

FALL 2017

# ANALYSIS OF PHILANTHROPIC OPPORTUNITIES TO MITIGATE THE DISINFORMATION/ PROPAGANDA PROBLEM

---

Prepared for:



Authored by:

Kelly Born, Hewlett Foundation

Nell Edgington, Social Velocity

## About the Foundation

The William and Flora Hewlett Foundation is a nonpartisan, private charitable foundation that advances ideas and supports institutions to promote a better world.

Launched in 2014, the Hewlett Foundation's Madison Initiative focuses on strengthening democracy and its institutions – Congress, in particular – to be more effective in a polarized age. The initiative is nonpartisan in its approach and supports organizations across the ideological spectrum – think tanks, advocacy groups, academic researchers and civic leadership organizations – who seek to understand and improve the political system so that our elected representatives are better equipped to solve society's greatest problems.

Learn more at [www.hewlett.org](http://www.hewlett.org).

*Note: The foundation does not lobby or earmark grant funds for prohibited lobbying activities, as defined in the federal tax laws. The foundation's funding for policy work is limited to permissible forms of support only, such as general operating support grants that grantees can allocate at their discretion and project support grants for nonlobbying activities (e.g., public education and nonpartisan research). Additionally, the foundation may fund nonpartisan political activities by grantees in compliance with the electioneering rules. The foundation does not engage in or use its resources to support or oppose political candidates or parties.*



---

# TABLE OF CONTENTS

---

I.	Executive Summary	4
II.	The Disinformation/Propaganda Problem and Its Evolution	7
III.	Potential Solutions: What's Known and Still to be Known	17
IV.	Preliminary List of Organizations Working on Disinformation/Propaganda Problem	30
V.	Preliminary list of Funders in the Disinformation /Propaganda Space	35
VI.	Appendix	38

# I. EXECUTIVE SUMMARY

While the problems of disinformation, misinformation and propaganda are not new, certain aspects of modern technology and communications appear to be contributing to a rapid polarization and democratic deterioration in the U.S. and abroad. Although the William and Flora Hewlett Foundation's Madison Initiative is currently supporting some experimental efforts aimed at addressing the disinformation/propaganda problem, this report explores whether a larger, more sustained investment from the foundation may be warranted.

For the purposes of this report, we have defined the current “information problem” as three-pronged:

- **Disinformation:** *intentionally false* or inaccurate information that is *spread deliberately*, historically by “unknown state officials” but increasingly by other politically interested actors, with a motive of undermining public confidence.
- **Misinformation:** *unintentionally promulgated*, inaccurate information, which differs from propaganda in that it always refers to something which is not true, and differs from disinformation in that it is “intention neutral”.
- **Propaganda:** Information, historically promulgated by state officials but today often also by political opponents, that may or may not be true, but which presents the opposing point of view in an unfavorable light in order to rally public support. This differs from traditional public diplomacy, in a healthy marketplace of political ideas, either when the information is false or, when it contains elements of truth, but relies upon *non-rational means of persuasion* – using emotional appeals not to convince or attract but to disrupt, divide, confuse, or otherwise damage target audiences’ understanding or political cohesion, as a recent *publication* by the National Endowment for Democracy notes.

Given the goals of The Madison Initiative, we have a particular interest in the role that this three-pronged information problem, in combination with social media, has come to play in increasing polarization and tribalism, thereby weakening our democratic systems.

To analyze the disinformation/propaganda problem and potential philanthropic interventions that the Madison Initiative could undertake, we held conversations with 45 experts in this realm. These were evenly divided between formal interviews and more informal conversations with leaders from academia, digital media platforms (both current and former employees), think tanks, and civil society organizations.

The full list of informants and their affiliations, along with the questions asked of them, can be found in the Appendix of this report. The list of informants was developed with staff from the Hewlett Foundation and the Hewlett Foundation Disinformation/Propaganda Advisory Committee based on their knowledge of the key thinkers and influencers in the space. Most informants were asked a subset (rather than the entire list) of questions based on the particular area of expertise the informant possessed.

Interviews ranged from 30-60 minutes. Kelly Born, Program Officer at the Hewlett Foundation, conducted 28 of the interviews alone. Nell Edgington, consultant from Social Velocity, conducted 13 of

the interviews alone. And Kelly and Nell together conducted four.

In exploring how modern disinformation, misinformation and propaganda differ from that of prior generations, several elements appear particularly problematic, including:

- The democratization of information creation and distribution;
- The socialization of information sharing;
- The atomization of news away from reputable brands to individuals;
- The anonymity of shared content;
- The increasing level of content personalization; and
- The privatization and sovereignty of the technology platforms.

All of these factors combine to elevate the internet and, in particular, social media platforms (Facebook, Twitter, increasingly Instagram, and others) as the most critical point for intervention. Most funders are focused on improving the quality of journalism, fact-checking, or news literacy. But the problem, we believe, lies elsewhere – online, and in particular across popular social platforms. The overwhelming insight from the interviews and background research conducted for this report confirmed this. It also highlighted that we simply do not yet know enough about the current problem, how technology impacts it, or the efficacy of proposed solutions, to craft a set of interventions. The area with most promise is investing in a research agenda and team of researchers who can help to further understand and define the problem, as well as explore potential interventions.

Several broad questions remain to be answered, including:

To what extent are disinformation and propaganda actually impacting the democratic process?

In particular, which content, which content sources, and which distribution channels have the greatest impact on the growing problem of polarization and tribalism online?

Which interventions are likely to be effective, and how can they best be distributed?

Can we, and if so *how* can we teach algorithms to distinguish among shades of truth while appropriately balancing free speech concerns?

How do we address human nature's response to sensational, biased, or self-affirming information?

But answering these questions necessitates access to more data from the platforms than is currently available given their concerns about both business interests and user privacy. At the same time, those researchers already working on these questions tend to be siloed between universities and sometimes even between departments within the same university. And those funders who are supporting research are also doing so in siloes. There is a real need, then, to create a comprehensive research agenda and data infrastructure that allows access to and interconnection among the data and those with the ability to gain insight from it.

Subsequently, the opportunity for the Hewlett Foundation may lie in investing in efforts to create and connect that data infrastructure and support its thoughtful analysis.

## II. THE DISINFORMATION/PROPAGANDA PROBLEM AND ITS EVOLUTION

Problems of propaganda, misinformation and disinformation are longstanding, and the forms they've taken have morphed over centuries as governments, communications technologies and policies, and understandings of human psychology have evolved. The problems of today's information environment are both old and new, but a framework for thinking about the critical elements of a healthy information environment for a modern democracy has yet to be developed. Here we seek to determine the hallmarks of a healthy information environment, and to identify what differentiates today's misinformation, disinformation, and propaganda challenges from those of prior decades, or centuries – with a particular focus on the role of the internet and of online social platforms in accelerating polarization and tribalism in democratic politics.

While challenges around propaganda, misinformation and disinformation are longstanding, one interviewee for this report noted that in recent history these issues have been concentrated outside the West:

---

“The issue of fake news, propaganda and social media, if you talk to people in the Middle East or South Asia, they will say Americans are just experiencing what has been a problem for a very long time. People are like, ‘Ah, right now Americans finally feel this pain.’”

---

Indeed, the 2016 presidential election brought the issue to the forefront of American consciousness and since then, there has been a growing desire to figure out what happened and why.

At its core, the disinformation/propaganda problem undermines our democracy. In order for a democracy to flourish, there must be a healthy information ecosystem. As another interviewee argued:

---

“I would say there is a sense of failed democratic discourse in our society and an attack on our core institutions – professional journalism, science, and law – as independent systems of validation around which democratic disagreement lives. Instead, we have competing alternative realities within competing political groups. We have suffered a loss of a shared framework.”

---

Media and information systems certainly play an important role in modern political systems. In democracies, where citizens elect their representatives, the media has long served to inform citizens about social, political, and economic events – helping to provide a common understanding of facts

upon which to debate public policy options.<sup>1</sup> Media, particularly as manifest via investigative journalism, have also helped to hold public officials accountable. The journalism of yore intended to inform citizens about their communities and leaders, not *persuade* them to some specific action or outcome.

Although the hallmarks and values of a healthy information environment in the internet era have yet to be clearly defined, most frameworks agree that, at the highest level, an ideal media environment should put society's welfare at its center. But historical debates around the elements of a healthy information environment have focused on *media ethics, editorial* standards<sup>2</sup> and, more narrowly, *journalism ethics*<sup>3</sup> - typically from the perspective of newsrooms' obligations rather than citizens' needs. More recent reflections on the *elements*<sup>4</sup> of a successful journalism field have expanded the conversation, as have the policies and principles of newer groups like *Wikimedia*,<sup>5</sup> which have likewise wrestled with issues of accuracy and bias. In 2009, *the Knight Commission on the Information Needs of Communities in a Democracy* wrestled with these questions,<sup>6</sup> as did the Federal Communications Commission in their 2011 *Working Group on Information Needs of Communities*.

---

<sup>1</sup> Scholars speak to “political information needs”: the “requisite access to information—typically via the mass media—that citizens need to fulfill their role—primarily informed voting— within a Western, representative, democratic political system. Whereas political information needs address the individual voter, they are less about his or her needs and more about the information necessary to sustain the political system.” They also note that when “institutional trust is high...there will not be significant community information needs.” Brendan R. Watson & Sarah Cavanah (2015) *Community Information Needs: A Theory and Methodological Framework*, *Mass Communication and Society*, 18:5, 651-673, DOI: 10.1080/15205436.2015.1059948

<sup>2</sup> The Ethical Journalism Network highlights five core editorial standards: 1) Accuracy and fact-based communications; 2) Independence; 3) Fairness and Impartiality; 4) Humanity (journalists should do no harm); 5) Accountability and Transparency.

<sup>3</sup> The Society of Professional Journalists highlights several approaches to ethical journalism: 1) Seek Truth and Report It; 2) Minimize Harm; 3) Act Independently; 4) Be Accountable and Transparent. <http://www.spj.org/ethicscode.asp>

<sup>4</sup> The 10 Elements of Journalism identified by Tom Rosenstiel, President of American Press Institute and founder of Pew's Journalism Research program, include: 1) Journalism's first obligation is to the truth; 2) Its first loyalty is to citizens; 3) Its essence is a discipline of verification; 4) Its practitioners must maintain an independence from those they cover; 5) It must serve as an independent monitor of power; 6) It must provide a forum for public criticism and compromise; 7) It must strive to keep the significant interesting and relevant; 8) It must keep the news comprehensive and proportional; 9) Its practitioners must be allowed to exercise their personal conscience; 10) Citizens, too, have rights and responsibilities when it comes to the news.

<sup>5</sup> Wikipedia's Core Content Policies include: Neutral point of view; Verifiability; No original research.

<sup>6</sup> In 2009, *the Knight Commission on the Information Needs of Communities in a Democracy* identified three core needs: 1) Maximize the availability of relevant and credible information to all Americans and their communities; 2)

Strengthen the capacity of individuals to engage with information; and 3) Promote individual engagement with information and the public life of the community. According to the Commission, the press can be seen as having four primary functions with regards to its community – coordinate activities, solve problems, public accountability and a sense of connectedness.



In the internet age -- where gatekeepers have been removed, information creation and distribution is highly fragmented, and citizens are inundated with information – four elements appear key: accuracy, commonality, relevance, and constructiveness.

Foremost across all of these frameworks is the need for accurate, fact-based information “in pursuit of the truth” – acknowledging that while the “truth” may not always be readily clear, verifying information and “*getting the facts right is the cardinal principle.*” Related are the needs for: independence (e.g., ensuring visibility into the political affiliations, financial arrangements or other connections that might constitute conflicts of interest); accountability and transparency (e.g., checking sources, providing audiences visibility into journalistic processes, ensuring anonymity only for those facing threats of danger, etc.); and fairness, impartiality or neutrality (e.g., pursuing a diversity of views, avoiding stereotyping, and providing context – though not necessarily in pursuit of “balance,” given concerns around creating false equivalency).<sup>7</sup>

Second, if obvious, is a need for information that is relevant – identifying and prioritizing key risks and issues of greatest societal importance (social, political, economic, environmental, etc.) while minimizing noise. As Tom Rosenstiel, President of the American Press Institute and founder of the Pew Research Center’s Project for Excellence in Journalism notes, “[Journalism] must strive to keep the significant interesting and relevant [and to] keep the news comprehensive and proportional.”<sup>8</sup> This includes reflecting the needs and concerns of diverse communities of interest beyond traditional media.

A third critical element, largely overlooked in frameworks that preceded the internet’s fragmentation, is the need for a common basis of information. Information can be relevant and even accurate (if incomplete) while being delivered to isolated echo chambers, creating a divergent view of societal issues. But as Daniel Patrick Moynihan famously said, “Everyone is entitled to his own opinion, but not to his own facts.” Scholars have long viewed communities forming as a product of individuals’ “common attention to...issue[s] of collective concern.”<sup>9</sup> In a democracy, a healthy information environment should help provide a forum for public criticism and compromise, helping to “clarify the goals and values of society [with an] implicit... appeal to avoid pandering to the lowest common denominator.”<sup>10</sup> Information environments play “a multifaceted, complex role in fostering a sense of community, communicating, for example, boundaries of who belongs and who does not and a ‘shared symbol system’” or culture... a community is defined by the effort to address some sort of collective

---

<sup>7</sup> According to a recent study by the American Press Institute, today only 31% of Democrats and 8% of Republicans believe that “news from media is very accurate.” [https://www.americanpressinstitute.org/publications/reports/survey-research/partisanship-and-media/?utm\\_source=20170713&utm\\_medium=email&utm\\_campaign=friends](https://www.americanpressinstitute.org/publications/reports/survey-research/partisanship-and-media/?utm_source=20170713&utm_medium=email&utm_campaign=friends).

<sup>8</sup> <https://www.americanpressinstitute.org/journalism-essentials/what-is-journalism/elements-journalism/>

<sup>9</sup> Habermas, J. (1974). The public sphere: An encyclopedia article (1964) (S. Lennox & F. Lennox.Trans.). New German Critique, 3, 49–55. Retrieved from <http://www.jstor.org/stable/487737>

<sup>10</sup> Straubhaar, LaRose & Davenport, p. 477

need and that effort itself helps define the boundaries of the community.”<sup>11</sup> Without some common sense of facts, identity, and view on what policy problems need to be addressed, democracy may prove exceedingly hard to maintain.

Finally, our information environment should be at least minimally constructive. This includes traditional notions of *minimizing harm* by supporting “open and civil exchanges of view...avoiding pandering to lurid curiosity, even if others do [and] considering the long-term implications of the extended reach and permanence of publication.”<sup>12</sup> It should also include newer conceptions of journalism acknowledging the need to not only report on problems, but also on effective potential solutions.

Unfortunately, many of these essential elements of a healthy information ecosystem in a democracy are undermined in our current media environment. While problems of propaganda, misinformation and disinformation are not new, the internet -- and particularly social media and search platforms -- is. Information is now distributed in a fundamentally different way than when print, radio, and television dominated because information was curated and disseminated by media gatekeepers. The new internet information age is different in (at least) six respects:

## 1. Democratization

Perhaps most importantly, both information creation and distribution have been fully democratized. There is no longer a gatekeeper, and with it we have lost any “common set of facts” upon which to found political discourse. As former DARPA Social Media and Strategic Communication lead Rand Waltzman noted in a recent report, *The Weaponization of Information*, any individual or group can communicate to and influence large numbers of others online. This is a feature, not a bug, and has many positive aspects. But it has also removed the gatekeepers, enabling massive fragmentation and inundation in the information environment.

## 2. Socialization

A direct byproduct of this democratization is the switch from “trusted institutions” as gatekeepers, signaling or endorsing the credibility of information, to peer-to-peer sharing. This presents three key challenges:

- a. **Accuracy:** Prior gatekeepers, however flawed they were in execution, were at least in spirit committed to journalistic standards of excellence. Journalists and editors have now been replaced with peer-to-peer sharing, which too often fails to elevate the most accurate, relevant, or constructive content.

---

<sup>11</sup> Brendan R. Watson & Sarah Cavanah (2015) *Community Information Needs: A Theory and Methodological Framework*, *Mass Communication and Society*, 18:5, 651-673, DOI: 10.1080/15205436.2015.1059948

<sup>12</sup> <http://www.spj.org/ethicscode.asp>

- b. **Diversity:** Because people can now choose to have their news filtered through their friend networks, concerns about filter bubbles and echo chambers are widespread. The jury here is still out, however, on how big a problem this is - largely because Facebook algorithms remain proprietary.
- c. **Moderation:** Finally, mixing “social” and news means that political information, previously consumed in moderation by most, now has uninhibited access to an online audience. Prior media modalities required readers to “opt-in” – to seek out MSNBC, Fox News, (or the John Birch Society, before that). Audiences for Fox News have thus remained in the 2-3 million range. But social platforms allow political extremists’ access to apolitical or moderate users, who may be on Facebook looking for cat videos, and are instead exposed to fallacious or polarizing content without seeking it out.

### 3. Atomization

Related to both democratization and socialization is the divorce of news from brand - into stories disconnected from their producer, and from any social cues regarding accuracy or source credibility. Even if citizens could produce their own content, and share it socially, distributors of disinformation could never have gotten a “fake news” story in front of sizeable audiences without this shift from the New York Times or Boulder Camera to primarily receiving news via social platforms. Whereas readers used to be able to readily “consider the source” when selecting news – recognizing that tabloids were located in the grocery store checkout line - now all information circulated on social media is, essentially, coming from a trusted source – one’s friends and family. Once news discovery was divorced from brand, much else followed.

### 4. Anonymization

At the same time that content creation has been democratized, it’s been anonymized. The dominance of name brand newsrooms is dwindling, much less reporters with known bylines or trusted news anchors. Now anyone can create and share content, and while platforms like Facebook require (however ill-enforced) that users use their real names, the rules and structure of platforms like Twitter lend themselves even more easily to anonymous accounts. Two problems emerge:

- a. **Accountability:** Anonymity has eradicated any visibility into conflicts of interest. Internationally, plausible deniability was introduced for state actors intervening in foreign information environments.
- b. **Authenticity:** Taking lack of accountability one step further, anonymity also facilitates the role of bots (“web robots” – essentially computer programmed “if/then” statements that are used to perform highly repetitive operations, such as responding to keywords like “guns” by tweeting out related, often biased, articles). One 2015 *study* found that bots generate around 50 percent of all web traffic. Of course there are “good” bots as well as “bad” which assist in answering customer questions online, or keep weather, sports, and other news updated in real-time. But the ability of bad actors to “game” online information systems allows more extreme or biased ideas to quickly rise to the foreground, amplifying biased perspectives and lending them the appearance of mainstream popularity.

## 5. Personalization

Online immediacy also enables a new level of content personalization. While TV and radio always traveled fast, online content creators can A/B test and adapt micro-targeted messages in real-time. It's well known that internet marketers, issue advocates, and political actors alike now collect unprecedented amounts of data about our online affiliations, interests, and behaviors. Claims abound that “by leveraging automated emotional manipulation alongside swarms of bots, Facebook dark posts, A/B testing, and fake news networks” groups like *Cambridge Analytica* are able to create personalized, adaptive, and ultimately addictive propaganda.

## 6. Privatization:

Finally, as Stanford Professor Nate Persily has observed, these platforms have an incredible degree of sovereignty. Unlike television, print, and radio, which are overseen (at least nominally) by the Federal Communications Commission and (around elections time) the Federal Elections Commission, no one is regulating them – they are self-regulating at best.

Of course, all of these technological changes have taken place in (and contributed to) a world of very different societal norms than existed even a few years ago. One Russian scholar recently noted that, whereas Gorbachev might've been deterred by credible accusations that he was distorting the truth, no such shaming serves to deter Putin (or President Trump). The norms around truthful communication have eroded badly amongst political elites; sharing of mistruths is no longer punished with the same level of public censure.

This is a messy and sprawling problem. Given this, the Madison Initiative is particularly concerned with whether, and how, social media and search platforms are furthering partisan polarization by:

- Exposing citizens to high volumes of factually inaccurate or biased information, thereby misinforming citizens about both politicians and policy issues;
- Filtering news to citizens ensconced in ideological “echo chambers,” thereby exacerbating ideological divisions, incivility, and partisan hostility;
- Sensationalizing news in order to drive “click-bait” traffic, thereby shifting national attention away from core social issues towards content that is more “shareable,” again furthering partisan conflict.

At the same time that these changes are undermining democratic institutions, the same tools and approaches are affecting all of the other policy issues the foundation cares about. In such an information environment, it is hard to imagine how democracy can thrive.

In fact, the rapid spread of new digital technologies has happened without a large-scale understanding of, or appreciation for, its impact on our democratic system, a key problem highlighted by one of the interviewees for this report:

---

**“The Silicon Valley-based firms that have designed our social media don’t fully appreciate their role in public life yet. They are all consistent about how they are just data or technology companies. They must rise to the challenge of appreciating the ways in which their technologies are actually part of modern democracy. We must engage with these firms and get them to be good civic actors before governments start regulating and censoring.”**

---

As Rand Waltzman put it in his testimony to the Senate Armed Services Committee, Subcommittee on Cybersecurity in April, these new digital technologies have fundamentally changed how disinformation /propaganda spreads and influences our democratic processes:

---

“These technologies have resulted in a qualitatively new landscape of influence operations, persuasion, and, more generally, mass manipulation. The ability to influence is now effectively ‘democratized,’ since any individual or group can communicate and influence large numbers of others online. Second, this landscape is now significantly more quantifiable. Data can be used to measure the response of individuals as well as crowds to influence efforts. Finally, influence is also far more concealable. Users may be influenced by information provided to them by anonymous strangers, or even by the design of an interface. In general, the Internet and social media provide new ways of constructing realities for actors, audiences, and media. It fundamentally challenges the traditional news media’s function as gatekeepers and agenda setters.”<sup>13</sup>

---

While some of the emerging interventions to disinformation and propaganda focus on content (improving journalism, fact checking) many interviewees agreed that our current problem is one of distribution based on new technology tools. As an interviewee argued:

---

“Content is not it. When you step back and get away from content that is when you find success. Fact checking is not going to do it. When you go through the data looking for actors you find something more important, you find these huge networks of bots. The best question is how did this happen, how did we allow our systems to become this way? We need to take a multi-disciplinary approach to this problem, an ecosystem view of the problem. The internet and media system has evolved to a different type of thing; maybe it’s time to look at it as a system.”

---

---

<sup>13</sup> Waltzman, Rand, p. 2.

And another interviewee described how the unique problem we are facing right now is one of distribution systems, rather than content:

---

“Journalistic quality is important, but a lot of the solutions will happen when you step away from content a bit. Disinformation/misinformation is a structural issue which is more based on systems that are converging between humans and technology and our understanding of reality and how they are manipulated.”

---

Still another interviewee seemed to agree that it is the technology platforms, and Facebook in particular, where any search for solutions must focus:

---

“The social media companies are best positioned to help understand and address these problems, and Facebook in particular. I say that for two reasons 1) scale – more people get more news from Facebook than any other platform and 2) technical infrastructure and the design of the algorithms. It is challenging right now that Facebook has the largest audience but its algorithms are likely the most susceptible to abuse and virality of misinformation. The algorithms that power Google and Google News and those that power Twitter or LinkedIn are less susceptible. And so you have the greatest likelihood of this problem emerging on the platform with the greatest audience.”

---

And indeed, because the platforms are so new and so critical to any potential solutions to this problem, there is an opportunity to help the platforms both understand their role in a strong democracy and take a more active role in it, as one interviewee pointed out:

---

“We have a unique opportunity to involve social media platforms in democratic experiments. Facebook could become a platform for experimenting with voter turnout, jury led systems, voter guides. Getting Facebook to imagine ways it could contribute, would mean the pressure to do something policy wise disappears. They are doing a lot with fact checking, and they could do more to get people to express themselves. For example, in several democracies the exit polling system is broken. This is one of the core things necessary to make democracy work, an external check on the outcome. Exit polling hasn’t worked well since 2000. Facebook would be a pretty good platform for doing an exit poll, even if Facebook distanced themselves from the outcome.”

---

Working with the platforms to understand better and contribute more fully to the health of our democratic systems is the most leveraged place where the Hewlett Foundation can have an impact.





### III. POTENTIAL SOLUTIONS: WHAT'S KNOWN AND STILL TO BE KNOWN

As discussed above, the most promising interventions lie around improving the current curation and distribution of information online. Here we explore a range of issues: which actors are most incentivized to address the problem, the need for caution in pursuing policy solutions without a clearer understanding of the problem itself, how varying solutions will play out in different geopolitical contexts, and the human and psychological motivations underlying readers' appetite for "fake news." We conclude by underscoring the need for more researcher access to data, as well as cross-institution collaboration to support these inquiries.

Over the last five to ten years, there have been some concerted efforts to address the disinformation/propaganda problem. These efforts can be categorized into 5 key areas:

- On the **creation** of content side, most efforts are focused on:
  - Improving journalistic quality
  - Identifying and holding accountable creators of dis/misinformation
- On the **distribution** of content side, most efforts are focused on:
  - Providing support and advice to, or increasing accountability of, the platforms (Facebook, Google, Twitter)
- On the **consumption** of content side, most efforts are focused on:
  - Improving/increasing fact checking options
  - Increasing citizen news literacy

Efforts to improve journalistic quality, while important, alone will not solve the problem. After all there is accurate, balanced information available on many if not all policy issues. It is simply falling on deaf ears.

The latter efforts to "clean up the mess once it's out there" face numerous challenges. In its initial years fact-checking itself became polarized, with liberals supportive and conservatives often highly skeptical. However, in recent months, fact-checking - labeled differently - has become increasingly common amongst conservatives (with liberals of course skeptical of the conservative efforts). Moreover, behavioral economists have long noted that fact-checking inaccurate information seldom corrects the misbelief, and often makes it worse. Even if it was effective, it is virtually impossible to scale, with the average fact-check taking 8 to 15 hours, compared with the pittance of time required to create and distribute falsehoods.

News literacy holds more promise, but is similarly hard to scale and requires a long time horizon. It is also hard to imagine the average citizen, who can barely be bothered to vote, or read the full article, investing the time to quality-check most information that comes their way - much less to develop the level of literacy necessary to successfully detect a well-resourced and relatively sophisticated foreign propagandist. News literacy efforts can only address part of the problem.

Moreover, several interviewees expressed concern about any solution that is funded by the platforms because it is at least partially focused on what is best for the business interests of the platform, not necessarily what is best for democracy. An interviewee said:

---

“There are a lot of organizations that are doing analysis of potential solutions, but I feel like there are a lot of ties between those organizations and the platforms. They are not independent from the greater problem; they are funded by tech companies. I really think the answer is an independent one. The independent angle is good. We need independent advocates and researchers who are outside of the reach of the friendly industrial business complex. Academics are good, except when institutes are funded by platform money.”

---

Indeed, most interviewees agreed that solutions would never come from the platforms themselves, and argued that there needs to be external, independent actors or entities (whether governments, public, or stakeholders) putting pressure on the platforms to make changes. As one interviewee described:

---

“To be completely realistic, Google or Facebook are never going to release or do an algorithmic audit; it’s how they make money. You can’t go after that. The best way to deal with that is to work with them directly. I think they are willing to make choices to deal with this, but it needs to be an external party that holds them accountable. There needs to be a UN for the Internet. The issues span across nations and borders. When a regulation happens in one country, people just move the servers to another country where it’s more convenient. And it needs to be outside of the industry, more of an academic entity.”

---

There have also been many efforts aimed at simply bringing people together to talk about the problems. Some interviewees argued that the Hewlett Foundation investing in more convenings of people working on the disinformation/propaganda problem would not be additive as there are already many convening efforts that are underway:

---

“My feeling is that we are doing a lot of convening; I’ve been to several events. I would encourage Hewlett to not do more convenings. The Harvard group and the Wilson group are the vehicles for convening with policymakers. And Stanford already does a great job of convening academics.”

---

Many interviewees also cautioned that any solutions to the disinformation/propaganda problem must be global in scope. This is not just an American problem and because the platforms have global reach (for example 90% of Facebook users are outside of the United States) the implications of any intervention that includes the platforms will also have a global reach. As this interviewee described:

---

“Disinformation/misinformation is not a problem that can be solved within one country, it’s a global issue and the platforms are global. Any measure taken in one political/regulatory context needs to be considered globally. One of the big problems with regulatory responses in Europe is a failure to care what they mean for users in repressive countries. This is not something you can circumscribe around one nation state. And particularly while platforms do respond to jurisdiction-specific regulations, their algorithms and terms of service are consistently globally applied – they don’t change across borders. So any approach needs to be a global approach.”

---

In fact several interviewees cautioned that any potential solution could have backfire effects that actually serve to limit freedom and democracy in other countries. Already authoritarian governments in other countries are using changing social media platform policies as an excuse to shut down social media accounts of political opponents and reformers. One interviewee put it this way:

---

“There is tremendous risk in a number of countries where governments are regulating in this space. The perils of regulation are not something people are talking about. It’s really real in a lot of countries (beyond the U.S.) where these conversations embolden tyrants.”

---

And while it is unlikely that the U.S. government will regulate the platforms any time soon, there is potential interest in regulation among European Union policymakers, as one interviewee pointed out:

---

“Somewhere between most promising and most likely is leadership from the European Union. The EU is the global political body that is most likely to show some leadership in this domain and what Europe does will have an impact on other technology markets. Silicon Valley doesn’t always like what Europe does, but they will learn they need engagement there.”

---

And some interviewees expressed concern about even democratic governments over-regulating technology to ill effects:

---

“The goal should be to not let Europe overregulate. It’s quite possible that Europe will go too far and that will stifle the good stuff of social media. The conversations about hate speech in Europe are quite different than in the U.S. People in Europe frame conversations about political speech as a balance between free speech and election

interference. And there is sentiment that Brexit and Trump are examples of free speech tipping too far. There is a danger that Europe will overregulate.”

---

But what became increasingly evident over the course of this report’s investigation is that before we can hope to effectively address the problem of misinformation, disinformation and propaganda, we must understand much more about the problem and its causes. An overwhelming message delivered by almost everyone interviewed for this report was that we simply don’t know enough yet about to what extent and how misinformation and disinformation have impacted our democratic systems, especially (given our focus) in ways that are exacerbating the problem of polarization. We cannot hope to focus on the right interventions if we don’t yet fully understand the problem. Therefore, the first step in finding a solution, or set of solutions, is to better understand what exactly the problem is.

First, we must better understand the extent to which misinformation, disinformation and propaganda are actually impacting the democratic process. Although we all suspect that misinformation, disinformation and propaganda dramatically impacted the 2016 U.S. presidential election, we don’t know that for a fact.

We also don’t know whether the problem of dis/misinformation that surfaced during the 2016 presidential election affects the entire right/left political spectrum equally, or if it is a fundamentally different problem at different parts of that spectrum, as one interviewee pointed out:

---

“Imagine that we came to believe that the problem of disinformation/misinformation in the 2016 election is fundamentally different than what the internet has done to discourse in the center and the left. How would we know that and how would it influence our investment in solutions if that were the case? My guess is that it should make a big difference.”

---

Fundamentally, we need to understand where the problem is caused by a lack of knowledge (so that greater news literacy might be a solution, for example, to help deter innocent sharing) versus when the problem is caused by devious manipulation of information (where literacy is not the problem, ill intent is). As one interviewee noted:

---

“The challenge for anyone who wants to do something about misinformation is being able to distinguish between innocuous ‘I don’t know’ vs. the more motivated form of ignorance where people claim certain facts that are clearly incorrect.”

---

And if dis/misinformation did impact the outcome of the election, we don’t know which content, which sources, and which distribution channels had the greatest impact on the election outcomes, as an interviewee explained:

---

“There isn’t any really good baseline metrics for making any descriptive claims about to what extent misinformation is a problem, to what extent we can make a clear distinction between misinformation and ignorance...It’s really important to come up with baseline metrics where we can make the distinction between the informed/uninformed/misinformed citizens. That is fundamental to this entire enterprise. We must pinpoint where people are motivated to provide incorrect answers as opposed to innocently answering incorrectly. That is a really important question, and we don’t have any good metrics in the aftermath of the 2016 presidential election and it’s really important someone do a baseline study which allows us to set down this empirical baseline which we can then use to track changes over the course of the Trump administration.”

---

While both civil society and the platforms themselves have already made some attempts to address the problem, we don’t know which, if any, of those efforts have been effective. While there are studies about the impact of fact-checking and fake news flags by the platforms, there have been no comprehensive studies analyzing the impact of the various measures platforms have taken. Even when some interventions have been debunked, people are still investing in them. An interviewee described it this way:

---

“The whole story of regulation of online communication in the last two years is about one solution being worse than the previous problem. When Facebook took humans out of the news process and instead used algorithms, it made it worse.”

---

Indeed, we don’t know how the algorithms help or hurt the problem. And we don’t know how to teach those algorithms what is truthful content and what is untruthful content. As one interviewee described:

---

“The move to social media as a principal source of news and information, and the fact that social media experiences are driven by algorithm, and the fact that algorithms are motivated by what user is likely to engage in and not by truthfulness of the content, these have gathered to be a perfect storm for the virality of untruthful content. The challenge and the needle hole that must be threaded is that those are the exact dynamics that make the virality of truthful content and gives rise to Black Lives Matter, the Tea Party, and the Arab Spring. So now you are getting into a distinction: we want all these dynamics in play for truthful content but not for untruthful content. While we can argue there is objectively something truthful, an algorithm is unable to attack that on its own. So while there may be objective truth, our identification of it is only possible through a human lens.”

---

And how we define misinformation and disinformation itself is part of the problem. An interviewee pointed out that the definition of misinformation and disinformation is not purely objective:

---

“If you are going to say, fake news is not okay, how do you define fake news? All journalists make mistakes. Are you going to say fake news is news that is deliberately false? How do you do that? What’s the role of people who are sharing it? Even if you identify a piece of news is deliberately wrong, how do you operate against that if it’s a post that includes a URL, can Facebook block the sharing of that URL on their own platform? Looking to Facebook and Twitter to solve this is unrealistic. Definitions of fake news are just too complicated for an algorithm to get it right.”

---

While the algorithms are contributing to the problem, as one interviewee argued, underlying this is human nature and humans’ inability to actually do what they say they want to do, so we have to figure out a way to bake that into the algorithm:

---

“When your entire business model is Pavlov’s dog hitting the button to get another biscuit, it doesn’t encourage you to start getting rid of some biscuits. It’s like tending to a child, if your entire parenting model is based on keeping the child happy, they would never eat vegetables. So that creates a challenge for Facebook. And you have the same challenge in other parts of society, for example if you ask someone on a telephone poll if they think voting is important 90%+ say ‘Yes.’ Then if you ask them if they intend to vote, 80%+ say ‘Yes,’ but only 60% show up. Facebook has these same dynamics – people say ‘I don’t want fake news,’ but they don’t want to take any responsibility for managing that. They want the app to magically show them the content they want. This creates the necessity to have the machine think for the user because the user will not self-curate. The machine has to be trained.”

---

Other interviewees agreed that this is fundamentally a human behavior problem, and it is crucial to investigate that piece of the puzzle. People may claim a distaste for “fake news” but still share it. In fact, recent research on fact checking actually demonstrated that stories deemed “untrue” on Facebook were actually shared more frequently than those that were “true.” Therefore, we must better understand how to hack human psychology to make the spread of disinformation and propaganda less appealing. One interviewee described it this way:

---

“At the end of the day, these algorithms are simply trying to create a machine version of human behavior. And the crux of the problem is the hypocrisy of human behavior. We say we don’t like fake news, but we really like fake news. For *The National Enquirer*, *News of the World*, and celebrity magazines their entire business model is based on this vice of human behavior. And so on some level as long as algorithms are programmed

to respond to other humans, and unless there is new thinking on how to hack human behavior, the algorithms will never fix the problem. Fact check flags are ineffective, content removal is unscalable, but the one thing that makes people less likely to distribute misinformation is when they are shamed by their friends. So how do you create that sense of shame or the sense of responsibility that would trigger shame? That has to be answered on a psychological level. What makes people eat their vegetables? There may be lessons from completely different fields of work that would have answers to this. Smoking cessation or obesity reduction programs are a lot closer to the problems we are trying to solve than some of the computer science or political science approaches to this problem.”

---

The vast majority of interviewees agreed that there is still much to understand about the problem itself. Most believe that the platforms, and Facebook in particular because of its reach, are of paramount interest to research.

But Facebook data is not currently as accessible as researchers would like because of Facebook’s legal, privacy and business concerns. While a couple of interviewees argued that using Twitter’s publicly available data would be enough, most argued that the small scale of Twitter users (only 24% of online adults use Twitter, whereas 79% of online adults use Facebook)<sup>14</sup> makes Twitter less important as a data source for really understanding the problem and its potential solutions.

Facebook allows both public and private posting, but private posts are where the bulk of highly targeted politically advertising and information transfers happen. Facebook has traditionally not been willing or able to release private post data because of privacy concerns. Similarly Google can provide aggregated search data, but not individualized data about what types of people are searching for what types of data.

To really understand the problem, one interviewee argued that researchers need access to both public and private data and the websites interacting with the platforms:

---

“We need access to both public and private posts. You can do a lot of research on Twitter, but not at the scale that we need. In addition to the platforms, you need to be crawling the individual website that put out this info, starting from well-known media sites (CNN, Fox) all the way to the obscure ones.”

---

Therefore, before we can research the problem more effectively, researchers need access to a much deeper set of data. As one interviewee explained:

---

<sup>14</sup> Pew Research, Social Media Update 2016

---

“If I were going to put \$10 million into the (disinformation / propaganda) problem, I would focus on the research side and build the data infrastructure.”

---

Indeed, many interviewees bemoaned the lack of access to comprehensive data about what is actually happening in the disinformation/propaganda ecosystem. As one interviewee put it:

---

“One issue is the availability of data and building both access to existing data and new data sources to facilitate the right kind of research. The first category of questions about what’s actually happening should be easy to answer. But the reason they are hard is that the data is hard to access; it’s fragmented, different people have different pieces, no one can see the full picture. We need to build a panel, an environment where there’s a large number of people we can track (like Nielsen does) and monitor what they are doing through surveys, experiments, incentives. Setting that up would be good.”

---

While there are many researchers and academics working on better understanding the disinformation/propaganda problem, this sort of comprehensive data infrastructure ecosystem simply doesn’t exist. As one interviewee succinctly stated: “If you are at Google or Facebook you have access to mountains of data, but not if you are at Stanford or Columbia.”

Still another interviewee argued that the research conducted to date has been fractured and too narrow. Rather, a broader data set and body of researchers needs to be created:

---

“I’ve seen groups that do good work, teams that have been working with Facebook and Google during Brexit and French elections, trying to find misinformation. I don’t know if those things are scalable and how targeted they are. They seem like whack-a-mole, one election, and a narrow focus. We need something much more scalable. But no one is doing that yet. You need to go to a place where people have an engineering commitment, to house the data and analyze it. We need to do something from outside the platforms. You need academic institutions that can own this.”

---

The reason this comprehensive data does not yet exist is that there are serious risks that the platforms feel they are facing in releasing their data, as one interviewee described:

---

“The other challenge affecting all the platforms is a political one... Right now tackling the issue of fake news post-election creates a belief that you are just part of the liberal conspiracy because it is believed as a way to clamp



down on conservative speech, but since the presidential election there has been a significant rise in deliberate misinformation from the left. When you have experiments like Facebook's partnering with fact checking, most of those organizations are perceived to be left-leaning by those on the right, so that gives the impression that Facebook is outsourcing content evaluation to left wing activists."

---

Indeed, concern about their users' privacy along with concerns about damaging their brand combine to make digital technology companies for the most part unwilling to share their user data. As one interviewee described:

---

"There is a big barrier in data sharing among the platforms. Facebook feels like there's a ton of risk to them if people look at their data and find something that looks bad for them, or cherry pick the data to make it look bad for them. The risk/reward calculation doesn't look good. There's not a huge amount of upside from people doing research from their data. The platforms are very worried about that stuff. Nevertheless, there are a ton of questions which the platforms have the ability to answer and no one else does."

---

But Facebook has grown to play such an important role in our information ecosystem that they must be made to understand and embrace that role, as one interviewee argued:

---

“Facebook has the data to know where the people who were subject to Russian attacks ended up going to vote or not in the 2016 election. But Facebook doesn’t want to be blamed for having created the situation that picked Trump by perhaps suppressing the black vote. Facebook will only work with outside researchers who will sign an NDA so that Facebook can control what they publish. Facebook has very strong control over the data and a very self-serving model of what they decide to publish. That needs to be remedied. They shouldn’t have the ability to sit on so much data and decide to only release it when it fits their interest. We need to create enough social pressure on Facebook that they decide to participate while creating enough protections in the system to address their legitimate concerns without addressing their illegitimate concerns. I don’t think they will do it unless they feel under serious social pressure.”

---

Another interviewee went even further suggesting that the biggest contribution the Hewlett Foundation or other philanthropists could make would be to encourage Facebook to share their data:

---

“There is data inside of Facebook that is not being shared, they have the biggest data set. One intervention the Hewlett Foundation could make would be to come up with some sort of a data trust where Facebook could get all the assurances it needs, that the data won’t help their competitors. If Facebook could commit to giving enough of their data to an independent research group, which can analyze everything Facebook has, that would transform everything.”

---

Some of those interviewed argued that because the platforms have disincentives to share data or make changes to their algorithms and many governments are unwilling to regulate the platforms, a solution could be found by forming a small group of high-level influencers who could quietly encourage the platforms to share data, participate in independent research studies, make changes, etc. As an interviewee described:

---

“A small, influential panel of people could call up Google, and say ‘Can we talk about how you index Twitter?’ If an individual researcher tried to do that he wouldn’t get past the receptionist. If instead the research effort were supported by a credible, legitimate group, the researcher could get some depth on these questions. A group of people with gravitas could have productive relationships with the platforms that could encourage proactive, as opposed to reactive, responses from them.”

---

To truly gain access to the data, one interviewee argued that you could fund a team of engineers inside the platforms who could run various tests and then scrub private information from the results:

---

“You need people who work for Facebook who can look at the data to run an experiment for you and share aggregated data. That’s something that these platforms could enable, a 10-person engineering team to handle these requests. You would have to pay them for their time and omit data that is not okay to share legally. You can imagine the platforms doing that, they might do it as part of their media outreach. Google News Lab makes publishers happy by providing tools and data sets. You could imagine them providing this kind of service on Google trends data. It’s fairly easy for them to build a class that is ‘probably Republican’. And then run a study, Independents did this, Republicans did this. That requires a team inside the platform that is dedicated to helping research groups.”

---

But this model assumes that platforms could somehow be expected to voluntarily share data even when it runs counter to their business or reputational interests.

Another interviewee suggested that a way to make access to data more acceptable to the platforms is to create a larger repository that they could all share, as the interviewee explained: “One way to make it more palatable is to have all the platforms contributing data to some sort of repository. This would be a little less intrusive, so we aren’t climbing all over their secret sauce. That might make it feel better to them.”

Beyond the need for access to data is also a need to connect all of the actors working to understand that data. Part of the goal of building a robust data infrastructure would be to connect the (until now) disparate information and actors engaged in understanding the problem. As several interviewees pointed out, there are tremendous silos between academics researching disinformation/propaganda, platform employees, and journalists, among others. Here’s how one interviewee described it:

---

“On the research side, most of the (platform) companies don’t share data, and Facebook doesn’t play well with academia. And academics don’t have a strong sense of what goes on at Facebook, yet we are pretty sure that Facebook data scientists do know what is going on. We need to get social scientists to pivot to do more investigative work and help academics to pivot to help data journalists. We need to provide a platform for more integration.”

---

Once the hurdle of gaining access to the data is overcome, interviewees argued that it should be housed within a University with a multi-disciplinary approach:

---

“Pick a University or a consortium, house it in a research center at a major university with the funding and technical infrastructure and the ability to curate

it properly. Set up very clear rules about how to get access and what security assures other researchers, while protecting the privacy of the users and the legitimate concerns of Facebook.”

---

So a critical element necessary to make the data more accessible is greater collaboration at a higher level between the platforms and academic researchers, where risks to the platforms are mitigated and greater research access to the data is achieved. One interviewee described it this way:

---

“Academic/private sector collaborations would have tremendous value.... A model that I might think about is from development economics. A huge amount of that work has involved randomized trials of different interventions, not things the researchers are making up, but things NGOs are doing on the ground. So you have a constant flow of interventions, - some that make a difference, others that don't - but the value for science is big. This teaches us a lot. If there is appetite for trying to build solutions of different kinds, having a partnership between Silicon Valley tech people and academics who know how to structure things could be huge...Connecting the private sector and academia could be really helpful.”

---

Finally, any attempt to better understand the problem needs to be forward, not backward looking. As one interviewee described:

---

“We need to think about what's coming down the line. For example, it's going to become possible in 5 years to create artificial videos of someone speaking that is indistinguishable from reality. As that becomes routinized, we'll be dealing with fake reality, not just fake news. That poses a serious challenge to propaganda, news and how we are going to deal with that. Researchers should be encouraged to investigate the next set of challenges. Research needs to be looking forwards not backwards. We need to have a funding strategy that is nimble enough going forward. That means thinking about how different actors are going to behave in the future. You have to think about virtual reality and the new technologies and how they will be utilized. And how all our devices will merge into one thing (phone, computer, TV).”

---

So the key role, many interviewees argued, for a foundation like Hewlett to play is in supporting (through money and network building) analysis by interdisciplinary groups of researchers to better understand the problem and its potential solutions. As an interviewee explained:

---

“This is how foundations can be most effective by funding academic groups that want to understand what is happening with the news ecosystem. Not only what’s going wrong and to what extent we are being influenced, and test solutions, but also what is the right strategy. But it all starts with having a large amount of data and the computer scientists and social scientists to understand it. The right place to do this is in academic institutions where these people co-exist.”

---

With greater access to both private and public data at all of the major platforms, academic researchers can begin to understand whether disinformation and propaganda is impacting our democratic systems, and if so how. Armed with this greater understanding of the problem, we can then better test and implement the most effective interventions. Interviewees overwhelmingly agreed that there likely is no silver bullet solution to the problem of misinformation, disinformation and propaganda. Indeed the recurring theme among interviewees is that once we understand the problem better, it will likely require a holistic solution.

## IV. PRELIMINARY LIST OF ORGANIZATIONS WORKING ON THE DISINFORMATION/PROPAGANDA PROBLEM

As mentioned earlier in this report, the problem of disinformation and propaganda has become paramount in recent months. Therefore, the list of organizations researching, thinking about and intervening in the problem has grown notably. We have initially identified almost 20 organizations working in this space – a non-exhaustive list, inclusive only of those centers with two or more researchers working on the problem, and with a principally domestic focus (there are of course still many more working abroad, or comprised of much smaller teams). Most have been doing this for less than five years, some launched as recently as late July. Below, they are broken down into three groups: academic institutions, think tanks, and policy advocates.

The list below is a non-exhaustive summary of entities in which foundations concerned about current disinformation problems might invest, with the expectation that different funders would focus on different groups according to their different emphases (e.g., the Hewlett Foundation would focus on the subset that have a particular emphasis on the problem of polarization). More detailed information on these entities and a list of individual researchers working in this space is included in the Appendix to this report. Note that we are certainly aware that there are additional groups and perspectives that could and should be included in this list. We welcome input on who these groups may be.

*Because many groups' strategies are evolving, the below information is accurate as of Summer 2017. Names are listed alphabetically.*

### Academic Institutions

- The Belfer Center for Science and International Affairs at Harvard Kennedy School  
Led by the former campaign managers for Hillary Clinton and Mitt Romney, alongside experts from the national security and technology communities including Facebook and Google, the project aims to identify and recommend strategies, tools, and technology to protect democratic processes and systems from cyber and information attacks.
- The Berkman Klein Center for Internet & Society at Harvard University  
Led by Yochai Benkler, Berkman's research is focused on how public discourse is influenced by the internet, how artificial intelligence impacts it, as well as how children use the internet and learn what is true. They have received funding from several foundations including the Ford Foundation, the Gates Foundation, the Knight Foundation, and OSF.
- The Knight First Amendment Center at Columbia University  
Launched in 2015 by the Knight Foundation (and a smaller grant from Hewlett) the First Amendment Center works to defend and strengthen the freedoms of speech and the press in the digital age through strategic litigation, research, and public education.
- London School of Economics

There are two entities at LSE. First is the ARENA Project within the Institute of Global Affairs, which Hewlett has funded and which seeks to analyze the root causes of distorted information and hate speech in the European electoral context. The second is The Truth, Trust and Technology Commission, which develops strategies for the transformation of information for the public good.

- MIT's Institute for Data, Systems, and Society (IDSS)

Their interdisciplinary team of 100+ researchers include academics from data science, political science, behavioral economics, anthropology, social psychology and decision theory. They are now seeking to launch a new research program, Social Networks, focused on how information impacts people's decision-making, social norms, ideologies, and the spread of ideas.

- The News Integrity Initiative, CUNY Graduate School of Journalism

The News Integrity Initiative at the CUNY Graduate School of Journalism is a \$14 million fund investigating the roots of mistrust in news and investing in solutions for building trust and public support for quality journalism. NII was launched by a consortium of funders including Facebook, the Craig Newmark Philanthropic Fund, Ford Foundation and others.

- NYU's Social Media and Political Participation (SMaPP) Lab

A very well-respected, small research team. Upcoming research plans focus on how the digital world impacted the U.S. election in terms of echo chambers, exposure to information, political knowledge, political participation, and interaction with traditional media. They are seeking \$1.4 million over 3 years and have received some funding from the Knight Foundation.

- NYU's Stern School of Business, Center on Business and Human Rights

Led by Michael Posner, head of the Stern School's business/human rights program, this two-year project for the World Economic Forum's Global Future Council on Human Rights is exploring steps digital platforms can take to address dissemination of extremist content, political propaganda and false information. Stern is now seeking to extend this work beyond its current 2-year mandate.

- Oxford Internet Institute

This multidisciplinary 45-faculty team explores issues including privacy, security, e-government/e-democracy, virtual economies, smart cities, and more. Their smaller program on Computational Propaganda examines the paths misinformation travels and potential digital interventions. They have been foremost in analyzing online political bots.

- Stanford Program on Social Media and Democracy

Stanford's well-respected multi-disciplinary research team, which includes key experts from law, comparative democracy, and computer science, will focus on understanding information integrity, the regulatory frameworks available, how artificial intelligence impacts opinions and behavior, and how it can be altered in more beneficial ways.

- Tow Center for Digital Journalism

Led by their new Director of Research Jonathan Albright (an expert on the emerging online disinformation ecosystem), their research focus is evolving. Past projects focused on understanding the benefits and challenges of automated reporting tools, large-scale data and impact evaluation metrics, story impact on audiences, etc. They have been funded by the Knight Foundation.



## Think Tanks

- Atlantic Council's Digital Forensic Research Lab

The Lab was established in 2016 to expose falsehoods and fake news, document human rights abuses and verify and report the authenticity of conflicts and global events.

- The Berggruen Institute

The Berggruen Institute's mission is to develop foundational ideas and shape political, economic and social institutions for the 21st century. In late 2016, they launched the Democracy for the Digital Society Project, led by Toomas Ilves, former president of Estonia. They have been convening Silicon Valley leaders, but next steps are unclear. Current plans appear to be underway for developing pilot projects in Italy, Canada and the City of London.

- Center for Democracy and Technology (CDT)

This well-respected think tank led by Michaelle Richardson is launching a project to understand how Americans perceive online news, how online communities influence interaction and identity, and how to integrate journalism standards into the digital world. They are seeking \$300,000 for this project.

- Data and Society

Data and Society's research team of 23 is focused on the social and cultural issues that have arisen from new technologies. They cover a range of issues, from the future of labor to "big data," and have a two-person research team looking at online information manipulation. They have received funding for their broader programming from several foundations including the Ford Foundation, MacArthur Foundation, the Knight Foundation and OSF.

- First Draft News

This small team, with funding from Google, develops fact-checking tools for journalists. They recently moved to the Shorenstein Center at Harvard and are in the midst of a strategic planning process to determine their goals and financial requirements going forward.

- Institute for the Future's Lab for Digital Intelligence

IFTF is launching a new project, the Lab for Digital Intelligence led by Sam Woolley (formerly of the Oxford Internet Institute), to conduct research and research-informed advocacy work at the intersection of technology and political communication with a special focus on social media and the Internet, with a particular focus on computational propaganda and politics.

- New America Foundation's "Viral"

"Viral," a new project which will be led by Peter W. Singer, Strategist and Senior Fellow at New America, and Emerson Brooking, a research fellow at the Council on Foreign Relations, will explore how social media is being used in politics, war, media, crime, terrorism, and beyond.

- Pew Research Center

Pew Research Center's Journalism Research team has recently partnered with Knight, Gates, and Ford on a new examination of how the public assesses the legitimacy of news and information and their sources.

- RAND Corporation

The RAND Corporation is launching a project called "Truth Decay" led by Jennifer Kavanagh to research the definitions, causes, and consequences of the erosion of respect for facts and evidence in politics and the blurring of the distinction between facts and opinions.

## Policy Advocates

- Electronic Frontier Foundation

This advocacy group focuses on user privacy, impact litigation, policy analysis, and grassroots activism, and is particularly engaged with the tech/engineering community. They employ 15 lawyers and have an annual budget of \$7-8 million.

- Free Press

This well-respected team focuses on internet and press freedom and has an annual budget of \$4-5 million. They receive funding from the Ford Foundation, OSF and Democracy Fund.

- Open MIC

Their small research team is focused on open and affordable internet access, privacy rights, and fair, open and transparent data analytics. They receive funding from the Ford Foundation and Media Democracy Fund, among others.

- The Open Technology Institute at New America Foundation

Similar to Free Press, this group focuses on internet and press freedom, but engages less on media and journalism issues. They have a staff of 20-30.

## V. PRELIMINARY LIST OF FUNDERS IN THE DISINFORMATION /PROPAGANDA SPACE

Just as the list of research entities and think tanks/civil society organizations working on the disinformation/propaganda problem has grown dramatically in the past several months, so too has recent interest among foundations.

Most funders in this space are focused on improving the quality of journalism, or supporting fact-checking or news literacy (this includes the Knight Foundation, Democracy Fund, MacArthur Foundation, and the McCormick Foundation). Many fewer are focused on the intersection of democracy and the internet, seeking to improve the distribution of information (like the Ford Foundation, Open Society Foundation and, potentially, the Hewlett Foundation). Of this latter group, the focus tends to be not solely on research but also on policy advocacy, supporting civil society interventions, or translating research into action. Very few appear to be focusing primarily on funding the research itself.

Below is a brief summary of key current/potential funders. *As with the nonprofits noted above, because funder strategies are evolving, this information is accurate as of summer 2017; please let us know of others who should be included:*

- Democracy Fund

DF is undertaking a range of interventions aimed at newsrooms, platforms, and the public in the United States. They have current grants to the Engaging News Project, the Trust Project, the Coral Project, and Free Press, with one pending for the Trusting News Project, as well as four grants in the Knight Prototype Fund focused on misinformation. Like the Hewlett Foundation, DF brings a bipartisan focus to their work.

- Ford Foundation

In 2017 Ford invested \$5 million to support a variety of the aforementioned research institutions and civil society organizations through their Internet Freedom and Creativity and Free Expression programs. Ford's interests and focus are highly aligned with the Hewlett Foundation's in this space, though their Internet Freedom work tends to have more of a social justice emphasis.

- Gates Foundation

With the goal of promoting evidence-based information and public discourse, Gates has funded several research projects (Pew, Gallup, MIT Media Cloud) in partnership with other funders, and is considering others. Through their existing media portfolio, they have funded fact-checking organizations globally including Africa Check and India Spend. Gates does not as of now have a dedicated budget for disinformation. They have a one year discretionary fund of \$1M at the New Venture Fund for research to understand the dis/misinformation ecosystem and efforts to build audience trust in fact-based information. Gates has a global focus, with disproportionate interest in Africa and India, as well as Germany, France, U.K. and the U.S.

- Hewlett Foundation

Through its Madison Initiative, the Hewlett Foundation has been investing roughly \$15M annually since 2014 – increased to \$20M per year through 2020 - on nonpartisan efforts to improve U.S. democracy, with a specific focus on making Congress more effective in an era of polarization. The foundation’s approach includes providing support to groups across the political spectrum. The Madison Initiative is considering an additional investment to support research on the role that the internet and social media in particular, are playing in increasing polarization and tribalism in the U.S.

- John D. and Catherine T. MacArthur Foundation

Investments have focused on supporting journalism and documentary film, with one or two grants in the information research domain (e.g., Tow at Columbia).

- Knight Foundation

Knight is supporting two lines of work focused on disinformation. First, helping to lead a research effort with several partners (Ford, Gates and OSF), focused on understanding trust in journalism, testing the public’s response to different kinds of content, and improving the flow of quality news and information. Second, additional efforts to understand and support trust in journalism including the recently announced Commission on Trust, Media and Democracy with Aspen Institute. Knight recently held a prototype competition focused on the same.

- Media Democracy Fund

MDF is a re-grantor focused on protecting the public’s rights in the digital age. MDF is currently exploring questions around antitrust and competition, including potential strategies to address public interest concerns and harms caused by dominant digital platforms.

- Mozilla Foundation

The Mozilla Information Trust Initiative (MITI) was launched in August 2017 to 1) advance research on how misinformation impacts Internet Health, 2) encourage product experimentation to address ideological polarization, civil discourse, and the spread of false information online, and 3) to grow a grassroots movement around solutions to misinformation through events like MisinfoCon. In 2016, the foundation’s budget was \$20 million. MITI started with a \$1 million commitment in 2017 to explore how Mozilla might engage in the space, leveraging its network of thousands of volunteers, 400 thousand individual donors, 2 million supporters, and 300 million Firefox users.

- The News Integrity Initiative, CUNY Graduate School of Journalism

NII is the one academic institution noted above that is also funding. Through grants, applied research, events and other activities, NII is supporting work that fosters trust between newsrooms and the public, and nurtures constructive, inclusive community conversations. NII also supports research focused on media manipulation. Grantees include Free Press, Listening Post Collective, The Coral Project, Center for Media Engagement, Data & Society, First Draft News, among others. NII’s international partners currently include Internews and the European Journalism Centre.

- Omidyar Network

ON recently announced a \$100 million investment to fight misinformation and hate speech and to support investigative journalism. They have also supported several fact-checking efforts including Full Fact and the International Fact-Checkers Network. ON's work tends to be global in orientation although they have also supported U.S.-specific efforts combatting online hate speech such as the Anti-Defamation League's Center for Technology and Society.

- Open Society Foundations

OSF is focused on both research on the manipulation of news and information, and supporting its translation to civil society organizations. They have recently allocated \$1 million in the U.S. to this problem, and future budgets are being determined. OSF is likely also to have a global focus.

- Rita Allen Foundation

RAF is focused on improving communications of science, and on understanding the behavioral science behind how audiences process information and form opinions.

- Robert R. McCormick Foundation

McCormick's focus is primarily on trust in media and primary grants have been to support news literacy, press freedom and supporting Chicago news entities.

- Skoll Foundation

Skoll is also considering work in this space, but is still in a learning phase. They are looking at the full "ecosystem" of news including both supply (digital and print content and distribution) and demand issues, and are likely to first focus on the former. Initial investments could focus on experiments exposing the public to diverse media diets, providing infrastructure for investigative reporters, helping to create a more sustainable high-quality news supply chain, etc.

## VI. APPENDIX

In the pages that follow are background materials for this report. These are:

- a. Emerging research agenda for the field
- b. A list of all report informants, their affiliations and types of expertise
- c. The script and list of questions used in the interviews

# A. POTENTIAL DISINFORMATION / PROPAGANDA RESEARCH AGENDA

Based on our interviews, the following is intended as a high-level, consolidated, draft list of key questions that must be answered for the field to be in a better position for sound policy decision-making in 3-5 years - setting aside for the moment both who would undertake what research, and whether current data access is sufficient to answer these questions.

## Explanatory Questions: What is the Problem?

1. **Primary Question:** How/to what extent is online disinformation/propaganda directly affecting the performance of U.S. democracy?

How successful has disinformation been in **affecting people's political beliefs / knowledge accuracy**?

What impact does disinformation/propaganda have on **partisan beliefs**? i.e., Are citizens **polarizing** because of exposure to biased information or are polarized citizens seeking out biased information - How strong is **causality**?

2. If the internet/social media do appear to be affecting U.S. democratic performance, then:

Who are the **primary producers of disinformation** in the U.S. and elsewhere (government officials, political parties, media companies, individuals on the internet, etc.)?

- i. Which **producers** of disinformation/propaganda are most effective at impacting electoral outcomes, and how concentrated are they?
- ii. \*Are different purveyors **more likely to propagate polarizing** disinformation/propaganda?
- iii. How do candidates, government officials and the legacy media **respond to or become accomplices in purveying falsehoods** spread through on-line disinformation?

Which **channels** (technology platforms, TV, radio, print, etc.) are most effective at spreading disinformation/propaganda?

- i. How do those **channels interact** with both online platforms and each other around disinformation/propaganda? e.g.,
  1. **Does social media follow traditional media** in emphasis on topics, or are issues discussed differently on social media than in traditional media?
  2. Has social media **changed how "legacy media" produces content** and/or the way in which this content is consumed?

What **strategies and tactics** (e.g., bot networks) are most prevalent/successful in distributing disinformation? e.g.,

- i. On **which internet platforms** is the spreading of falsehoods most pervasive, and why?

- ii. How and when is **computational propaganda (bots) effective** in purveying falsehoods dangerous to democracy?

How does **audience** behavior increase/decrease the impact of disinformation/ propaganda?

- i. How significant are **echo chamber/filter** bubble challenges?
- ii. In what ways can news audiences be meaningfully **segmented**? e.g.,
  1. How much does internet information consumption differ by segment?
  2. How does the spread of disinformation/propaganda vary along the right/left political spectrum?
  3. Are particular segments more prone to sharing polarizing content?
- iii. What cues do people use to **assess the credibility** of an online news story (e.g., brand, byline, expert citations)?
- iv. How does a story's **perceived credibility influence willingness to read**, watch, or share it with other people?

What **types of content** are most common/problematic, particularly in relation to political polarization (e.g., stories supporting in-group party narratives, attacking voters from the opposing political camp, discouraging opposition party members from voting, etc.)?

- i. What factors of disinformation/propaganda **content** determine its **virality**?
  - ii. What are the **dominant narratives** across different audience segments and distribution channels, particularly in terms of driving polarization?
3. What can be learned from other countries/movements/moments in history? e.g.,

**Historical /anthropological/ international/ social science analogs** (Nazi /Soviet/ Western propaganda, biology, Guttenberg bible, smoking cessation programs)?

**Different countries' historical efforts to regulate** disinformation?

Foreign efforts to **regulate disinformation in the digital age** – e.g., Are these efforts different or merely an extension of previous efforts (e.g., Germany vs South Korea)?

## Experimental Questions: Where / How Should We Focus Our Efforts?

1. How do different platforms' roles differ? e.g.,
  - a. What is the distinct **impact of different platforms**, particularly in spreading polarized content?



- b. What are the **economic incentives of each platform**?
  - c. What is the role of platforms in **personalization of content** and political **micro-targeting**?
  - d. What **policies have the platforms adopted and rejected** to deal with the spreading of disinformation?
2. Which of the many **proposed interventions work** with which audience segments, to change attitudes/behaviors? e.g., What should the platforms do differently to:
    - a. Help **elevate “quality” content**?
    - b. \*Help **reduce polarization and/or partisan incivility**?
  3. How can the platforms be **incentivized** to do these things?
  4. What are the **human rights implications of platform changes/policies in different political contexts**?
  5. How can researchers understand and **communicate the manipulation mechanisms at play** in legible ways to media, tech, industry, academic, and civil society communities?

## Practical, Legal and Regulatory Questions:

1. **Primary question:** What is the universe of legal/reform options for platform responsibility for disinformation/propaganda?
2. General questions:
  - a. How and **where do we draw the lines** between accurate, overly biased, inflammatory, and/or patently false information?
  - b. What are the **necessary ingredients for social information systems** to encourage a culture that values and promotes truth?
  - c. How do we **balance privacy and accuracy, authenticity, and accountability**?
3. Specific legal questions:
  - a. How/should traditional views of **free speech** be updated to address current technologies and their applications?
  - b. How does/should regulation of disinformation **compare to regulation of other types of prohibited speech** (e.g., terrorist recruitment/incitement, hate speech, obscenity, intellectual property, bullying, etc.)?
  - c. What are legal obstacles to **intermediary liability** for self-regulation of fake news?
  - d. In what new ways might **antitrust laws** be applied to the platforms in a manner that would improve the information environment?

- e. Do, or how can the platforms begin to **better coordinate** efforts (e.g., common definitions of and standards for dealing with disinformation/propaganda)?
  - f. Would it be feasible/useful to expand who is able to pursue **liable claims** (e.g., beyond individuals)?
  - g. Could/should **“takedown” requirements for IP violations** be extended to other forms of content?
4. Enforcement questions:
- a. Could **FTC regulation** of “deceptive acts or practices” be extended beyond commercial goods?
  - b. Could **FCC regulations** of false information be expanded beyond broadcast TV/radio (and could/would they be meaningfully enforced)?
  - c. Could/how could **FEC regulations** requiring disclosures of paid online elections communications be better enforced?
  - d. Could/how could **platforms** be incentivized to effectively self-regulate?

## B. LIST OF INFORMANTS

<b>Name</b>	<b>Affiliation</b>	<b>Expertise Type</b>
<b>Adam Berinsky</b>	Political Scientist – MIT	Academic
<b>Adam Sharp</b>	Former Head of Twitter News	Platform
<b>Alexios Mantzarlis</b>	Lead of the International Fact-Checking Network – Poynter	Reformer
<b>Alice Marwick and Robyn Caplan</b>	Researchers, Data and Society	Think Tank
<b>Amber French and Ben Lennet (consultant)</b>	Media Democracy Fund	Reformer / Funder
<b>Anne Applebaum</b>	Co-founder - ARENA at London School of Economics	Academic / Journalist
<b>Barbara Raab</b>	Program Officer, Journalism - Ford Foundation	Funder
<b>Ben Nimmo</b>	Atlantic Council	Academic
<b>Ben Scott</b>	New America Foundation; SVN	Academic / Reformer
<b>Brendan Nyhan</b>	Political Scientist - Dartmouth College	Academic
<b>Chancellor Williams</b>	Program Officer, U.S. Democracy - Open Society Foundations	Funder
<b>Craig Aaron</b>	President & CEO - Free Press	Reformer
<b>Janet Haven</b>	Director of Programs - Data & Society Research Institute	Academic
<b>Jeff Jarvis and Molly deAguiar</b>	CUNY, News Integrity Initiative	Academic / Reformer

<b>Name</b>	<b>Affiliation</b>	<b>Expertise Type</b>
Jessica Clark and Vince Stehle	Media Impact Funders	Funder Affinity Group
Jonathan Albright	Research Director - Tow Center on Digital Journalism, formerly Assistant Professor of Communications - Elon University	Academic
Jonathan Kartt	Program Officer - Rita Allen Foundation	Funder
Josh Stearns	Associate Director for Informed Participation - Democracy Fund	Funder
Joshua Tucker	Professor of Politics, NYU and Co-Director NYU Social Media and Political Participation (SMaPP) lab	Academic
Krishna Bharat	Founder - Google News	Platform
Lori McGlinchey	Senior Program Officer, Internet Freedom - Ford Foundation	Funder
Manami Kano	Deputy Director for Global Media Partnerships - Gates Foundation	Funder
Matt Perault, Andy O'Connell, Katherine Brown	Global Policy Development - Facebook	Platform
Michael Connor	Open Mic	Reformer
Matthew Gentzkow	Professor of Economics - Stanford University	Academic
Michael Posner	Human rights lawyer, runs the Stern School business/human rights program	Academic
Nate Persily	Professor - Stanford Law School	Academic
Peter Pomeranzev	Co-founder - ARENA at London School of Economics	Academic / Journalist
Phil Howard	Professor of Sociology, Information and International Affairs - Oxford Internet Institute	Academic

<b>Name</b>	<b>Affiliation</b>	<b>Expertise Type</b>
<b>Rebecca MacKinnon</b>	Director, Ranking Digital Rights Project - New America Foundation	Academic / Reformer
<b>Richard Gingras</b>	Vice President - Google News	Platform
<b>Sam Gill</b>	Vice President for Communities and Impact - Knight Foundation	Funder
<b>Sam Woolley</b>	Director of Research, Computational Propaganda Project at Oxford Internet Institute, incoming Researcher at Facebook	Academic, Platform
<b>Scott Carpenter</b>	Managing Director, Jigsaw - Google	Platform / Reformer
<b>Shanto Iyengar</b>	Professor of Political Science - Stanford	Academic
<b>Solomon Messing</b>	Director of DataLabs - Pew Research Center, formerly Facebook	Platform / Academic
<b>Stacy Donohue</b>	Partner, Governance & Citizen Engagement - Omidyar Network	Funder
<b>Stephen Lewandowsky</b>	Stephan Lewandowsky, Professor, Cognitive Psychology - University of Bristol	Academic
<b>Susan Morgan</b>	Susan Morgan, Senior Program Officer, Information Program - Open Society Foundations	Funder
<b>Tim Hwang</b>	Fellow - Data & Society	Platform / Reformer
<b>Tom Black</b>	Senior Program Officer, Measurement Learning and Evaluation - Gates Foundation	Funder
<b>Tom Glaisyer</b>	Program Director - Democracy Fund	Funder
<b>Yochai Benkler</b>	Faculty Co-Director - Berkman Klein Center for Internet and Society at Harvard	Academic

## C . INTERVIEW SCRIPT AND QUESTIONS

[*Note: Most informants were asked a subset (rather than the entire list) of these questions based on the particular area of expertise the informant possessed.*]

**Interview Introduction:** “Thanks so much for taking the time to talk with us today. As we mentioned in the email, as part of the Hewlett Foundation’s Madison Initiative we have long been concerned with how new technologies and the changing media landscape are affecting U.S. democracy. While we are supporting some early grant experiments in the disinformation/propaganda arena, we are also exploring whether the Hewlett Foundation should consider a sustained investment to address these information problems. We are in the very early stages of exploring this, so we’d like to ask you some questions.”

1. How are you defining the current misinformation, disinformation, and propaganda problem?
2. What are the key outstanding questions to more clearly define this problem?
3. How should we go about answering these questions, and who is best positioned to answer them?
4. What do we know and what do we still need to know about potential solutions to this problem?  
What research would be helpful to better understanding how to solve the problem?
5. What do you see as the most promising potential solutions?
6. What do you think are the key levers to addressing the disinformation/propaganda problem?
7. What are the research or learning questions that need to be answered in order to find effective solutions, for example: the impact of the media environment on people’s views, do the platforms become more open or more closed over time, how receptive are they to support, does showing someone information contradictory to their world views change their opinions, what impact antitrust litigation would have on the media environment?
8. Is there a role, and if so what is the role, for philanthropy to play in attempting to address these problems?
9. What are the risks (at systemic and individual grantmaker level) to the solutions you find most promising?
10. What would be the key indicators of progress in the solutions you’ve identified as most promising?
11. What are the biggest challenges from a platform perspective?
12. What are you doing already to address this issue?
13. How are you handling this in other countries with different free speech / privacy considerations?
14. As a funder, how do you define the disinformation/propaganda problem?
15. Where are you focusing your efforts to find solutions and why?
16. Which solutions do you think the Hewlett Foundation is uniquely positioned to support in this area?

17. Who else do you think we should be talking to or what resources should we be reading?
18. Any other thoughts?

