

# Stanford

Geopolitics, Technology,  
and Governance  
*Cyber Policy Center*

## Stanford Communication

## How to Report Responsibly on Hacks and Disinformation

*10 Guidelines and a Template for Every  
Newsroom*

### About the Authors

**Janine Zacharia** is the Carlos Kelly McClatchy Lecturer in Stanford's Department of Communication. In addition to teaching journalism courses at Stanford, she researches and writes on the intersection between technology and national security, media trends and foreign policy. Earlier in her career, she reported extensively on the Middle East and U.S. foreign policy including stints as Jerusalem Bureau Chief for the Washington Post, State Department Correspondent for Bloomberg News, Washington Bureau Chief for the Jerusalem Post, and Jerusalem Correspondent for Reuters.

**Andrew Grotto** is director of the Program on Geopolitics, Technology and Governance and William J. Perry International Security Fellow at Stanford's Cyber Policy Center and teaches the gateway course for graduate students specializing in cyber policy in Stanford's Ford Dorsey Master's in International Policy program. He is also a Visiting Fellow at the Hoover Institution. He served as Senior Director for Cyber Policy on the National Security Council in both the Obama and the Trump White House.

The run-up to the 2016 U.S. presidential election illustrated how vulnerable our most venerated journalistic outlets are to a new kind of information warfare. Reporters are a targeted adversary of foreign and domestic actors who want to harm our democracy. And to cope with this threat, especially in an election year, news organizations need to prepare for another wave of false, misleading, and hacked information. Often, the information will be newsworthy. Expecting reporters to refrain from covering news goes against core principles of American journalism and the practical business drivers that shape the intensely competitive media marketplace. In these cases, the question is not whether to report but how to do so most responsibly. Our goal is to give journalists actionable guidance.

Specifically, we recommend that news organizations:

- Adopt a playbook—we present one below—of core principles and standards for reporting on newsworthy events involving false, misleading, and hacked information;
- Establish a repeatable, enterprise-wide process for implementing the playbook;
- Commit their senior leaders to ensuring the success of these initiatives across vast newsrooms siloed into distinct areas of responsibility—from political coverage to national security reporting to social media teams.

## Background

[UNESCO](#) offers the following definitions of disinformation, misinformation, and mal-information.

- Disinformation is “false and deliberately created to harm a person, social group, organization or country.”
- Misinformation is “false but not created with the intention of causing harm.”
- Mal-information is “based on reality, used to inflict harm on a person, social group, organization or country.” Doxing, or the release of hacked materials, is a form of mal-information.


For the sake of simplicity, we refer to these three forms of untruths collectively as propaganda—information, especially of a biased or misleading nature, used to promote or publicize a particular political cause or point of view.<sup>1</sup>

Of course, propaganda is an ancient phenomenon, but it has become an especially hard problem in this era of social media and political polarization. There’s been much useful work in this realm. [First Draft News](#) has been fighting it online since 2015. The organization trains journalists and publishes useful guidelines on reporting responsibly “[in an age of information disorder](#).” Data and Society’s report on the manipulation of the media by Alice Marwick and Rebecca Lewis is a masterful overview of the problem and a must read on this topic, especially the case studies of reporting leading up to the 2016 presidential election. [The American Press Institute](#) has issued strategies for truth-telling, and many others, including our Stanford colleague [Renee DiResta](#) and the [National Endowment for Democracy](#), are sounding the alarm to journalists that they need to adjust to this new reality.

We recognized during our participation in [Stanford's Information Warfare Working Group](#) that newsrooms still could benefit from a straightforward protocol for news situations involving various forms of propaganda. Andy and Janine decided to pair their expertise—his, as a former White House official in 2016, and hers, as a long-time journalist. After months of research and consultations in California and many conversations with reporters from national and local news outlets, in February 2020, we met with senior editors and reporters in New York and Washington DC at the Associated Press, The New York Times, NBC, The Washington Post and National Public

---

<sup>1</sup> See, e.g., entries for “propaganda” at [Dictionary.com](#) and the [Oxford English Dictionary](#).



Radio. What we found was a confirmation of the usefulness of the endeavor, eagerness to engage with our suggestions, and valuable feedback that helped inform this end product.

Our practical, actionable playbook is designed for editors and reporters to use when the provenance of material they receive is an essential dimension of its newsworthiness. Examples where provenance matters include hacked material from newsmakers, content associated with campaigns to spread disinformation, and bot-driven conspiracy theories in social media about newsworthy people or events.

Naturally, each newsroom will make its own judgments. But we wanted to empower editors and reporters with concrete guidelines and help ensure that decisions aren't made in an ad hoc manner, especially in a competitive, breaking news situation. The more news outlets that embrace a new set of norms, the more resilient American media will be against exploitation by malicious actors.

## **Some core conundrums**

Reporting on information of murky provenance takes time, which can clash with the competitive instinct to publish as quickly as possible. To deal with this, our playbook includes suggestions for how to minimize the pause and to create story placeholders that allow for necessary reporting.

Another problem is the danger of bots like those used during the 2016 Russian campaign. "They can do something even the most energetic person can't do, which is send you the same message, or a variance of the same message, a hundred times," Lee Ross, a Stanford professor of psychology and expert on human judgment, said at a recent meeting of Stanford's Information Warfare Working Group. Repeat exposure wins, Ross said, and "when something is congruent with our beliefs, we don't second-guess it." The working group includes psychology researchers, engineers, political scientists, former U.S. officials and other area experts.

Fact-checking and labeling as false an item of propaganda risks exposing more people to the information than would have ordinarily seen it. Psychology research shows people tend to choose

to believe information that accords with their beliefs even if it is debunked. Editors need to be mindful of this when deciding when and whether to cover certain kinds of viral propaganda.

Another challenge, we recognize, is that the threat can evolve. New tactics to dupe journalists and lure them into coverage will emerge. This doesn't excuse newsrooms from preparing for the known threats. They just need to be willing to revise and update their protocols as tactics change.

Finally, in our consultations with news organizations, we observed a "prisoner's dilemma" dynamic that highlights the imperative for norm development around disinformation. The prisoner's dilemma demonstrates how even when actors share a common interest, coordination problems can emerge among them that, if left unresolved, make all of them worse off. The coordination problem for respectable news organizations is this: they share a common


interest in informing their readers and facilitating quality civic discourse, but they also face intense competitive pressures for scoops, professional accolades, and, ultimately, readers' eyeballs in a crowded field where novelty sells. There is a tension between being first to report on the one hand, and, on the other hand, taking the time to get the full story.

Professional journalism's answer to this coordination problem has been to establish a corpus of more or less consensus principles of responsible reporting. These principles are comprised of the written standards that news organizations produce and unwritten professional norms of conduct. And while there are certainly subtle differences in standards and norms across respected news organizations and reporters, these differences amount to minor variations of what are essentially consensus principles. The principles help reporters and newsrooms navigate the tension between being first and getting the

### **The Prisoner's Dilemma**

The prisoner's dilemma is a classic of game theory and involves a pair of co-conspirators who value liberty above all else and have been separately detained for interrogation. The police present each of them with the same deal: betray your co-conspirator, and you will get off easy while your co-conspirator bears the brunt of punishment. Or, stay silent and hope that your co-conspirator doesn't betray you, because if they do you bear the brunt of punishment and thus end up worse off than if you confessed. If both co-conspirators stay silent, however, the police's case against them falls apart and they both walk. Under these circumstances, the co-conspirators face strong incentives to betray each other, unless there are mechanisms available to them that facilitate cooperation despite them being detained separately and unable to communicate in the moment.





story right, while at the same time serving as accountability criteria for what respectable, professional journalism looks like.

When it comes to propaganda, the corpus of relevant principles is immature, at best, and we have seen no evidence of it being institutionalized in a manner that gives them consistent, predictable weight across respected news organizations. The result is that news organizations are operating in a norms vacuum. This must change.


## **How we got here**

The 2007 Virginia Tech shooting coverage was the first time The Washington Post used a live blog for a breaking news event. “Individuals were all of a sudden publishing things with far less than the usual oversight,” recalled R.B. Brenner, who was Metro editor at the time. The newsroom’s mentality shifted from prioritizing the print paper to prioritizing the online edition. Editors had to make decisions much more rapidly and adjust on the fly to changing norms.

News publishing today is this phenomenon on steroids. While the major mainstream news outlets remain most influential, the bevy of online news outlets competing for eyeballs has further sped up the editorial process. The competition for clicks can at times cloud the risks of moving forward with a story involving information of dicey origin. The explosion of digital propaganda means that newsrooms must now integrate new norms into their decision-making process on how to deal with this new reality.

## **Journalistic self-restraint**

There is a long history of journalists refraining from publishing, particularly in the national security realm. In 1958, when New York Times military affairs reporter Hanson Baldwin spotted an unusual plane on a German base and later determined it was a secret U.S. U-2 spy plane, The Times never published the story despite its obvious newsworthiness. Other news outlets including The Washington Post knew of the plane and didn’t publish either, at the urging of the U.S. government. Recently, the major news outlets refrained from publishing the name of the whistleblower who informed Congress of concerns about President Donald Trump’s call with the




Ukrainian president, even as right-wing news outlets and others—including Trump himself—tweeted the alleged name. Examples of self-restraint extend to coverage of sexual assault; often, names of victims aren’t reported without consent. In reporting on suicide, journalists routinely abide by expert [recommendations](#) that urge reporters to avoid sensational headlines and to avoid focusing on the means of death so as not to foster suicide contagion. News organizations filled with seasoned journalists—who are trying to be sensitive to national security interests while being cognizant of the public’s need to know—can model similar good behavior when reporting on propaganda.

Legacy news outlets can chart the course for everyone else who wants to be a credible, fact-based news organization. Even if their role as conventional gatekeepers has eroded, what The New York Times or Associated Press or The Washington Post does often dictates whether something will be talked about endlessly on cable news or makes it into the national zeitgeist. In 1961, a Stanford professor wrote in The Nation about U.S. training of forces in Guatemala for an invasion of Cuba. The New York Times followed the story four months later, prompting President John F. Kennedy to remark: “But it wasn’t news until it appeared in The Times.” Even today, publication in a news outlet like The Times focuses national attention on an issue even if a story is previously circulated on 4chan, Twitter or a less prominent news outlet. Yes, the material may be out there. But that is not a justification to automatically publish it on your site.

We recognize that in the national security realm, reporters often have the luxury of time if they are dealing with an exclusive. This is different than a WikiLeaks-type dump. It’s a harder decision for reporters, who are innately wired to want to be first. Reporters win kudos from their colleagues and competitors when they break a story. When Janine worked at Reuters in the late 1990s, she was judged against AP and AFP on the speed of her news alerts. The explosion of online competitors has made this priority even more urgent for the major news outlets.

What is needed is a speed bump built and overseen by newsroom leadership, with buy-in from across the organization. The norm of responsible reporting shouldn’t be to be the first and the loudest and the most prolific on social media. It should be to tell a story in the most responsible way that is also in the public interest. “The single most important thing is to fight the impulse to publish immediately,” Philip Corbett, The New York Times standards editor, told us.



In the early 2000s, amid a series of scandals at [The New York Times](#), [The New Republic](#) and [USA Today](#), involving prominent journalists who invented sources, newsrooms began tightening the rules on anonymous voices in stories in an effort to restore their damaged credibility. They also resolved that reporters should be more transparent in their stories about the motivations of sources, a point that is especially relevant now. “There was a big push at one point to address all those issues and include whatever you learned in the story [about the source’s motivation], so long as you weren’t exposing the identity of the source,” said former New York Times Washington bureau chief Philip Taubman. Discussions focused on the question of “how much do we owe the readers to tell them about the provenance of what we got,” he added.

There is a lively debate occurring within newsrooms and among reporters about how to cover disinformation campaigns, address misinformation, and report on leaked materials in stories. What has proven elusive so far is a consensus set of best practices and, just as important, a template for how to implement them. The playbook and implementation template we present below are intended to help fulfill this need.

# Newsroom Playbook for Propaganda Reporting

## 10 Guidelines

1

**Develop newsroom social media guidelines—and require all reporters to abide by them.** It is critical in these situations to fight the impulse to publish—or tweet—immediately. Commit instead to being first, responsibly. For example in the event of an extremely newsworthy hack, have the top editor send an organization-wide email instructing all staff not to live-tweet the content. Instead, indicate to readers that you are aware of the development and your reporters are working to determine the provenance of the material.

2

**Remember that journalists are a targeted adversary and see yourself this way when digesting disinformation or hacks.** Ask yourself: Are we being used here? Be on the lookout not only for obvious email dumps but also for direct messages sent via social media from dubious sources who may not be who they purport to be. Familiarize everyone in your newsroom with this minefield so they are aware of the risks.


3

**Beware of campaigns to redirect your attention from one newsworthy event to another — and don't reflexively take bait.** In 2016, the one-two punch of the Access Hollywood tape, followed less than 60 minutes later by Russia beginning the drip-release of John Podesta's emails, illustrated that news organizations want to be on high alert for stories intended to redirect the news cycle. This doesn't mean ignoring the late-breaking event; rather, it means covering the event in a manner that appropriately contextualizes the timing and substance of the event as potentially part of a disinformation campaign.

4

**Break the "Pentagon Papers Principle:" Focus on the why in addition to the what. Make the disinformation campaign as much a part of the story as the email or hacked information dump. Change the sense of newsworthiness to accord with the current threat.** Since Daniel Ellsberg's 1971 leak of the Pentagon Papers, journalists have generally operated under a single rule: Once information is authenticated, if it is newsworthy, publish it. How it was obtained is of secondary concern to the information itself. In this new era, when foreign adversaries like Russia are hacking into political campaigns and leaking material to disrupt our





democracy and to favor one candidate, journalists need to abandon this principle. That is not to say reporters ought to ignore the hacked material if it is newsworthy. But high up in the story they need to focus on the material's provenance. The why it was leaked as opposed to simply what was leaked. "You're in a whole different universe where foreign governments are trying to game the American democracy, especially the First Amendment privileges of the press, to benefit themselves and the candidate that they want to support," Taubman said. "And news organizations have to recognize that the Pentagon Papers Principle cannot apply in those cases. They have to have a different standard."

In other words, authentication alone is not enough to run with something.

5

**Build your news organization's muscle for determining the origin and nature of viral information.**

A responsible newsroom would never take the authenticity of leaked or other non-public content at face value because the authenticity of the content goes to the very heart of its newsworthiness. In the digital age, the same is increasingly true of provenance: the who, why, when and how of content's journey to the public domain may be an essential dimension of its newsworthiness. Establishing provenance, however, will in many cases require technical skills that few reporters possess. News organizations have options for filling this need, which range from establishing a dedicated, in-house digital provenance team with the necessary skills, to forming partnerships with other organizations to pool resources and build shared capability. The latter may sound like a stretch. But news organizations are already collaborating in areas like fact-checking.

6

**Learn how to use available tools to determine origins of viral content.** Reporters do not need advanced skills or degrees in data science to perform basic digital provenance analyses. Still lacking is a dream tool that could automatically tell reporters who first put something up on the Internet. But applications like [Hoaxy](#), [Graphika](#), [CrowdTangle](#) and [Storyful](#) help interpret trends and content on social media. The learning curve for these tools is not steep, and reporters who invest time in developing basic proficiency with them will often be able to develop a first-order approximation about provenance that could inform story development.

7

**Be explicit about what you know about the motivations of the source and maintain that stock language in follow-up stories.** Make sure that this guidance comes down from the top editors and is on a checklist of desk editors and copy editors so there are layers of oversight. There should be equal guidance to reporters who are active on social media that they prominently feature the provenance of the material and its goals in their distribution of this information. If the provenance isn't immediately known, focus your teams on answering that question. When there's a news imperative to cover a story, acknowledge that provenance is a question mark and explain in the story why the origin of the material is critical.


8

**So the provenance doesn't get lost in follow-on stories or sidebars, consider having a box or hyperlink attached to every story on the topic with stock language reminding the reader of the motivation of the leak and why the news outlet is publishing the information. Extend this practice to any accompanying photos, videos or other content.**

For example, stock language for the 2016 DNC hack reporting might have read something like this: "These emails were hacked by Russian operatives to undermine Hillary Clinton's campaign. The xxx is reporting on the portions that are deemed to be in the public interest and is refraining from reprinting those messages that are solely personal in nature."

9

**Don't link to disinformation. If you do, make sure it is a no-follow link.** We noted in our consultations with major news outlets that most are already independently deciding not to link directly to disinformation, an example of the kind of organic norm development that we seek to promote with this report. When news outlets link to disinformation, the content and its source (e.g. site, group or user) get amplified in people's feeds and in search engine algorithms. To avoid such amplification, refrain from linking to questionable content. Instead, describe the information with text and explain to the reader why you aren't linking to it. Alternatively, link to the content using a "no-follow link." Technologist Aviv Ovadya has explained how to do this in First Draft's report [here](#). Actions such as these signal to search engines to not count the link as a "vote" in favor of the target page's quality, which would improve its ranking and exposure. As Cornell Tech University expert on online information technologies Mor Naaman warns, however: "Remember that search engine and social media platforms may consider reader clicks on the link as a signal for interest, thereby contributing to the direct propagation of the



linked page.” For content that is authentic, he suggests, bring the file under your own domain name, instead of linking to a third party whose web platform and associated content you can’t control. This has the added benefit of drawing and keeping traffic on your site.

**10**

**Assign a reporter to cover the disinformation/propaganda beat if you haven’t already.**

Especially in the run-up to the election, having a reporter writing about information manipulation is recommended.

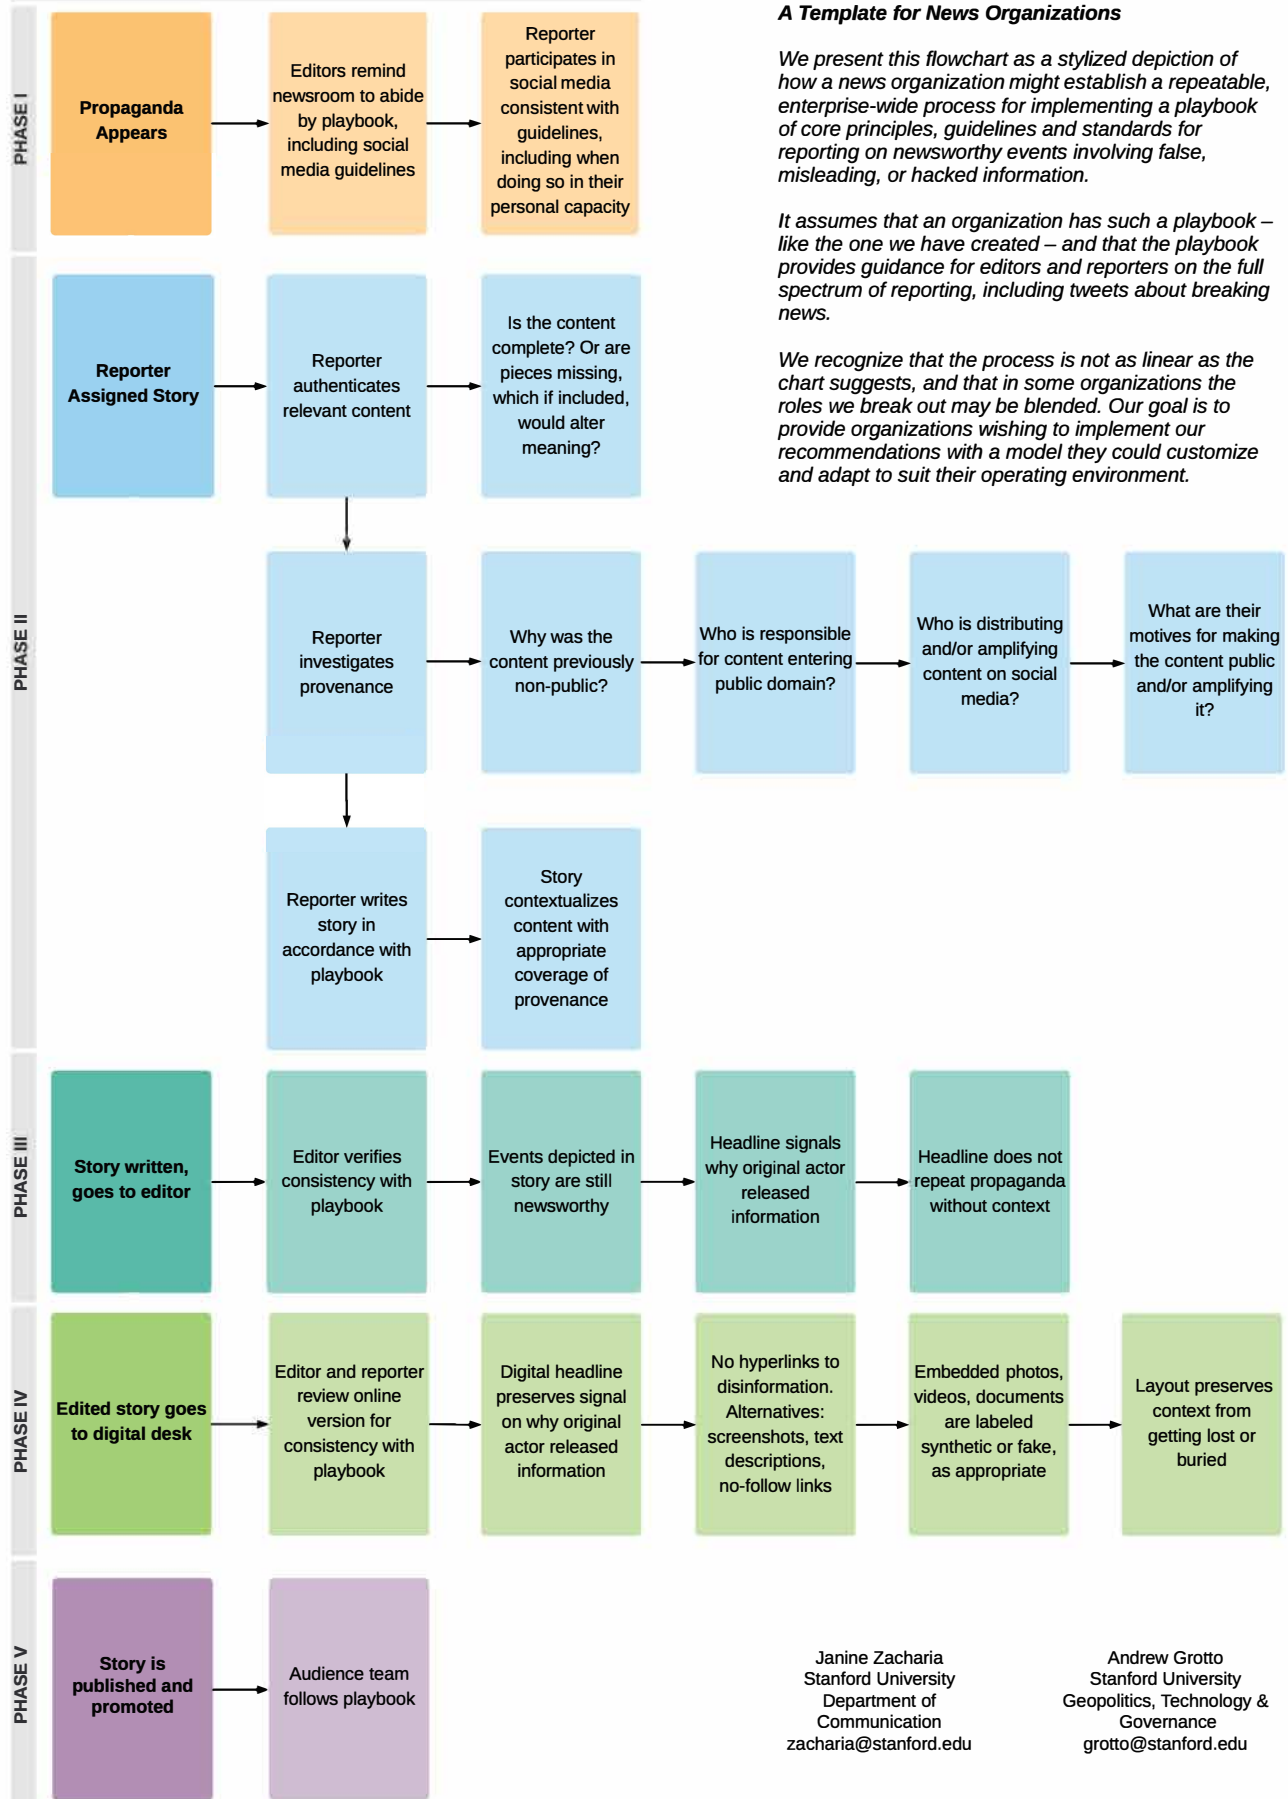
# IMPLEMENTING THE PLAYBOOK

## A Template for News Organizations

We present this flowchart as a stylized depiction of how a news organization might establish a repeatable, enterprise-wide process for implementing a playbook of core principles, guidelines and standards for reporting on newsworthy events involving false, misleading, or hacked information.

It assumes that an organization has such a playbook – like the one we have created – and that the playbook provides guidance for editors and reporters on the full spectrum of reporting, including tweets about breaking news.

We recognize that the process is not as linear as the chart suggests, and that in some organizations the roles we break out may be blended. Our goal is to provide organizations wishing to implement our recommendations with a model they could customize and adapt to suit their operating environment.



Janine Zacharia  
Stanford University  
Department of  
Communication  
zacharia@stanford.edu

Andrew Grotto  
Stanford University  
Geopolitics, Technology &  
Governance  
grotto@stanford.edu

## Appendix: Annotated Template

### PHASE I

#### Propaganda appears

- 1. Editors remind newsroom to abide by playbook, including social media guidelines**
- 2. Reporter participates in social media consistent with guidelines, including when doing so in their personal capacity**

#### Informative reference

- DiResta & Grossman (2019).

### PHASE II

#### Reporter assigned story

##### **1. Reporter authenticates relevant content**

- a. Is the content complete? Or are pieces missing, missing which, if included, would alter meaning?

#### Informative references

- “Using Attribution as Cover,” in Chapter 3: The Sound of Silence: Strategic Amplification (Benkelman, 2019).
- “Learn basic skills to identify manipulations,” in Chapter 2: Sea of Falsehoods (Benkelman, 2019).
- “10 tips for verifying viral social media videos” (Funke, 2018).
- “First Draft Toolkit” (First Draft, 2018).
- “Verifying Online Information” (Urbani, 2019).

##### **2. Reporter establishes provenance**

- a. Why was the content previously non-public?
- b. Who is responsible for the content entering the public domain?
- c. Who is distributing and/or amplifying content on social media?
- e. What are their motives for making the content public and/or amplifying it?

#### Informative references

- “Digital Voids” (Boyd & Golebiewski, 2018)
- Chapter 3: The Sound of Silence: Strategic Amplification (Benkelman, 2019)
- “Field Guide to Fake News” (Bounegru et al, 2017)
- “Source Hacking” (Donovan & Friedberg, 2019)



### 3. Reporter writes story in accordance with standards

- a. Story contextualizes the content with appropriate coverage of provenance.

#### Informative references

- “The Truth Sandwich,” in Chapter 3: The Sound of Silence: Strategic Amplification (Benkelman, 2019).
- "Responsible Reporting in an Age of Information Disorder" (Kwan, 2019).

### Story written, goes to editor

1. Editor verifies consistency with playbook
2. Events depicted in story are still newsworthy
3. Headline signals why original actor released information
4. Headline does not repeat propaganda without context

#### Informative references

- “The Cover/No Cover Tension: Vaccine Hesitancy,” in Chapter 3: The Sound of Silence: Strategic Amplification (Benkelman, 2019).
- Chapter 5: Responsible Headlines, in “Responsible Reporting in an Age of Information Disorder (Kwan, 2019).

### Edited story goes to digital desk

1. Editor and reporter review online version for consistency with playbook
2. Digital headline preserves signal on why original actor released information
3. No hyperlinks to disinformation. Alternatives: screenshots, text text descriptions, no-follow links
4. Embedded photos, videos, documents are labeled synthetic or fake, as appropriate
5. Layout preserves context from getting lost or buried

### Story is published and promoted

1. Audience team follows playbook

## Bibliography

- Benkelman, S. (2019). *Getting it Right: Strategies for truth-telling in a time of misinformation and polarization*. Retrieved from <https://www.americanpressinstitute.org/publications/reports/strategy-studies/truth-telling-in-a-time-of-misinformation-and-polarization/single-page/>
- Bounegru, L., Gray, J., Venturini, T., & Mauri, M. (2017). *A Field Guide to “Fake News” and Other Information Disorders*. Amsterdam: Public Data Lab.
- Boyd, D., & Golebiewski, Mi. (2018). *Data Voids: Where Missing Data Can Easily Be Exploited*. Retrieved from [https://datasociety.net/wp-content/uploads/2018/05/Data\\_Society\\_Data\\_Voids\\_Final\\_3.pdf](https://datasociety.net/wp-content/uploads/2018/05/Data_Society_Data_Voids_Final_3.pdf)
- DiResta, R., & Grossman, S. (2019). *Potemkin Pages & Personas: Assessing GRU Online Operations, 2014-2019*. Retrieved from <https://fsi-live.s3.us-west-1.amazonaws.com/s3fs-public/otemkin-pages-personas-sio-wp.pdf>
- Donovan, J., & Friedberg, B. (2019). *Source Hacking: Media Manipulation in Practice*. MA. Retrieved from <https://datasociety.net/research/>
- First Draft. (2018). *First Draft Toolkit*. Retrieved from <https://start.me/p/vjv80b/first-draft-basic-toolkit>
- Funke, D. (2018). *10 tips for verifying viral social media videos*. Retrieved from <https://www.poynter.org/fact-checking/2018/10-tips-for-verifying-viral-social-media-videos/>
- Kwan, V. (2019). *Responsible Reporting in an Age of Information Disorder*.
- Urbani, S. (2019). *Verifying Online Information*.

## Acknowledgements

The authors thank their colleagues in the Stanford Information Warfare Working Group for catalyzing this collaboration and inspiring many of the ideas presented. Janine road-tested some ideas at a closed session hosted by the Federal Election Commission and the Stanford Global Digital Policy Incubator in September 2019. Danielle Jablanski provided valued research and programming support. The authors are grateful for feedback received from technology researchers and seasoned reporters including R.B. Brenner, Philip Taubman, Mor Naaman, Aviv Ovadya and Rosanna Guadagno that we convened at Stanford in December. And the authors would especially like to thank the senior editors and reporters who took the time to meet with them in New York and Washington DC and helped inform revisions to this document including John Daniszewski and Karen Mahabir at the Associated Press; Rebecca Blumenstein, Philip Corbett and Erin McCann at the New York Times; Marian Porges and Christopher Scholl at NBC News; Tracy Grant, David Cho and Glenn Kessler at the Washington Post; and Terence Samuel and roughly 20 other editors and reporters at National Public Radio.