should consider including captions from the story and/or other contextualizing information within the image itself so it can't be hijacked and spread by manipulators as easily.
- When approaching breaking stories about identity-based harassment and violence online, reporters should **be especially conscientious about the histories of this violence**, and the histories of the activists who have been working to combat it.
- Reporters should make an effort to **talk to people who have direct, embodied experience with the interpersonal, professional, and/or physical implications** of a given issue, while also being sensitive to the emotional labor of the ask.

TIPS FOR REPORTING ON SPECIFIC HARASSMENT CAMPAIGNS OR OTHER COORDINATED MANIPULATION ATTACKS

- If a story includes, or seems like it may include, polluted information, reporters shouldn't just quote from, but should **actively consult with experts who have studied computational and/or networked propaganda**, or other forms of media manipulation.
- Reporters should consider their personal relationship to the story, and how their own experiences factor into the conversation. They should **take preemptive protective and/or self-care measures – both at the level of individual mental health support and/or the kinds of safety trainings** and resource-sharing offered by organizations like the International News Safety Institute and the International Women's Media Foundation – if a story is likely to trigger post-traumatic or secondary stress responses.

TIPS FOR REPORTING ON SPECIFIC MANIPULATORS, BIGOTS, AND ABUSERS

In the instance that reporters are asked to cover specific manipulators, bigots, or abusers, several guidelines will help bolster the positive impact of the reporting, as well as protect the reporter:

- While stories must address the manipulators, bigots, and abusers involved in particular attacks, **reporting should avoid framing bad actors as the center of the narrative.**
- When framing a story about specific white nationalists and supremacists, reporters and editors should **run a newsworthiness calculus on each personal detail** they are considering reporting.
- Reporters and editors should be aware of how strategic many groups of white supremacists and nationalists are in their communications and messaging, which is geared toward maximizing recruitment. Similarly, reporters and editors should be aware that **extremist groups, along with other groups of media manipulators, are eager to use journalistic norms as a weapon against journalism.**
- Building on longstanding best practices in journalism, reporters and editors should **respond with heightened vigilance** when antagonists, bigots, or other stripes of manipulator reach out with a tip or unsolicited commentary. Ask whether or not the apparent agenda can be verified.
- In cases when a reporter is inclined to reach out directly to a specific antagonist, manipulator, or abuser, they should **first reflect on whether the story absolutely requires quotes from these individuals.**

TIPS FOR REPORTING ON SPECIFIC MANIPULATORS, BIGOTS, AND ABUSERS

- If the story does warrant an interview (because it helps establish context, because it more clearly illustrates what exactly the individual is advocating, because it serves a counter-argumentative function), reporters should situate bigoted or manipulative sources' statements historically and ideologically, and **minimize the inclusion of euphemistic dog whistles**.
- Reporters should be aware that all communications in emails and in interviews, in fact **anything reporters say publicly or even semi-privately about a particular story and/or subject, may be used** against the reporter and/or their publication.
- Whether subjects are directly interviewed or are observed on social media, **reporters should weave the performative nature of manipulators' actions** into the story.
- Stories should **minimize focus on individual motivations** or personal psychology.
- Stories should **avoid deferring to manipulators' chosen language, explanations, or justifications**; for example, when violent white supremacists claim they are just trolling in order to deflect personal responsibility for spreading hate. They may say it's "just trolling," but stories should describe the behaviors, and their impact on targeted communities, as accurately and as free of euphemism as possible. Just as importantly, stories should not employ the aggressors' insider lingo to describe specific actions or targets.

TIPS FOR REPORTING ON THE INTERNET

The following are a series of suggestions applicable to all stories with online elements. By standardizing these more general strategies, reporters not specifically assigned to digital culture or technology beats, as well as the editors overseeing these stories, will be better equipped to navigate the increasing overlap between "internet" and "non-internet" subjects:

- Rather than pointing to the fact that something on the internet exists, stories should **focus on how a particular interaction, technology, or community works and why** that matters to a broader audience.
- Editors should **avoid assigning breaking stories about online communities or behaviors to people unfamiliar with those communities** and behaviors.
- Stories should be framed especially carefully when they address behaviors on and around 4chan following a tragedy. Reporters assigned to these stories should conduct background research on the history of 4chan and other online subcultures, and should **assume that claims emanating from 4chan, particularly in the wake of a tragedy, are pointedly performative**, and almost certainly false.
- Given the tendency for people online – from reporters to massive YouTube audiences

TIPS FOR REPORTING ON THE INTERNET

– to equate “poor quality” with “authenticity,” manipulators have increased efforts to play into this confirmation bias. Reporters should therefore **take caution when anyone claiming to have a tip sends over files with overly amateur and/or analog stylings**; for example, poor photo quality, shaky video, images sloppily annotated using iPhone markup tools, and anything photocopied.

- Reporters and their editors should internalize the idea that **social media does not constitute a “person on the street” scenario**, nor is an embedded tweet or Facebook post akin to a pulled quote. Regardless of the kind of story being reported, reporters should avoid linking to a handful of social media posts and then attributing that perspective, positive or negative, to “the internet.”
- Reporters should **talk to sources for digital culture stories at length, ideally face-to-face**.
- Reporters, editors, and publishers should be reflective about **how social media and other digital tools complicate established reporting practices**.
- Reporters, editors, and publishers alike should prefigure every professional decision with the recognition that **individual journalists are an integral part of the news being reported, and further, that the institution of journalism is, itself, the system being gamed by manipulators**.

ENDNOTES

1 For more on the many problems associated with the contemporary troll frame, particularly when used to describe bigoted, identity-based antagonisms, see Phillips (2016) and Phillips and Milner (2017).

2 Sometimes used interchangeably with “internet culture,” meme culture refers to the aesthetic, behavioral, and rhetorical traditions associated with the sharing and remixing of memetic media across social platforms; see Milner (2016) for more on the logics of memetic spread. Although it is frequently used as shorthand in academic and journalistic circles, “meme culture” (along with “internet culture”) belies the fact that multiple, sometimes overlapping, sometimes outright conflicting, cultures employ similar expressive traditions and practices.

3 See Phillips (1974); Gunn and Lester (2012); Bohanna and Wang (2012); Tufekci (2015).

4 See McBride (2017); “Recommendations for Reporting on Mass Shootings” (2017).

5 See Beckett (2016); “Terrorism and the Media: A Handbook for Journalists” (2017); “Ankara Memorandum on Good Practices for a Multi-Sectoral Approach to Countering Violent Extremism” (2016).