

Fighting Fake News

Workshop Report

hosted by

[The Information Society Project](#)
[The Floyd Abrams Institute for Freedom of Expression](#)

On March 7, 2017, the Information Society Project at Yale Law School and the Floyd Abrams Institute for Freedom of Expression hosted a workshop intended to explore the ongoing efforts to define fake news and discuss the viability and desirability of possible solutions.

The discussion encompassed attempts to identify the particular harm associated with fake news; the many First Amendment questions that arise in any attempt to create governmental regulations on specific kinds of speech; and the pros and cons of self-regulation by those involved in the digital ecosystem.

This workshop was meant to be a first step towards encouraging interdisciplinary conversation and work on these issues. There were twenty-one participants from various disciplines, including members of academia, the practicing bar, news organizations, information intermediaries, data scientists, computer scientists, and sociologists.

The workshop was held under Chatham House Rules. This report highlights some of the many points raised during the day-long discussion. It does not represent the views of the individual participants, their affiliated institutions, nor the sponsoring organizations. Nor is this report a transcript; many points raised by participants have been rearranged by subject matter for readability.¹

¹ Sandra Baron and Rebecca Crootof prepared the initial draft of this report, based in part on notes provided by Anna Gonzalez, YLS '18. Participants were given the opportunity to review it and make corrections and suggestions before publication, but not all participants did so. With awareness of the irony, citations to relevant studies have been excluded to avoid inappropriate associations between statements and participants.

Table of Contents

Session 1: Defining the Problem of “Fake News”	3
Session 2: How Misinformation Spreads.....	5
Session 3: Identifying Players and Pressure Points	6
Session 4: Proposals for and Problems with Government Solutions.....	7
Session 5: Proposals for and Problems with Non-Governmental Solutions.....	10
Routes to Solutions	11
Participants.....	15

Session 1: Defining the Problem of “Fake News”

In retrospect, the issue that proved the most challenging for the workshop participants—defining “fake news”—was never satisfactorily resolved. Instead, participants relied heavily on First Draft’s taxonomy, which identifies seven types based on degrees of falsity and intentionality, which was recognized as being helpful but incomplete.²

The discussion began with the question of whether “fake news” has now been used in so many different contexts that it is now fundamentally worthless. Workshop participants did distinguish between information and articles that are false from inception from information and articles that may or may not be false but are framed in ways to make them highly charged and often misleading. The latter category often is characterized as propaganda.³ Rather than spending an inordinate amount of time attempting to create a more precise definition of “fake news” or to debate the variety of forms it takes, participants focused instead on identifying its specific harms.

Participants determined that the most salient danger associated with “fake news” is the fact that it devalues and delegitimizes voices of expertise, authoritative institutions, and the concept of objective data—all of which undermines society’s ability to engage in rational discourse based upon shared facts. From this perspective, the distinguishing between Macedonian teenagers who distribute false stories for profit and those who engage in ideological propaganda may be focusing on a distinction without a difference, given how both contribute to creating societal chaos. The intent of the creator is less relevant than the fact of the harm—the insidious damage is the fact that the proliferation of false information discredits sources of relatively accurate and credible information regardless of what a specific “fake news” story is intended to accomplish.

Three corollary harms were noted: first, the problem of increasing fragmentation and politicization; second, the promotion of “safe news” at the expense of difficult or challenging news stories; third, the need for credible sources to allocate ever-diminishing resources to debunking inaccurate information (which poses both financial and reputational costs).

One participant observed that, if the primary harm of “fake news” is that it undermines trust, the common solution of “more news” doesn’t address this underlying problem at all.

Once these harms were raised, participants identified a number of structural reasons why these problems are particularly prevalent now:

- The exchange of information is now democratized, thanks to social media platforms and digital content production technologies (like Photoshop). Anyone is now able to produce credible “noise” that is difficult to distinguish from high-quality information.

² Claire Wardle, *Fake News. It’s Complicated*, FIRST DRAFT NEWS, Feb. 16, 2017, <https://medium.com/1st-draft/fake-news-its-complicated-d0f773766c79>.

³ While the term “propaganda” is often used as a synonym for “fake news,” it should be distinguished: propaganda need not be false; rather, it achieves its intended effect by emphasizing in-groups and outcasts or by creating dystopic realities. The colloquial use of “propaganda” emphasizes that it is often perceived as being used by a powerful few to rally or shape the understanding of a weaker many.

- The demand for “fake news” may be a natural byproduct of faster news cycles and increasing consumer demand for shorter-form content.
- While there is a general awareness of the existence of “fake news,” there is widespread disagreement over what comprises “fake news.” Merely labeling something as “fake news” can itself be considered mere propaganda, making it all the more important that journalists cite sources and “show their work.”
 - Press-branding campaigns that attempt to distinguish between traditional journalism or respectable new sources of media and propaganda or outright lies have not been an effective means of reestablishing the authority of the press. This is primarily due to social reasons to prioritize peer-determined “truth” over previously authoritative voices, the psychological realities of tribalism, the power of confirmation bias, and the dopamine surges associated with outrage.
- Traditional gatekeepers are less effective or visible. For example, traditional news organizations lack the institutional authority they once enjoyed. (This is also true for many other historically influential and authoritative voices, including medical professionals, scientists, religious leaders, and academic institutions.)
 - That being said, “fake news” often presents as traditional journalism, borrowing the authority of traditional journalism while simultaneously undermining it.
- Current gatekeepers are more likely to view news production and dissemination as a business enterprise than as providing a public service. Additionally, the public perception of mass media as a corporate, profit-driven entity has further diminished its authority.
 - This profit-driven approach may be partially due to the fact that most content distributors are no longer generally owned by a small group of families possessing a kind of *noblesse oblige*. While diversification is to be welcomed, a side effect of how this diversification has played out is that profit has been emphasized to the detriment of other aims.
- New respected and trusted gatekeepers have yet to be established.
- Ownership of news distribution has shifted from traditional content creators to digital distributors. Digital distribution allows for highly efficient micro-targeting and limited exposure of users to challenging content. In contrast, when content creators also were responsible for distribution, diverse content was often bundled together for a mass audience, fostering the development (either voluntarily or serendipitously) of a common set of shared facts. (One participant referred to this as having to eat your broccoli with your twinkies.) Digital distribution also tends to favor popularity, engagement, and “shares” over expertise and accuracy.

It is worth noting that, over the course of the workshop, some participants questioned our focus on fake news, expressing the opinion that the real problems lie elsewhere.

- One participant observed that, rather than being its own problem, fake news is actually merely a symptom of much deeper structural problems in our media environment. This participant questioned whether we should focus on those problems first, but simultaneously noted that doing so might not be tractable.

- One participant suggested that fake news poses a relatively trivial problem for various reasons: (1) it is competitive; (2) it is visible to users; (3) it is subject to confirmation bias; and (4) its impact is determined entirely by how digital distribution platforms—such as Facebook, Google, and Twitter—rank stories, resulting in the power to rank being far greater than the power of inaccurate content. This participant suggested that new means of online manipulation that are not competitive or visible but nonetheless cause shifts in an individual’s opinions, purchases, or voting preferences are of far greater concern. For example, biased search results can shift voter preferences dramatically, without anyone’s knowledge or awareness of why their opinion is changing.

Session 2: How Misinformation Spreads

In this session, participants considered how misinformation spreads and the role of online social media in creating and exacerbating echo chambers and filter bubbles. The discussion leader began with two premises: (1) An individual’s opinions and beliefs are influenced by what he or she reads; and (2) Most people choose to interact with those who share similar opinions (and avoid or “unfollow” those with whom they profoundly disagree). As a result, content consumers end up occupying segregated and polarized groups.

Because human beings are more likely to believe there is a reason for something if we see others promoting it, retweeting or sharing information alters how that content is perceived by subsequent content consumers. If we see a crowd of people running, our natural inclination is to run as well. Historically, this response may have helped us avoid predators; in today’s digital world, it makes us vulnerable.

People often use the number of retweets or shares as a proxy for credibility, even though there are many reasons to be skeptical of those numbers. First, the literature on signaling (especially Dan Kahan’s work), highlights how people repeat phrases—or retweet or share—to signal their membership in a certain group, and regardless of whether they personally believe or endorse the content.

Furthermore, bots are often used to falsely promote a piece. The practice of “astroturfing”—creating a false grassroots movement—builds on this by strategically distributing a specific piece of news through a variety of sources (such as front groups, sockpuppets, and bots) to give the impression that numerous sources are discussing the article. These practices help spread misinformation, manipulate what items appear to be trending, and ensures that “fake news” looks more popular than its more credible counterpart.

This bias also helps explain why “fake news” persists despite fact-checking. Not only may fact-checking articles not reach the same people who view the original piece, but the reiteration of the original claims by fact-checkers may lend them credence. Meanwhile, when contrasted with widely-shared misinformation, the fact-checking response might appear to be a minority and therefore less credible opinion.

One of the ways platforms contribute to this problem is by creating an environment of high levels of information unmatched by people’s limited attention. People retweet or share an article based on its headline and without ever having clicked on—and therefore without ever having actually read—it. This allows misinformation to be seen, accepted, and promoted just as much, if not more, than higher-quality information.

In short, any attempt to “grade” the information quality of a given work or to flood the “marketplace of ideas” with more information would not be effective solutions, as it is difficult for high-quality information to crowd out low-quality information.

This session concluded with two questions:

- Assuming one can identify an objective truth, how do you give people the tools to get to that truth?
- Even if you could get people tools to distinguish truth from fiction, would people care enough to use those tools?

Session 3: Identifying Players and Pressure Points

Over the course of the discussion, the primary players and pressure points were identified as:

- Content consumers
- Content creators (journalists, bloggers)
 - Some would include newspapers and broadcasters with content creators, on the grounds that they exercise some control over the created content and are not covered by a safe harbor.
- Content distributors
 - There was some disagreement as to how best distinguish between different kinds of content distributors.
 - Some favored distinguishing between traditional content distributors (newspapers, broadcasters) and digital content distributors (wikis, blogs, social media platforms, search engines, online news aggregators).
 - Others favored divisions based on whether a content distributor has an editorial process (newspapers, some blogs) or relies on algorithmic selection in determining what content is foregrounded (search engines, some social media platforms).
- Norm guardians (institutional fact checkers, trade organizations, and “name-and-shaming” watchdogs)
- Economic supporters (advertisers, foundations, and users)

Consumers play a large role in what content is created and how it is disseminated. As reported by the discussion leader, the United States has now reached almost full digital penetration: approximately 88% of American adults are online, and 85% are getting news from online sources. Almost 75% of American adults are now accessing news on their phones (as compared to 50% a few years ago). Social media is a significant provider of information: it is just as common for an American adult to get news through social media as through a news organization's website or app.

Participants noted that people tend to trust our networks of friends and family for news, and these organic formations are reflected and exacerbated by social media platforms. When people receive news and information through social media, they are less likely to be aware of the source of the information. They are more likely to remember a news source if they receive a link through email, a text, or a news alert.

Additionally, when content consumers go online, traditional news distributors lose revenue. As one participant put it, the economics of online journalism is brutal. For every \$1 gained in online clicks, \$15 in print revenue is lost. Historically, newspapers hired professionals who investigated, wrote, fact-checked, double-checked, and then had their work reviewed by an experienced editor. There's no value in this process in online journalism, because it takes too much time—and if you publish late, you might as well not publish at all. One participant noted that we've created a business model that destroys what we purport to desire.

One participant noted that it might be most helpful to approach this problem by thinking about how data cycles through our communications systems, how different kinds of data are promoted or abandoned by those systems, and how these different systems are more or less harmful to democracy and democratic discourse.

Session 4: Proposals for and Problems with Government Solutions

The discussion leader noted that, when discussing governmental regulation, it is important to distinguish between “the negative state” and “the positive state.” The negative state involves the government engaging in coercive actions, such as fining, taxing, and imprisoning. The positive state involves creating institutions and incentives, like land grant colleges or tax subsidies. The government has far more leeway when it takes positive action than when it takes negative action. Historically, governments took a negative state approach to speech regulation and regulated speakers; modern governments tend to take a positive state approach and regulate the infrastructure that enables the flow of information.

The discussion leader argued that “fake news” would generally fall into the category of public discourse and receive substantial First Amendment protection, regardless of its accuracy.⁴ Of course, not all speech is “public discourse.” Professional speech, commercial speech, and court

⁴ The Supreme Court has repeatedly held that false speech enjoys full First Amendment protection. *See, e.g.*, *United States v. Alvarez*, 567 U.S. ____ (2012).

testimony are not considered public discourse and so are more subject to regulation. “Fake news,” however, would likely fit into the “public discourse” category.

Some participants disagreed and argued in favor of testing the depth of First Amendment protections for “fake news.” It was also suggested that we consider ways in which “fake news” might be subject to rules that apply in non-public discourse frameworks.

In general, however, most participants were reluctant to propose negative state regulations for “fake news.” Some argued that the difficulty of defining “fake news” raised the attendant risk of overbroad government regulation. Others worried that opening the door to permitting government punishment of certain kinds of public discourse would grant the government too much power to control speech in areas of public concern. There were similar reservations to state-level regulations.

The option of using government funding or other economic incentives to indirectly promote legitimate news and information outlets was floated, but this was critiqued on similar grounds as those associated with government intervention to penalize certain kinds of speech—we simply do not want government actors determining what speech is true or worthy.

With a nod to cable regulation as a structural model, it was suggested that social media and search engines could be required to put alternative views on consumer feeds or in responses to queries. However, apart from the obvious challenge of determining what constituted an oppositional voice in a multitude of issues and ideas, there is also the likely chance that users would simply shift to a different service that was able to evade or disregard such regulations.

Some favored developing “whitelists” of articles or news sources, based either on user or an independent institution’s ratings. This proposal was critiqued on the grounds that government-regulated “whitelisted” media often becomes a proxy for state-sponsored or government-approved news.

Given the issues inherent in governmental regulation of content, participants then considered governmental regulation of technological architecture.⁵ Proposals included labeling bots, requiring that shared content reflect subsequent corrections or revisions, and permitting third party enforcement of platform terms of service regarding false speech. These and other suggestions are presented in greater detail in the “Routes to Solutions” section below.

Participants acknowledged that distributor liability (which has been analogized to intermediary liability) is not absolute. Digital content intermediaries have generally been afforded greater protection than other distributors as a result of §230 of the Communications Decency Act,⁶ but while that protection might be considered good policy, it is not constitutionally required.

⁵ The German proposal to regulate “fake news” by regulating information intermediaries was acknowledged, but the details had not been made public at the time of the workshop and were not discussed. *See* Anthony Faiola & Stephanie Kirchner, *How Do You Stop Fake News? In Germany, with a Law*, Apr. 5, 2017, WASH. POST, https://www.washingtonpost.com/world/europe/how-do-you-stop-fake-news-in-germany-with-a-law/2017/04/05/e6834ad6-1a08-11e7-bcc2-7d1a0973e7b2_story.html.

⁶ 47 U.S.C. § 230 (1996). Section 230 provides: “No provider or user of an interactive computer service shall be treated as the publisher or speaker of any information provided by another information content provider.”

Some suggested lessening the liability protections afforded intermediaries, possibly by permitting defamation suits. It is rarely easy or worthwhile for a defamed individual to sue the content creator—even if his or her identity can be ascertained, the defamer might well be impecunious or outside of national court jurisdiction. Creating remedies against domestic intermediaries offers those defamed by “fake news” a potential route to a recourse. The same might be argued about claims for invasion of privacy or other content-based torts. Presumably, increased intermediary liability would encourage intermediaries to better screen the content they permit on their platforms.

One benefit to governmental regulation is that it can address problems that platforms would like to address but feel unable to do so, either because they do not wish to be tarred as censors or because the economic incentives run contrary to their interfering with user access to content unless their competitors do so as well. Another potential benefit to increased intermediary liability would be that it might encourage content intermediaries to view themselves as the media companies of the 21st century, with all of the normative obligations that entails. (One participant noted that some platforms already take this view.)

However, many participants strongly disagreed with the utility of increasing intermediary liability as a matter of policy, even if it is constitutionally permissible. Some argued that a stronger notice and take-down enforcement regime—the likely outcome of greater liability—would result in too much information being censored or that the process would likely lack transparency. Others said that it was already difficult for information intermediaries to comply with existing law, even with §230 protections—there is simply too much information being uploaded to monitor it all. There was also skepticism that government regulation of information intermediaries would ever be sufficiently tailored, given the wide variety in types and structures of different platforms and their different user populations.

While nearly all proposals for government interventions were hotly debated, most participants favored the idea that the government could indirectly minimize the impact of “fake news” through promoting critical thinking training in public schools or through government-sponsored computer and content literacy tools and training programs.

It was acknowledged that social media platforms and other related tech companies are in an odd position: they are arguing that their contribution to the spread of fake news is purely happenstance, an unfortunate combination of human traits and neutral information-sharing platforms. However, they are simultaneously arguing that their algorithms are due First Amendment rights as a type of speech. This presents a challenge on how best to view the results of algorithmic sorting: is it the voice of the platform, or not speech at all?

Finally, participants acknowledged that one of the potential benefits of pursuing governmental regulation was that its threat might spur tech companies to develop their own internal means of addressing the problem—which provided an excellent transition to the subsequent session.

Session 5: Proposals for and Problems with Non-Governmental Solutions

This session began with an analysis of why search engines (as currently constructed) are particularly ill-suited to addressing the “fake news” problem. The discussion leader argued that the driving force behind fake news is the desire to attack the opposing side’s truth indirectly; not by confronting it, but rather by muddying the waters with too much information. Search engines don’t address this problem—they view themselves as virtual card catalogues for the universe of information, programmed to seek out the most relevant and highest quality information in response to a user’s query. Of course, how a given search engine determines what constitutes relevant, high-quality information varies and is often not transparent, but it is generally known that the popularity and the number of citations are often relevant factors in determining a search result’s likelihood of promotion. Unfortunately, these criteria are also consistent with the distribution characteristics of “fake news.” The practice of using the number of clicks as a proxy for relevance and the number of links as a proxy for quality should be reconsidered in light of the need to also employ signals about accuracy and truthfulness.

Many participants favored reevaluating platform algorithms as a way of addressing the problems associated with fake news. That being said, there is an acknowledged risk associated with just stating that we should “improve the algorithms.” Introducing any corrections into the sorting system invites (1) the introduction of biases and (2) increased gaming (i.e. search engine optimization). One participant suggested that this circle might be squared by strengthening algorithms to prevent gaming, which might reduce bias more than trying to write anti-bias code. Additionally, while algorithms may be improved to identify and demote blatantly false material based on certain shared characteristics, it is going to be far more difficult to address material grounded on a kernel of fact that is buried under speculation, hyperbole, defamation, and spin.

The concept of “whitelisted” articles, content creators, or content distributors arose again. In contrast to the government-approved whitelists discussed earlier, a participant noted that certain brands are already creating their own whitelists of acceptable or generally reliable content, and independent companies are offering services to assist advertisers with neutral ad placement. While this might minimize patently false clickbait, it will not address many other kinds of “fake news.” Politically motivated speech, for example, is unlikely to respond to the economic pressure associated with the withdrawal of ad revenue. Furthermore, depending on who creates them, whitelists may simply exacerbate echo chambers.

One participant observed that the missions of journalism and platforms are fundamentally at odds: Journalism attempts to provide needed but sometimes uncomfortable facts and the context necessary to make sense of them to keep their users informed. Platforms aim to provide desired and entertaining information to keep their users happy and interested. It is not clear there is a way to link these separate missions such that the new digital content distributors can assume the role of 21st century media providers.

Routes to Solutions

As the workshop drew to a close, nearly all participants agreed on one overarching conclusion: that **reestablishing trust in the basic institutions of a democratic society is critical to combat the systematic efforts being made to devalue truth**. In addition to thinking about how to fight different kinds of “fake news,” we need to think broadly about how to bolster respect for facts.

As one participant noted: Democracy is a strange idea, in that it puts complex problems to a plebiscite. To work best, it requires an informed citizenry with critical thinking and reading skills. The United States – among many other countries – is confronting a genuine breakdown of the basis upon which democracy rests. There is a real movement to delegitimize institutions, undermine rational discourse, and promote tribalism in the service of increased chaos and power consolidation. Meanwhile, human beings are susceptible to fake news that confirms desires and beliefs—and while that has always been true, new technologies cater to these traits and create echo chambers that ultimately destroy the informed discourse upon which a functioning democracy depends.

Unfortunately, there is no quick, permanent, or easy fix to the problems associated with the various kinds of “fake news.” That being acknowledged, participants proposed a number of ways in which different problems might be addressed over the course of the day. These have been roughly grouped into two categories: suggestions that focus on different players (content consumers, creators, and distributors) and suggestions that focus on direct and indirect means of regulation (through social norms, technological design, markets, and law).

Suggestions by Player

- **Content consumers** must be better educated, so that they are better able to distinguish credible sources and stories from their counterparts.
 - Creating more critical news consumers could be approached at a variety of levels. Consumers could be educated about how news information propagates in today’s world, the harms of fake news, and how to identify it. Another approach would be to enable the consumer to learn more about the actors and stakeholders in a given story, better understand their affiliations, and to identify and question unverified details.
 - Critical news consumption could be incorporated into school curriculums or promoted through government-sponsored computer and content literacy tools and training programs.
 - While this long-term solution was widely endorsed, one participant noted that education is hardly a panacea. Well-educated individuals often default to ideologies, and the lack of news literacy can hardly be the reason for the relatively recent fake news problem, since it is unlikely that news literacy has gotten significantly worse in recent years. Instead of focusing on education, this participant argued, we need to double down on supporting facts.

- **Content creators and content distributors** need to adapt their professional norms to new technologies of distribution, better explain these norms to the public, and protect them from erosion. They should also engage in practices that increase the likelihood of producing trustworthy, high-quality information.
 - Any given article can be erroneous or sloppy, but mainstream media gains its authority from voluntary compliance with certain norms, including using headlines that accurately reflect the content of a piece, identifying double- or multiple-sourced factual assertions, highlighting reliance on unnamed sources, including photos of reporters and links to their bios, and acknowledging and publicizing inaccuracies and correcting them. Editors often play a critical role in enforcing these norms.
 - Fact-checking by itself accomplishes little: when a statement is fact-checked and found to be false, a person holding an underlying opinion shaped by the false information rarely changes that opinion. In some cases, fact-checking may have the counterintuitive effect of strengthening a false claim, as it is repeated and human beings' belief in a claim tends to increase the more they hear it. Instead of another fact-checked article, content creators could focus on providing counter-narratives, stories that debunk the claim without explicitly restating it.
 - To the extent statements are labeled as false, it is preferable for content distributors to present fact checks as the product of the organization (like an unattributed editorial), rather than relying on individual journalists, speaking in their personal capacities, to do so. When a fact checking piece is presented as the work of an individual journalist, it dilutes the power of the fact check.
 - It is especially important to distinguish between objectivity (“given everything we know, this is wrong”) and neutrality (not taking a side).
- **Content distributors**—particularly digital content distributors—should engage in practices that minimize the spread of fake news and promote the dissemination of trustworthy, high-quality information.
 - This will sometimes require enforcing the norms of professional newsgathering discussed above. Content distributors who deal with non-edited content might consider whether and how they should employ individuals to serve an editorial function.
 - Content distributors should consider different tactics for countering click bait, to avoid being the conduits by which inaccurate information is disseminated.
 - More individual journalists seem willing to label statements as false in their own voice, but this may have the unfortunate side effect of causing consumers to dismiss such writers and their associated forums as propaganda promoters. It would be helpful for content distributors to take on more of a watchdog role or have a means for flagging inaccurate content.

Direct and Indirect Regulatory Suggestions

- **Accreditation systems:** There might be some form of accreditation system for content creators and/or content distributors.
 - One benefit of an accreditation system might be that it would cement and publicize existing norms and best practices, at either the content creator or content distributor level.
 - Professional organizations and coalitions might take on the task of creating codes of conduct or accrediting systems. Alternatively, or additionally, independent accrediting institutions could be created, provided they do not undermine First Amendment protections for speech.
 - One participant cautioned that formalized codes of conduct could be used against media organizations and speakers in libel, privacy, and other suits.
 - Any accreditation system must be designed in a way to avoid it becoming a means of silencing the “little guys” or a means of promoting government-approved news.
- **Technological design:** Design might be harnessed to both promote more credible content and/or to present content in a way that fosters consumer skepticism and critical analysis.
 - Digital content distributors should be aware that everything has a tradeoff: making some content more easily available also segregates users or reduces the likelihood that users will stumble across contradictory information. All design decisions need to be approached with this awareness.
 - For example, one participant critiqued Google’s “box” on the grounds that when Google answers a query in a highlighted box, virtually no one explores the search results.
 - This might include algorithms that identify bogus content—for example, the articles being written by Macedonian teenagers for economic gain—or which flag questionable pieces.
 - However, there was also some concern that this would just create a form of shadow censorship, whereby credible but unusual news providers would never be able to distribute their findings.
 - Additionally, it was pointed out that there only needs to be one false positive for the flagging entity to be accused of censorship.
 - Some suggested restrictions on the use of bots or requiring the identification and explicit labeling of bots and other means of boosting the perceived popularity of a given piece of content. For example, it might be possible to use CAPTCHA to prevent bots from sharing news stories.
 - However, one participant noted that labeling poses a huge practical problem, especially for smaller or newcomer digital operations, because it is both expensive and difficult.

- Platforms could require users to click on a link before sharing or retweeting it, to increase the likelihood that the user actually reviewed the content he or she was promoting.
- Content intermediaries could also be more transparent about the factors that result in some articles being promoted and others languishing.
- Certain practices—like unfollowing or blocking someone—promotes polarization. Platforms could make it harder to avoid confronting information from people with whom we disagree.
- Search engines should consider using signals that better correlate with quality, accuracy, and trustworthiness than number of clicks or number of links.
- **Market incentives** can be modified, either through legal or social mechanisms.
 - At present, funding incentives favor short-form, national news. We need to develop a system that better funds local journalism and deeply-researched, long-form journalism.
 - On the flip side, we should better understand why the current structure incentivizes for-profit false speech, like platform ad placements and programmatic advertising systems, so that we can determine how best to respond to or dismantle them.
 - Inspired by digitally organized consumer boycotts, advertisers have recently shown an increased interest in ad placement, and they are insisting that they or their advertising agencies exercise more control over what content their advertising supports. Some companies have created their own whitelists of sites they deem to have reliable content.
- **Law:** There may be relatively few legal tools capable of addressing aspects of the “fake news” problem, but the few that exist can be used to neutralize some of the worst, most blatantly false and profit-driven pieces.
 - Many of the suggestions noted above could be bolstered through direct legislation or indirect incentives that encourage certain best practices.
 - We need to develop legal tools that target methods of information distribution rather than content.
 - Platform terms of service (such as rules against impersonation and abuse) can be better enforced, either by the platforms taking the initiative to do so themselves or by permitting third-party enforcement of terms of service.
 - One participant argued for focusing on eliminating gerrymandering—some of the incentives for political echo chambers will evaporate if politicians are forced to represent a diverse set of constituents.

There was general agreement that it would be helpful to have more social research on questions of how to break down tribalism barriers and encourage cross-tribal communications. We need to zoom out and better understand why there is a demand for news that confirms existing desires and beliefs—and whether that human bias can be corrected.

Participants

Jack Balkin

Director, Information Society Project
Knight Professor of Constitutional Law and
the First Amendment, Yale Law School

Sandra Baron

Senior Fellow, Information Society Project
former Executive Director, Media Law
Resource Center

Robyn Caplan

Researcher, Data & Society

Rebecca Crootof

Executive Director, Information Society
Project
Research Scholar and Lecturer in Law, Yale
Law School

Steve Crown

Vice President and Deputy General Counsel,
Human Rights, Microsoft Corporation

Robert Epstein

Senior Research Psychologist, American
Institute for Behavioral Research and
Technology

Kyle Gibson

Yale Poynter Fellow in Residence
Founding Executive Producer, Women in
the World

Ellen P. Goodman

Professor of Law, Rutgers Law School

James Grimmelmann

Professor of Law, Cornell Tech

Srinandan Kasi

Co-Founder and CEO, iPipal Inc.
former Vice President and General Counsel,
Associated Press
former President and CEO, NewsRight

Kate Klonick

Resident Fellow, Information Society
Project
PhD in Law Candidate, Yale Law School

Jim McLaughlin

Deputy Counsel & Director of Government
Affairs, Washington Post

Cass Matthews

Senior Legal Counsel, Google

Filippo Menczer

Professor of Informatics and Computer
Science, Center for Complex Networks and
Systems Research, Indiana University
Network Science Institute

Amy Mitchell

Director of Journalism Research, Pew
Research Center

Kenneth Richieri

former Executive Vice President and
General Counsel, New York Times

Jacob Rogers

Legal Counsel, Wikimedia Foundation

David Schulz

Director, Media Freedom of Information
Access Clinic
Floyd Abrams Clinical Lecturer in Law,
Yale Law School
Partner, Levine Sullivan Koch & Schulz
LLP

Jason Stanley

Jacob Urowsky Professor of Philosophy,
Yale University

Nabiha Syed

Assistant General Counsel, BuzzFeed

Robert Post

Dean and Sol & Lillian Goldman Professor
of Law, Yale Law School