

Addressing false claims and misperceptions of the UW Center for an Informed Public's research

Mar 16, 2023



The researchers at the University of Washington's [Center for an Informed Public](#) (CIP) are recognized leaders in the study of rumors, conspiracy theories, and mis- and disinformation. Over the past decade, our research has made significant strides towards understanding and addressing these problems. Unfortunately, some of the projects CIP researchers have contributed to have become the subject of false claims and criticism that mischaracterizes our work, a tactic that peer researchers in this space are also experiencing. As mis- and disinformation researchers, it's distressing — though perhaps not surprising — to see some of the very dynamics and tactics we study being used to disrupt and undermine our own work and its impact. That includes our work with the nonpartisan [Election Integrity Partnership](#) research collaboration that we helped launch in 2020 with the [Stanford Internet Observatory](#) and other partners.

We're incredibly proud of our work. We appreciate the University of Washington's support as our team faces these [false claims](#) and [conspiracy theories](#).

The criticism of the CIP's research and team members is part of a larger effort that seeks to undermine work to understand and address online

misinformation, disinformation and other forms of strategic manipulation. This effort aims to equate work to understand and address these challenges with “censorship” — functioning to cast doubt on research investigating mis- and disinformation and to undermine interventions that attempt to create more trustworthy information spaces. The rhetoric is similar to that employed in support of attempts to reframe the events of January 6, 2021, and to counter [the findings of the U.S. House’s select committee](#) that investigated what led to the violent attack that day on the U.S. Capitol.

One of the challenges of addressing misinformation is that corrections can often do more harm than good by bringing additional attention to false claims, and exposing new audiences. Indeed, it is [well established](#) that just hearing a false claim repeated, even within a correction, can make it more familiar, more memorable, and ultimately more believable for some audiences. Our team is very aware of the risks of giving oxygen to false claims. At the same time, we recognize the need to provide factual information that refutes some of the worst falsehoods and contextual information about how our work has been profoundly mischaracterized.

Our research team has been studying online rumors and conspiracy theories for a decade. One thing we have learned is that some of the most effective false narratives work not by spreading outright falsehoods, but by selectively seizing upon and mis-contextualizing bits of factual information, layering those with exaggerations and distortions to create a false impression. Unfortunately, these false impressions aren’t easily refuted through facts that counter individual claims. Often, those rebuttals just provide more ammunition for additional misrepresentations. So we thought we might take a slightly different tact and engage at the level of the false impression to explain how misperceptions of our work are being weaponized to fit into established political narratives.

False Impressions of the Election Integrity Partnership

Many of the misleading narratives and consequent misperceptions focus on our work with the [Election Integrity Partnership](#) (EIP). In the summer of 2020, researchers from the Stanford Internet Observatory (SIO), the University of Washington’s Center for an Informed Public (CIP), Graphika and the Digital Forensics Research Lab (DFRLab) embarked upon a collaborative “rapid-response” effort, which

would become known as the Election Integrity Partnership, to surface, analyze, and communicate about rumors and misleading claims about election processes and procedures.

False impression: *CIP researchers were acting outside of the mission of the university.*

In mid-July of 2020, researchers at Stanford pitched the EIP to our team at the University of Washington's Center for an Informed Public as a collaboration between four research organizations who would share resources and expertise to help identify and address the spread of harmful false and misleading information that might threaten the integrity of the U.S. election. As university researchers, we are encouraged both to contribute to scientific knowledge and to have "broader impact" on society. For example, CIP co-founder Kate Starbird participated, both as a PhD student and UW faculty member, in real-time "crisis mapping" projects that used crowdsourcing techniques to support disaster response after the Haiti earthquake, during the Deepwater Horizon oil spill, and around dozens of other events. The UW CIP team has specialized skills in social media data analysis and in July 2020, we accepted the invitation to join the EIP, offering to provide data analysis and communication support to the project. **Our participation in the EIP directly aligns with the CIP's public service mission and the UW's commitment to public scholarship.**

False impression: *The EIP is a partisan political project.* This incorrect impression stems from an attempt to frame the empirical findings of the EIP — i.e., that misleading claims about election processes and results spread more among Republicans than Democrats — as reflecting political motives of the partnership's mission. When we agreed to join, we ensured the effort was explicitly non-partisan.

Plans included collaboration with an office within the Trump Administration (the Department of Homeland Security's Cybersecurity and Infrastructure Security Agency) and outreach to both major political parties, who were invited to contribute to a crowdsourced tip-line. **The founding mission of the EIP focused on protecting the integrity of elections, not on supporting any specific political outcome.** The EIP's work sought to mitigate the impact of false claims that might interfere with election processes as well as false claims about election interference. After the events of the 2016 election, we believed this mission to be a nonpartisan one.

False impression: *The EIP was a "secret" project.* This false impression first emerged in summer 2022 after online activists purported to "discover" the EIP's work — in a

[peer-reviewed paper](#). This delayed discovery was not due to any nefarious effort by the EIP to obscure our work. On the contrary, we made every effort to share the products of our work and to accurately describe the processes underneath. In the weeks leading up to and following the November 2020 election, we published [a number of blog posts, graphics and data visualizations](#) showing how certain misleading narratives spread online, hosted news briefings, and fielded numerous interview requests from news organizations. In March 2021, our team published "[The Long Fuse: Misinformation and the 2020 Election](#)," a nearly 300-page report that, in addition to describing our methods, documented the narratives and information dynamics of the "Big Lie." We are incredibly proud of this research, which is part of the historical record, cited in the [U.S. House select committee's final report](#) into what led to the January 6, 2021, attack on the U.S. Capitol. Our efforts from 2020 led to follow-up peer-reviewed research published in journals and conferences, including [Nature Human Behaviour](#). And the paper that was "discovered" in the summer of 2022 was published in an open journal and promoted through our CIP website and social media accounts. **The EIP's work was conducted openly and transparently.**

False impression: *The EIP was a "censorship" operation.* At the core of most of the false impressions of our work is a rhetorical argument that seeks to equate efforts to understand and counter false and misleading information with "censorship." This argument has increasingly been employed against social media moderation efforts — as though these companies do not routinely act to limit spam, pornography, harassment, impersonation, and other harmful content on their networks. In 2020, some social media platforms put into place "civic integrity" policies to mitigate the spread of false claims about the 2020 election, including content that could disenfranchise voters by confusing them about when or where to vote and content that delegitimized election results. One dimension of the EIP's work was to alert social media platforms to misleading claims about election processes, discovered in the course of our analysis efforts, that may have violated their policies. Our understanding is that the social media platforms the EIP worked with provide similar reporting mechanisms for other researchers and organizations, in part because they do not currently have the internal capacity or expertise to do that work alone. Platforms also provide reporting mechanisms for all users to utilize should they encounter content that goes against community guidelines. The EIP's reports to the platforms were voluntary contributions to these companies efforts to mitigate election misinformation, and the platforms were not bound to act on the

recommendations of our researchers. **We disagree with the framing of the EIP's work as "censorship" — and are troubled by broader efforts to equate research about misinformation and disinformation with "censorship."**

False impression: *The EIP orchestrated a massive "censorship" effort.* In a recent tweet thread, Matt Taibbi, one of the authors of the "Twitter Files" claimed: "According to the EIP's own data, it succeeded in getting nearly 22 million tweets labeled in the runup to the 2020 vote." That's a lot of labeled tweets! **It's also not even remotely true.** Taibbi seems to be conflating our team's post-hoc research mapping tweets to misleading claims about election processes and procedures with the EIP's real-time efforts to alert platforms to misleading posts that violated their policies. The EIP's research team consisted mainly of non-expert students conducting manual work without the assistance of advanced AI technology. **The actual scale of the EIP's real-time efforts to alert platforms was about 0.01% of the alleged size.**

False impression: *The EIP's purpose was to route moderation requests from outside organizations to social media platforms.* This misimpression relies on three distortions of our reported work. First, though the EIP reported content to platforms, **alerting platforms to content that violated their policies was only a small part of the EIP's mission** — and not equally shared across the four collaborating teams. Other activities included publicly communicating in "real time" about misleading claims and narratives through tweets and blog posts, documenting the wide range of misleading claims and narratives about the election in our final report, and publishing a dataset mapping tweets to hundreds of distinct claims. Second, though the EIP included a "crowdsourced" tip-line where external partners could share pieces of content for us to consider for review, our researchers made independent decisions about what to pass on to platforms, just as the platforms made their own decisions about what to do with our tips. Third, the [majority](#) of our work focused on content surfaced by our own, internal research team. **The EIP's purpose was to support U.S. democracy through independently organized efforts to identify, analyze, document, communicate about, and correct false rumors and disinformation about election processes and procedures.**

False impression: *The EIP operated as a government cut-out, funneling censorship requests from federal agencies to platforms.* This impression is built around falsely framing the following facts: the founders of the EIP consulted with the Department

of Homeland Security's Cybersecurity and Infrastructure Security Agency (CISA) office prior to our launch, CISA was a "partner" of the EIP, and the EIP alerted social media platforms to content EIP researchers analyzed and found to be in violation of the platforms' stated policies. These are all true claims — and in fact, we reported them ourselves in the EIP's March 2021 final report. But the false impression relies on the omission of other key facts. **CISA did not found, fund, or otherwise control the EIP. CISA did not send content to the EIP to analyze, and the EIP did not flag content to social media platforms on behalf of CISA.**

False impression: *The EIP collaborated with and worked to support the Biden Administration.* This impression builds upon documentation of the EIP's partnership with CISA (an office that sits within the Executive Branch of government) and is used to promote a "weaponization of government" narrative aimed at the Biden Administration. However, this is easily corrected by a glimpse at the timeline. The EIP was founded in 2020 and its collaboration with the CISA office took place between July 2020 and November 2020. During that time, CISA was run by an appointee of President Trump. CISA's association with the EIP was reviewed and approved by Trump Administration attorneys as compatible with CISA's congressionally approved authorities. **When the EIP collaborated with an organization within the Executive Branch of the U.S. government (CISA), it was during the Trump Administration.**

False impression: *Researchers at the University of Washington were paid and directed by the U.S. government in their work with the EIP.* This misimpression builds from factual, public information about UW researchers' federal grant funding, but integrates a mischaracterization of this funding being designated for EIP operations and conflates those operations with "censorship." It then expands to include a false allegation that the agencies within the U.S. government and specifically the National Science Foundation (NSF) intentionally funded the University of Washington to "censor" specific voices. In 2020, UW participation in the EIP was predominantly funded by foundational and philanthropic funding. UW personnel funded by Kate Starbird's NSF CAREER grant did participate in post-election period analysis of EIP data for the partnership's final report and for subsequent peer-reviewed publications — and that grant is publicly acknowledged in that work. However, **research grants from the U.S. government did not significantly fund nor did**

U.S. government funding agencies direct or encourage participation by UW students, staff, or faculty in the platform-alerting functions of the EIP.

False impression: *The EIP purposefully targeted conservative political speech.* This false impression is created by underemphasizing the narrow scope of our research and highlighting specific elements of our empirical research findings (that more misinformation spread on the political right) without context (that these findings are unsurprising and align with [other research](#)). The EIP's work was narrowly focused on content that 1) interfered with voting by misleading about when or where to vote; 2) encouraged others to commit fraud; 3) used intimidation or threats of violence to deter voting; or 4) delegitimized election results through the spread of false, misleading, or unsubstantiated claims. In the lead up to the 2020 election, the EIP reported on misleading claims spreading through left-leaning audiences as well as right-leaning ones (e.g., [here](#) and [here](#)). After the election, the vast majority of false claims about the election emerged and spread among supporters of President Trump (a fact underscored by the January 6 violence at the U.S. Capitol), which is reflected in our data and reporting. **The EIP exclusively tracked and reported on false, misleading, and unsubstantiated claims about election processes and procedures. In 2020, those claims were far more prominent among supporters of President Trump (and the president himself) than other political groups.**

False impression: *The EIP orchestrated content moderation decisions by social media platforms around the story of Hunter Biden's laptop.* This false impression has few facts or even details behind it, but takes shape through repeated speculation and insinuation. **The story of Hunter Biden's laptop was out of scope for the EIP's work and the EIP did not play any role in: 1) decisions by Twitter (or any other platform) to limit spread of the laptop story; or 2) attributions of the laptop story to foreign influence operations.**

False Impressions of Dr. Kate Starbird's work on CISA's external advisory committee

In December 2021, CIP co-founder and faculty director Kate Starbird, a UW Human Centered Design & Engineering associate professor, was asked by CISA Director Jen Easterly to serve on CISA's external advisory committee (CSAC) — and to chair the MDM subcommittee. "MDM" is an acronym used by the U.S. government to refer to misinformation, disinformation, and malinformation. Starbird agreed to chair the MDM committee, which released [a first set of recommendations](#) in June 2022, and a

[second set of recommendations](#) in September 2022, concluding the committee's work.

False impression: *Members of CISA's MDM advisory subcommittee worked as part of a "censorship regime."* This false impression combines the argument that "moderation equals censorship" with false speculation about the nature of the MDM subcommittee's work. **The MDM subcommittee did not participate in or recommend for others to participate in any activities related to social media platform moderation or other activities that could be construed, even broadly, as "censorship."** The subcommittee was initially tasked with addressing — through written recommendations — challenging questions about how CISA should structure their work, engage with external stakeholders, and address privacy concerns related to addressing mis- and disinformation. The subcommittee limited the scope of their work to the context of elections. The subcommittee's recommendations focused not on how government or platforms should limit communication, but on how the CISA office should use their own communication, for example through public awareness campaigns, debunking falsehoods, and helping to amplify factual information from local and state election officials. The subcommittee did not recommend or discuss what actions social media platforms should take pertaining to specific content or types of content. Subcommittee members Kate Starbird and Vijaya Gadde did not discuss activities that platforms have taken or should take regarding specific content or policies more generally — neither within their roles on the committee or outside them. **The CSAC MDM subcommittee did not discuss whether or how social media platforms should moderate content, either in specific cases or more generally.**

False impression: *CISA's MDM subcommittee recommended that the federal government participate in monitoring of social media platforms and other information spaces.* This misimpression emerged following an unfortunate November 2, 2022, article in The Intercept that included a misleading edit and broader mischaracterization of the subcommittee and its work. The article stated that the subcommittee had recommended CISA to "*closely monitor* 'social media platforms of all sizes, mainstream media, cable news, hyper partisan media, talk radio and other online resources.'" To be clear, **the MDM subcommittee explicitly never advocated for CISA to monitor or "closely monitor" anything.** During internal discussions, the members noted that questions about social media monitoring by CISA and other government offices were beyond the capacity of the subcommittee

and, though initially tasked with providing recommendations on that aspect of CISA's work, the subcommittee intentionally did not provide recommendations about this. The Intercept added "closely monitor" to a section of the report that was instead encouraging CISA to consider the challenge of MDM as broader than just social media. **This misleading misquote of the report has contributed to a lasting, widespread, and harmful misperception about the MDM subcommittee's work.**

What comes next and what's at stake?

At the University of Washington's Center for an Informed Public, our research team has developed unique expertise, including methods and infrastructure, for rapid analysis of social media information flows during fast-paced and high-visibility events. Our researchers have a long track record of studying rumors, conspiracy theories and mis- and disinformation that not only pre-dates the formation of the Election Integrity Partnership in 2020, but also our own research center, which was established at the University of Washington in 2019. This vital work will continue.

As [multiple public opinion polls](#) show, Americans are very concerned about mis- and disinformation, which can be harmful in certain contexts and lead to poor decision making during crises and emergencies. Disinformation can manipulate individuals and societies in harmful ways. Pervasive disinformation can undermine our trust in information, in science, in our foundational institutions, and in each other. Online mis- and disinformation have real-life consequences, as we saw on January 6, 2021.

Our work at the CIP, including our Election Integrity Partnership collaboration in 2020, has been transparent, research-driven, and rooted in support for democracy and support for a more informed public. This is the CIP's mission. This work will continue. We're currently working on a final report around our research on the 2022 U.S. midterm elections, expected for release in the coming months, but you can explore [blog posts and other analysis we published](#) and [shared via Twitter](#) last fall.

As researchers who study the dynamics of rumors, conspiracy theories, and mis- and disinformation online, we're well familiar with Brandolini's Law, or the [bullshit asymmetry principle](#) where the "amount of energy needed to refute bullshit is an order of magnitude bigger than to produce it." As we've been responding to this

slurry of false claims, distortions and misunderstandings, we've learned that all the attention on our research and researchers underscores their importance and impact. We will continue to stand behind and defend our work.

Kate Starbird

Associate Professor, University of Washington Human Centered Design & Engineering
Co-Founder and Director, UW Center for an Informed Public

Ryan Calo

Professor, UW School of Law and UW Information School
Co-Founder, UW Center for an Informed Public

Chris Coward

Senior Principal Research Scientist, UW Information School
Co-Founder, UW Center for an Informed Public

Emma S. Spiro

Associate Professor, UW Information School
Co-Founder, UW Center for an Informed Public

Jevin D. West

Associate Professor, UW Information School
Co-Founder, UW Center for an Informed Public

Other News



A chatbot exercise in 'BSing the BS principle'

Apr 4, 2023

In a March 31 opinion in The Seattle Times about AI chatbots, University of Washington Center for an Informed Public co-founder Jevin West says that chatbots will be “vectors of propaganda,” make it harder to discern truth and further erode trust in institutions.

[Read More](#)

Responding to recent questions about Kate Starbird's participation on a CISA external advisory committee

Apr 4, 2023

Over the past several months, our team at the University of Washington Center for an Informed Public has been responding to inquiries related to our research and UW associate professor Kate Starbird's participation on an external advisory committee for the Cybersecurity and Infrastructure Security Agency (CISA)

[Read More](#)

At MisinfoDay 2023 events across Washington, high school students and educators learn valuable skills

Mar 28, 2023

Approximately 700 Washington high school students, teachers, librarians and other educators participated in MisinfoDay 2023 programs in March across three in-person events at the University of Washington in Seattle and Washington State University in Pullman and Vancouver.

[Read More](#)[Join our Mailing List](#)[Contact Us](#)[Privacy Policy](#)[Terms and Conditions](#)

© 2021 Center for an Informed Public at UW