

You Are Here

A Field Guide for Navigating Polarized Speech, Conspiracy Theories, and Our Polluted Media Landscape

Whitney Phillips and Ryan M. Milner

The MIT Press
Cambridge, Massachusetts
London, England

© 2021 Massachusetts Institute of Technology

All rights reserved. No part of this book may be reproduced in any form by any electronic or mechanical means (including photocopying, recording, or information storage and retrieval) without permission in writing from the publisher.

Library of Congress Cataloging-in-Publication Data

Names: Phillips, Whitney, 1983– author. | Milner, Ryan M., author.

Title: You are here : a field guide for navigating polarized speech, conspiracy theories, and our polluted media landscape / Whitney Phillips and Ryan M. Milner.

Description: Cambridge, Massachusetts : The MIT Press, 2021. | Includes bibliographical references and index.

Identifiers: LCCN 2020007621 | ISBN 9780262539913 (paperback)

Subjects: LCSH: Internet—Moral and ethical aspects. | Social media—Moral and ethical aspects. | Fake news. | Disinformation. | Propaganda. | Media literacy.

Classification: LCC TK5105.878 .P45 2021 | DDC 302.23/1—dc23

LC record available at <https://lcn.loc.gov/2020007621>

10 9 8 7 6 5 4 3 2 1

d_r0

6 Choose Your Own Ethics Adventure

Our network climate is in crisis. It's inundated with pollution not just old, and not just new, but new as a catalyst to old, and old as a catalyst to new. It pulses without end underfoot. It blights the land we harvest. It swells and swirls in the clouds above. And it overlaps. Storm runoff filters into forest beds. Cyclones suck up slowly pooled poisons and dump them far afield. Fruits emerge from toxic soil. Just recognizing a pollutant's source can be vexing enough. Trying to quarantine that source when it's coming in from everywhere, when half the time it's invisible, when efforts to help so easily backfire, is Sisyphean. As the filth piles up and we gasp for air, the question looms: why even try?

The answer is simple: because there's hope. By embracing a communitarian approach to information—achieved through ecological literacy and network ethics—citizens of good faith can help clean up the pollution already present, minimize the new pollution produced, and, most profoundly, cultivate a different way of being in the world—one that can, if enough seeds are planted by enough people, grow into industry, education, and government. In some networks, those seeds have already begun sprouting. As the COVID-19 pandemic exploded, journalist April Glaser chronicled how everyday citizens created mutual aid groups by using text threads, Google Docs, and Facebook pages to assist community members in need.¹ In the most dire moments, these neighbors recognized the value of communitarian care: that when the people around you are healthier and safer, everyone is healthier and safer. The energies we need to cultivate already exist; they just need nurturing. That's our collective job moving forward.

The work ahead isn't restricted to any one profession or political affiliation; nor is it dictated by the size of our platforms. The information ecosystem is such that whatever we do for a living, whatever we believe, however many people might be listening to us, what we do and say can have sweeping effects. We might not all be

journalists, but we all sow content across and between networks. We might not all be influencers, but we all hold sway over someone. We might not all be educators, but we all can teach by example. We might not all be students, but we all have something to learn. We all have a part to play in what the world is like. It's time to start working together to ensure the healthiest possible future for all. Here, we offer some first steps on that journey.

One: Pull Out the Network Ecology Map

The first task is to triangulate our respective “you are here” stickers on the network map. Offline, we benefit greatly from knowing where we're standing. When we can assess what's happening around us, we're able to make the best possible choices about what we can or should do next—which turns to take more slowly or avoid completely, where we need to watch our footing, what we need to bring. The same holds true for our online roamings. When we know where we are in relation to everything else, we're in a better position to consider the consequences of what we do and say (or don't). Offline, we do this by looking up, looking down, and looking side to side. Online we can do the same thing, just with a more poetic eye: by charting the roots below, the land around, and the storms above.

Looking down at the roots beneath our feet helps us trace where polluted information came from and how it got there. What networks has the information traveled through, and what forces have pinged it from tree to tree, grove to grove? This pollution might be economically rooted, grown from profit-driven corporate institutions. It might be interpersonally rooted, grown from long-held assumptions about what's acceptable to share, or funny to share, or necessary to share. It might be ideologically rooted, grown from the deep memetic frames that support people's worldviews. The trick is to remember that polluted information never just appears. Pollution filters in, driven by countless catalysts. When trees start dying, when they fail to flower, when their leaves go brittle—that's the symptom. The cause

is how and why that pollution spreads. Neither symptom nor cause makes any sense without the other.

Looking around at our vast tracts of farmland helps us pinpoint the impact people have on their networks. Some of these impacts are the result of deliberate choices. Some are the result of simply existing. The trick is to remember that motives have very little to do with outcomes. Like Kansas farmers in the 1920s, people with the best intentions and the worst intentions can harm the land in equal measure. The environmental damage inflicted by the people setting out to sow chaos and confusion is obvious. Much less obvious, but just as sweeping, are the cumulative effects of all the pollution everyday people kick up without trying. Searching for the source of this pollution doesn't necessarily mean searching for villains. It means highlighting effects and identifying causes—most critically when that cause is you.

Looking up at the storms raging overhead helps us identify the holistic, evolving forces driving a story forward. Pizzagate, QAnon, and the Seth Rich assassination theory didn't spread solely because some people believed—or at least promoted—those particular theories. They spread because social media monetized, incentivized, and surfaced harmful content. They spread because journalists reported incessantly on far-right narratives during and after the 2016 election. They spread because audiences read and responded to what journalists published. They spread because these things, and so many others, influenced and were influenced by everything else. The trick is to remember, first, that no hurricane can be reduced to any of its parts, and second, that hurricanes are verbs. They absorb everything that caused them to form, and all the responses they generate. Each moment, each thing, becomes yet another source of energy. None of us exist outside the storms we track. So all of us must be careful about what we feed them.

Two: Take What You Don't Know as Seriously as What You Do

Triangulating your position on the network map is as much about identifying what you don't know as it is about establishing what you do. Critical yet frequently missing information can include where exactly a claim, attack, or campaign originated, or who exactly has joined in since. Additionally, why a particular poster or group is doing what they're doing is often difficult if not impossible to verify; this is the consequence of a media landscape governed by Poe's Law. Something might look a certain way online, but that doesn't mean you can know what's fueling a claim, attack, or campaign. It could be organic, driven by actual, embodied people sincerely participating. It could be artificial, driven by sock puppet manipulators, cynical propagandists, or automated bots. It could be the result of some dizzying combination of both. All that missing information leads to the most vexing unknown of all: who or what might benefit from a story.

Not knowing the answers to these questions means not knowing the underlying ethical stakes. Not knowing the underlying ethical stakes means not knowing the best course of action. Not knowing the best course of action almost guarantees that any amplification, of any kind, will send pollution flying in unpredictable directions. Assessing what we don't know is therefore a critical first step before publicizing anything.

Mainstream news coverage of the QAnon-infested Trump rally, discussed in chapter 4, exemplifies how faulty premises and limited observations can warp a story—and how warping a story can help supercharge that story as it roars up the coast. In this case, reporters made sweeping pronouncements about how widespread the QAnon conspiracy theory had become among Trump supporters, conclusions based on how many people showed up to the rally wearing Q shirts and holding Q signs. The result was a deluge of QAnon coverage. The small detail that the people brandishing QAnon paraphernalia had coordinated beforehand—indeed, that they had been coordinating for weeks at other rallies as well, with the specific intention of tricking reporters into declaring QAnon the zombie that ate the Trump campaign—wasn't included in the stories. This omission allowed

disinformation to masquerade as fact, much to the delight of the people seeding the pollution.

Of course, it's not always possible to know what isn't known, especially given that manipulation strategies—like those of QAnon proponents—are almost always developed outside the public spotlight. Not having access to those secluded corners, or even knowing for sure that there are secluded corners to look for, makes anticipating traps all the more difficult. That said, tracing the network map helps, at the least, to block out what areas *are* known. Based on that information, it's easier to identify what's missing.

Surveying the known landscape also helps pinpoint the areas on the map that are ripe for manipulation. The most common targets are unfolding news stories about human suffering, from acts of extreme violence like mass shootings to pandemics like the COVID-19 outbreak. In these cases especially, you can be certain that any media manipulation strategies you can identify, the manipulators have identified too. As you try and head the bad actors off at the pass, keep your eyes glued to the network map, and do what you can to learn what you can from what the map can't tell you.

Three: Remember That Affordances Have Consequences

The information ecosystem is not a natural environment. It was developed by people, none of whom stood outside their creations as objective, neutral observers.² They created certain things for certain purposes, encoded with certain assumptions about who their users would be and what those users would do. So while digital media allow a great deal of freedom, they're also limited by a curated set of tools for navigation and play. The affordances of these tools direct what we are able to do—indeed, what even occurs to us to try to do—as we traverse the landscape. Understanding the affordances that surround us online and what impact those affordances have on our behavior allows us to move with greater care and self-awareness. That digital media are often designed to *preclude* care and self-awareness only underscores the need for a manual override.

The most foundational affordances are the ones we employ without thinking. These include the ability to edit digital content, to use a segment of the content without destroying the larger whole, to store and index content, and to access content quickly and easily. Separately and together, these affordances work to sever individual texts—words, images, audio, GIFs, and video—from their broader personal, political, and historical contexts.³ The result, gesturing back to why we need to take what we don't know seriously, is to obscure crucial information like the text's origins, the circumstances of its collection, and the impact it had on the people who created it or were featured in it—essentially, everything that would help a person assess the consequences of sharing. With so much information missing, people tend not to make ethical choices. Not because they're *unethical*, but because they don't realize there's anything to be ethical about. This is a trap; there's always something to be ethical about.

Platform-specific affordances further shape our behavior online. The most basic functions of social platforms, namely, the ability to post, like, comment on, and share, and to see what others post, like, comment on, and share, might not seem worthy of much reflection. That's just social media. However, those basic platform affordances—coupled with how all that content is moderated, or not, and how algorithms promote it, or don't—lay an even more treacherous set of ethical traps.

Most pressingly, design choices predicated on the liberal assumption that information wants to be free prime the landscape for context collapse and its messy commingling of audiences. Different users might encounter the same words or image or audio or GIF or video, but what that content means for—or does to—each of those users can vary wildly. A single post can occupy the entire emotional gamut, from obvious joke to obvious attack to just about anything in between, all depending on whose eyes are seeing it. Context collapse further exacerbates and is exacerbated by Poe's Law. The original intended audience for any given post might—and this is a big *might*—know what something “really” means. Outside the intended audience, however, meaning is in the eye of the beholder; for every

deep memetic frame a person might be standing behind, and for every embodied reality that shapes a person's understanding of the world, there are different ways of seeing.

As the unfettered-speech, share-at-all-costs ethos baked into social media sends content screaming across different audiences with different assumptions, the very notion of *meaning* is jumbled beyond recognition. And when meaning is jumbled, so too is any clear sense of how best to respond to questionable content. Jokes are especially fraught online; they require a clearly carved-out play frame to function as a joke—one that signals to all involved, “Hey don’t take what I’m about to say seriously; I don’t really mean it.”⁴ That signal is the first thing chucked out the window by context collapse and Poe’s Law.

The effect doesn’t just ruin someone’s good joke. If it’s not clear whether a message—whatever form it might take—is meant to be a joke or a threat or an earnest observation, that message can be adopted, reframed, and potentially weaponized by others. These second- and third- (and forth-, and fifth-) order sharers can then import their own meanings to the message, spinning it off into increasingly far-flung corners of their networks. When that message contains even the slightest trace of falsehood or dehumanization, it’s almost guaranteed to leave a trail of pollution in its wake. Because new audiences don’t know to look for it, they are unlikely to notice this trail and therefore are unlikely to be wary of the message being shared. All they see is the content in front of them—not where it came from, not what it has done, and certainly not what it started out as.

The ethical pitfalls inherent to context collapse and Poe’s Law are epitomized by the “Donald Trump’s Dank Meme Stash” Facebook group discussed in chapter 3. The young reporter who joined the group for laughs assumed that the racist and sexist anti-Clinton, pro-Trump memes posted there were satirical, just another bit of lulzy fun. What the reporter didn’t know, what Facebook’s platform affordances obscured, was where the memes came from, who made them, and why. Had the reporter known that many of those memes

were created and spread by white supremacists, conspiracy entrepreneurs, and hostile state actors, he would have been less inclined to laugh. He would certainly have been less inclined to pass them on.

This young reporter, like so many of the other young reporters who covered the internet culture beat during the 2016 election, like so many reading now who have likely had similar experiences, didn't share or laugh at harmful and dehumanizing content with the specific intent to harm and dehumanize. Some people did, of course. But most people simply didn't see what they were doing, because the platforms they were navigating didn't allow them to—or want them to. The result was a windfall for the social media companies who hosted and monetized all that pollution.

The underlying problem is that there was not then, and is not now, much incentive for these companies to foreground network ethics or build protective guardrails into site design—because ethics and guardrails aren't good for business.⁵ What's good for business is people not thinking very hard about the content they share, and sharing as much of that content as possible. As users of these sites, and conversely, as people being used by these sites, we must understand that we don't own the land we traverse. Just being there comes at a very high, if largely unseen, price. Namely, we have been set up to pollute, because corporations and their shareholders benefit when we do. We must therefore work that much harder to navigate our landscapes carefully, to not allow the tools we use to restrict our sight, and, when we encounter particular texts, to actively, doggedly, seek out context. The alternative is to stand there, smiling and oblivious, as we kick over our own personal holding tank of accidental toxins.

Four: Adjust Your Understanding of Harm

Typically, conversations about harm online focus on people who relentlessly taunt, harass, and dehumanize others. These are the lions and tigers and bears at the top of the biomass pyramid: smaller in

relative weight and number than other animals in the ecosystem, but outsized in the dangers they pose. Assessing harm based on their predatory intent makes perfect sense. Logistically, it's difficult to imagine how ongoing hatred and harassment would even be possible unless someone was actively trying to hate and harass.

What makes perfect sense for apex predators, however, makes increasingly less sense as you move down the pyramid. In fact, positing intent *to* harm as a fundamental criterion *of* harm can inadvertently encourage damaging behavior by stifling ethical self-reflection. If harm is defined as something an abuser or manipulator or bigot does on purpose, and people don't consider themselves abusers or manipulators or bigots and certainly aren't *trying* to harm anyone, then they're almost guaranteed to give their behavior a pass. Harm is something that bad people choose to do. I'm not a bad person, and I'm not making that choice. So everything's good.

Of course, just because people believe they're in the ethical clear doesn't mean they actually are. They can spend all day justifying what they did. They were just trolling. They were just joking. They were just playing devil's advocate, no offense, you're being so sensitive. Fine; that's what they believe. But the second those actions harm another person, the wide-eyed insistence that, but-but-but, I didn't *mean* anything bad by it, doesn't change, can't change, and shouldn't be expected to change the impact of what that person said or did. If that impact was harmful, then by definition, they caused harm.

Patrick Davison, reflecting on the legacy of the MemeFactory performance trio discussed in chapter 2, epitomizes this tension. Looking back, Davison explains, he has regrets about the fetishized memes and jokes MemeFactory set their audiences up to laugh at—memes and jokes Davison also helped publicize through other venues like Know Your Meme. But, he says, "It's not that I look back and regret having attempted to cause harm; I look back and regret that my attempts to contribute (whatever it was) likely ended up contributing to lots of different kinds of harm."⁶ You can't regret making a choice you never actually made, and Davison didn't choose

to hurt anyone. Still, the jokey antagonisms MemeFactory helped normalize, the jokey antagonisms internet culture helped normalize, the jokey antagonisms the two of us as young researchers helped normalize, together shifted the media ecosystem. All this seemingly harmless play allowed, even encouraged, bigotry and falsehood to climb into the sheepskin of trolling.⁷ The result was to flatten unwilling people into fetishized memes, spread lies far and wide, and reinforce lulz as an aspirational register. It may have been unintended, but the damage was still sweeping. It still herded prey of all shapes and sizes into the mouths of all those lions and tigers and bears.

These are the harms we miss when we train our eyes on the top of the biomass pyramid. To reduce the pollution we spread, to reduce the pollution we create, we must lower our eyes to the bottom of the pyramid, to the rabbits and worms and fungi. In other words, to *us*. There the harms aren't just larger in number and weight. There the harms have something the top layers will never have: our ability to intervene, right now, with our own fingertips.

Five: Spread Information Strategically

As we think carefully about whether and how to respond to information, particularly when that information is already polluted or has the potential to mutate into pollution once it leaves our feeds, the call is not for people to keep their mouths shut. It certainly isn't for people to "remain civil," an all-too-frequent gaslighting tactic deployed by those whose behaviors are blaringly immoral but who don't like being criticized for it. Indeed, as the journalists interviewed in chapter 3 explained, there can be as many arguments *for* responding to polluted information as there are for *not* responding. Not responding prevents people from telling the truth, educating those around them, showing the embodied effects of online harms, and pushing back against bigots, manipulators, and chaos agents. Not responding can signal complicity. Not responding doesn't make the problem go away.

Given these risks, silence isn't always possible. Silence isn't always advisable. The challenge is to be strategic about the messages we amplify.⁸ More than that, the challenge is to approach amplification with ecological literacy. The question isn't just "to amplify or not to amplify?" The question—to be asked anew case after case, click after click—is: What are the environmental impacts of my choices? What pollution might my actions generate, for whom? What pollution might my actions mitigate, for whom? Most important, whose bodies might my actions nourish, at the expense of whom and what else? As ethical precepts, the following guidelines are far from a list of "thou shalt nots." Not every case poses the same risks. Not every case will have a clean outcome, particularly when the information storms loom especially large. Sometimes the best we can strive for is *less* pollution, not *no* pollution. These guidelines reflect that underlying ambivalence. So, rather than offering false comfort ("when *x* happens, just do *y*"), they emphasize the deep reciprocities between audiences and institutions and affordances. To the extent that there are solutions, they live within those connections.

Give Yourself Some Credit

First and foremost, whatever our professions or ideologies, we are all part of the amplification chain. This is obvious for people with large platforms, who are well aware, even painfully aware, that their words can travel across the globe in an instant. Everyday folks, on the other hand, with everyday follower counts, can feel like their words don't matter. That simply isn't true; everyday folks have enormous power. Everyday folks might get handed a menu of options by algorithms, but that's because algorithms are trained by the patterned actions of social media users. Everyday folks might get media narratives shoved in their faces, but that's because journalists are trained by the habits of their readers. Everyday folks might get products and ideas peddled to them by influencers, but what those influencers peddle is shaped by the support, needs, and preferences of their audiences. Recognizing that what we do has consequences across networks and even entire industries is the first step in making more strategic

choices—because it reminds us that people are watching, often in ways we can't confirm or predict.

Everyday people play an especially crucial role within the biomass pyramid. In the natural world, the lions and tigers and bears wouldn't last a season without the bunnies and worms and fungi. Likewise, the apex predators of the online world, whose actions are most toxic, are able to do what they do thanks to the energies provided by everybody else. From our retweets to our comments to our laughter, what the rest of us do helps keep bad actors well fed. Understanding just how much they depend on our resources is the first step toward minimizing their harms. They need us, which is one power we have that they don't.

This is not a set up to the command, “don't feed the trolls,” often presented as the only viable defense against the dark arts online. “Don't feed the trolls” implies that if you, as an individual, feed the trolls, then what happens to you, as an individual, is your fault. For one thing, that's victim blaming. The perpetrator is the harmful one and shouldn't be doing what they're doing (so cut it out, you assholes). For another thing, framing hate or harassment or dehumanization as an isolated problem obscures the structures that support and even incentivize those actions. Simultaneously, framing the *target* of hate or harassment or dehumanization as isolated, as a singular receptacle for harm, obscures the interconnections between biomass strata—and, most important, the collective power of people at the bottom of the pyramid to reshape their networks.

In other words, the call—unlike the highly individualistic “don't feed the trolls”—isn't to more effectively fend for ourselves when someone harms us. The call is to protect the people around us from harm and, more broadly, to make the landscape less amenable to those who are harmful. By making life harder for the apex predators, we ensure, together, that more people can be more free, more safe, and more empowered within our shared spaces.

Understand the PR Goals of Bigots, Abusers, and Manipulators, and Do Not Help Them Do Anything

Asking people to protect one another from apex predators is one thing. Explaining exactly how to do that is another.

Those who work for social media platforms have a unique opportunity, and from a communitarian perspective, a unique responsibility to affect these kinds of changes. Efforts to cultivate healthy community norms are especially crucial to ensuring the safety of users and signaling to abusers, “you are not welcome here.” As computer scientist J. Nathan Matias shows, steps as straightforward as clearly and consistently reinforcing community guidelines can reduce harmful behaviors and allow more people to participate more meaningfully online.⁹

But of course, relatively few of us work for social media platforms. For everyone else—even those with the smallest social media following—one of the most effective strategies is to reconsider where we’re pointing our cameras (metaphorically and sometimes literally). In particular, we should avoid focusing too intently on the motives of bigots, abusers, and manipulators, or too intently on the community they originated from, or too intently on anything that frames them as the protagonists of the story. The bad actors are part of the story, certainly, but they’re not the whole story.

The risk isn’t just that we’ll tell a worse and smaller story by highlighting bad actors. Treating them as inherently interesting and deserving of our undivided attention incentivizes future bad actions and, by extension, future harms to others. Because from the perspective of those bad actors, it makes sense to use the same weapon again; it worked so well the last time. Intense focus on what bad actors think and do and feel also risks replicating their marginalizations by implicitly saying, you know what, this terrible bigot or abuser or manipulator is right; it *is* their world, and we really are just living in it. You, targeted person—who are statistically more likely to be a woman or a person of color or both—are little more than a bit player in their drama. Please stand aside while we take their picture. Don’t give bad actors that satisfaction, and don’t provide them tools for future abuse. Most important, don’t be complicit in their harms. If a response seems warranted—and

sometimes it is—whip that camera around to the rest of the landscape, making sure to train your lens on the people targeted.

That said, don't merely to show the effects of those harms and then stop filming—particularly when the person holding the camera is white and the person they're documenting is not. Debra Walker King warns of the unintended consequences of reducing marginalized people to abstract representations of pain.¹⁰ A person who has been harmed is so much more than the violence they've been forced to endure. Narratives about their experiences should strive to foreground the agency, resilience, and unique subjective experiences that allow them to navigate a systemically hostile environment. That's where the counternarratives are, including the subtly powerful reminder that people who choose to set the world on fire, who don't care about the lives of others, who are violent and hateful and inhumane, are not the only people on the internet. Their targets also exist; the rest of us also exist. And not just exist, but are foundational to the ecosystem. How can we reframe the story so that the overwhelming number of citizens of good faith, and more pressingly, the overwhelming number of community organizers, activists, and heroes who selflessly work for the benefit of others, are the ones in the spotlight?

Consider What Light Might Do Besides Disinfect

When confronted by violence, bigotry, and lies, many people ascribe to the maxim that “light disinfects.” As explained in chapter 5, this assertion can be made while standing behind two different frames. The first is the light of liberalism, which maintains that we must show what the bad actors are doing so that rational observers can recognize, analyze, and reject their harms. The second is the light of social justice, which maintains that we must show what impact the bad actors have on their targets so that public opinion can shift and usher in needed social change.

This is an important distinction. *And*, the difference between the light of liberalism and the light of social justice online can crumble

under the weight of platform affordances designed to maximize sharing, scramble audiences, and generally affix question marks over what anything really means. As a result, both kinds of light can be quite volatile. The light of liberalism, which trains its spotlight on the worst actors, is most obviously so; for one thing, many of the people illuminated have perfected the art of weaponizing light. But the light of social justice can pose similar challenges. For instance, while social justice interventions might elevate a target's experiences for some audiences, other audiences might simultaneously flatten those experiences into a series of consumable agonies, even entertaining agonies. They might also single that target out for more and worse abuse.

Online, we can't escape this ambivalence. All the light shining on one side of the network might disinfect a particular toxin admirably—or at least serve as a beacon of solidarity for those who already know the toxin is hazardous. But in the corners infested with bigots and abusers and manipulators, where disinfectant is most desperately needed, the same light can grow something worse. It can serve as proof of concept for even more egregious behavior. It can unify otherwise disorganized groups. It can make bigots, abusers, and manipulators feel good about themselves and excited to see what else they can ruin.

That's not the only thing light can do. In addition to helping cultivate poisoned gardens, light can also fix things in place. Dentists use these kinds of lights—typically high-powered halogen or LED bulbs—to set composite fillings. In the context of polluted information, spotlights can serve a similar function. Whether we're shining the light of liberalism or the light of social justice, when we point at something online through our comments, retweets, or hot takes, we help amplify that thing. Maybe we have no other alternative; maybe the benefits of amplification outweigh the risks of silence. At least, maybe they do as far as we can tell in that moment. Regardless, our amplification makes that thing stable and searchable, particularly if other participants join in and pile on.

The bigger your audience, the brighter your light, the stronger your stabilizations. But as above, so below: even people who have very few followers can generate a chain reaction of illumination, especially when coupled with the power surge provided by hashtags, curated feeds, and other bits of algorithmic docenting. Unfortunately there are no simple, universal answers for questions about when to illuminate and when to leave something in the shadows. Whether we're the first or the thousandth person to encounter that thing, there are no guarantees about what our light might do. But there are questions we can ask ourselves as we weigh our options. What audiences might the post or comment or hot take reach, both intended and unintended? Is clarifying or condemning a point for one audience worth emboldening, validating, or otherwise delighting another? Whose interests might I ultimately be serving by what I share? There may not be perfect answers to these questions; there may be no way to avoid spreading some pollution. Still, the takeaway is simple: shine your high beams wisely.

Remember That Facts Are Not Cure-Alls

Like many assumptions about how best to respond to polluted information, the "light disinfects" maxim is premised on the old-growth liberal presumption that when people are exposed to lies and dehumanizing attacks, they will reject them. All we need to do is present those harms unvarnished, and critical thinking will do the rest.

Unfortunately, as case after case throughout the book has shown, facts and facts alone are not why people believe in and do things. It's not that we *never* believe things because of facts. That happens too. But very often, when we commit ourselves to something strongly enough to act with gusto, we have arrived at that point not because of facts but because the belief or action lines up with our deep memetic frames. If facts aren't how we got there, then facts won't change our minds.

So while the impulse might be noble, throwing a fistful of facts at someone who is wrong about an empirical truth is highly unlikely to solve the problem. The best-case scenario is that the person will reject our evidence out of hand, because from their vantage point, seeing through their frame, we are unintelligent, misinformed, or downright disturbed. The worst-case scenario is that our fact checks inadvertently reinforce their false beliefs, particularly when the person hearing the fact check already believes that we are biased, hostile, or a representative of a vilified *them*.

That doesn't mean we should look the other way when someone says that a pandemic is a hoax, any more than we should look the other way when confronted by violent bigotry. It means we should intentionally craft thoughtful responses that minimize unintended consequences while remaining conscientious, always, of the complexities of human psychology. The question is, how do we do that? If facts aren't enough, if facts might even make things worse, how can any of us hope to clean up all this pollution? How can we stop the pollution from flowing in the first place? There are no simple, one-size-fits-all answers to these questions either. But there is one basic place to start, and it returns us to the declaration: "you are here."

Six: Know Thy Politically Situated Selves

Determining where someone stands on the network map begins with assessing their deep memetic frames. This can be tricky, as deep memetic frames are often difficult to detect, especially for the people standing behind them. For these people—indeed, for everyone—the frames we see the world through aren't frames. They're just how the world is. In some ways, that's right: deep memetic frames are both constructed realities *and* how the world is, at least for the person whose frames they are. The frames might not be true, but they are certainly *real* to the person peering through them.

This is not to fall back on slippery, noncommittal, beige-tinged relativism, the moral equivalent of shrugging and saying, "people are

different” as someone commits a crime against humanity. Not all deep memetic frames are created equal. Some are explicitly damaging and false. The point of identifying a person’s frame is not to make excuses for the person or their frame. It’s to better contextualize what that person believes and how they came to believe it. It’s to get the lay of that person’s land.

Doing so helps us better target our responses to false beliefs, most critically when the beliefs pose a threat to public health. Stephan Lewandowsky’s research team, as well as studies by Brendan Nyhan and Jason Reifler, lay out how this can work.¹¹ As highlighted in chapter five, both sets of researchers emphasize how essential narrative coherence is to our worldviews. Fact checks that merely pull the narrative rug out from under someone are unlikely to be successful. And so, when a corrective is issued, it must not merely reject a certain detail. It must instead present an alternative coherent story that explains why something is the way it is.

The trick is figuring out what alternative coherent story would be needed. That requires assessing the stories that people are already telling themselves, which in turn requires assessing how those stories connect within the person’s deep memetic frames. These stories and their supporting frames reveal contextualizing information like who the *them* of that person’s conspiracy theories are, and where that person sees the most pressing dangers. Merely yelling at someone about how wrong they are—about COVID-19, about the Deep State, about climate change—*isn’t* going to tell the sort of alternative story that might, just might, get them to start seeing, or at least start being open to the possibility of seeing, things differently.

This approach has two immediate benefits. First, presenting a frame-appropriate alternate narrative minimizes the polluted splashback that our most well-intentioned fact checks, along with our less well-intended insults, can generate when they collide with someone else’s beliefs about the world. Instead, it zeros in on the coherency gaps that people don’t realize are there, which emerge from internal contradictions within their own stories or because some element of those stories doesn’t line up with the norms the person

otherwise accepts. In other words, these interventions are like a water gun pointed at a hole in a fence. As long as your aim is steady, the water goes one place and one place only. This is a much more exacting approach than broadly fact checking, which is like throwing a bucket of water against the fence; some of the water might go where you need it, but a lot will splash all over everything else.

Second, the very act of acknowledging someone else's frame and aligning the discussion with where that person is standing offers a basic affirmation of their identity, an approach that, as Lewandowsky and his team highlight, helps increase a person's receptivity to factual counterpoints.¹² Relatedly, starting with the logics and vocabulary of a frame, even when pushing back against it, guides the conversation toward meta-reflection.

The satanic conspiracy theorists discussed in chapter 1 show how this might work. These theorists alleged that the similarities between satanic abuse narratives proved the truth of those narratives and therefore justified efforts to out Satanists within local communities. How could so many people tell the same basic story if there weren't an extensive network of Satanists committing the same kinds of ritualized crimes in the same kinds of ways? An eye-rolled "Oh Jesus, there are no child-sacrificing Satanists" wouldn't just *not* convince a conspiracy theorist. It would likely serve as proof that they were onto something. In contrast, a discussion about the echo-systems that push information across networks—with each network seeming to corroborate the others but actually just reflecting back the same sources, details, and atrocity stories—would help guide believers to the gaps in their frame, because, in less loaded circumstances, with a subject matter they're less invested in, they would likely be able to understand how echo-systems work and what impact such systems would have across networks. The focus of the discussion, then, wouldn't merely be the facts being checked. It would be the coherence and explanatory power of the alternative explanation.

Such a discussion is certainly not guaranteed to dislodge false beliefs about Satanism or whatever the conspiracy theory might be. We are, all of us, strange and stubborn creatures. But if we guide

someone's attention to the coherency gaps they've never needed to notice, while at the same time offering an alternative story explaining not just *that* something is the case but *why*, that person may begin to see the deep memetic frames they experience the world through.

This is, of course, slow and subtle work, infuriatingly so when responding to harmful or dehumanizing frames. It can be outright cruel to tell habitually targeted people to wait patiently while someone holds their attackers' hands and tries to get them to see the world a little differently ("you need to be more gracious; the person who hates that you exist is *learning*"). It would be better if we had a switch to flip, and better still if marginalized people were not always asked to bear such disproportionate burdens. But encouraging people who cause harm to think about *how* they think and *why* they think is at least somewhere to start. Ideally, efforts to access deeper shared truths will arc toward justice—even while those efforts themselves highlight how much injustice there is.

Such outward effort requires inward reflection as well. As we stand there, considering what frames another person is seeing through, we too are seeing through frames. Those frames might correspond to objective reality. They might be ethically robust. But they're still filters that shape our experience of the world. At the least, what we see and know about the world can directly influence—and directly interfere with—what we're able to see and know about another person's experiences. Before we turn our attention to others, we must therefore ask ourselves: What frames am *I* seeing through? What cultivates *my* sense of reality, of goodness, of justice? What do *I* feel in my bones to be true about the world and my place within it? What do I think about how I think?

For white people, particularly white people who are middle-class, able-bodied, cisgender, and straight, this requires a careful look at the deep memetic frame that guides every aspect of life. Our own two lives very much included. This is, of course, the white racial frame, which directs sight to certain people, places, and things while muting or outright pathologizing other people, places, and things. Because this frame is so normalized and so central to mainstream life in the

US—and life in the Western world more broadly—it’s a particularly difficult frame for people seeing through it to detect. This is no accident. The more invisible the white racial frame is, the less questioned it is, and the less questioned it is, the more successful it can be. As sociologist George Lipsitz says, whiteness is nothing if not possessively invested in preserving itself.¹³ It doesn’t want to be seen. Otherwise it would be dismantled.

Of course, the invisibility of whiteness only works in one direction. Those who aren’t seeing through the frame have a much easier time detecting it—largely because they are harmed by it. This is an uncomfortable truth that many white people haven’t fully contended with. It’s easier to keep telling yourself the story that if you’re not a cross-burning bigot, if you’ve never targeted or harassed anyone because of their skin color or religion, you’re one of the good ones. It’s easier to see racism as an ugly relic of the past, one that *you’ve* gotten over, so why can’t everybody else?

You might not see your whiteness; you might not feel its effects. But whether you personally see it or not, the white racial frame harms people of color, and has for centuries. And not just people of color: the white racial frame harms everyone. The oceans of pollution that roared across the landscape because of the unexamined whiteness of early internet culture, and later because of the perfect confluence of white abstraction and white irony when faced with white supremacist violence, are a testament to what the white racial frame can do to the landscape. It must be seen. And it must be dismantled.

The act—and clearly it can be a distressing act—of exploring our own deep memetic frames is the final, most critical step in “you are here” network mapping. Triangulating our personal relationship to the technologies we use, the people we interact with, and the networks we navigate is critical. But the map will only ever be static and one-dimensional if it doesn’t address our own ideologies and experiences. To truly steward the land, we must understand how we shape the land—and how that land has shaped us.

You Are Here

The dense networks connecting everything to everything else online can be a source of profound destruction. They can also be a source of profound resilience. In her reflections on the natural world, Robin Wall Kimmerer reminds us what resilience looks like. In so doing, she reminds us that we already have the answers we need. All plants and all animals are linked within an intricate, interconnected gift economy. Humans often refuse and exploit these gifts. And yet, still, the natural world is guided by giving—between families, across species, throughout the entire ecosystem. “Such communal generosity might seem incompatible with the process of evolution,” Kimmerer observes, “which invokes the imperative of individual survival. But we make a grave error if we try to separate individual well-being from the health of the whole.”¹⁴

Our information ecosystem is no different, and the stakes are just as high. Too many people exploit and refuse the gifts of connection, instead embracing the bundles of rights that jealously guard what’s theirs. Too many people fixate on their own freedoms *from*, burying their responsibility to cultivate freedoms *for*.

This is certainly a grave error. How we tend our own soil affects the soil of those around us, and those around them, as the camera zooms out, revealing glowing connections across neighborhoods, counties, states, nations, and the entire globe. From such heights, there are no individual parcels of land to carve out, no property lines to defend, no untouched acres from which to bellow *not in my backyard*. This, in the end, is what the “you are here” map reveals. The networks of networks comprising our intertwined lives, the speed with which information travels, the effect of that information on so many unseen others, beats with reciprocity, and therefore with responsibility. If we choose instead only to see the rights we have, rather than the space we share, then the forest will die, and we will die with it.

It doesn’t have to be that way. We can choose, instead, to tend the land wisely, acknowledging always that our fates are connected.

Quietly, slowly, and with sincerity, we get there working side by side,
piling small things up until they're big.

Chapter 6

- . April Glaser, “Communities Rally around One Another—and Google Docs—to Bring Coronavirus Aid,” NBC News (March 20, 2020), <https://www.nbcnews.com/tech/tech-news/communities-rally-around-each-other-google-docs-bring-coronavirus-aid-n1164126>.
- . A growing corpus of studies has explored how bias is baked into the technologies we use. For studies of the inequality encoded into algorithms, see Cathy O’Neil, *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy* (New York: Random House, 2016); Safiya Noble, *Algorithms of Oppression: How Search Engines Reinforce Racism* (New York: NYU Press, 2018); and Virginia Eubanks, *Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor* (New York: St. Martin’s, 2018).
- . Whitney Phillips and Ryan M. Milner, *The Ambivalent Internet: Mischief, Oddity, and Antagonism Online* (Cambridge: Polity Press, 2017).
- . We borrow this framing from anthropologist Gregory Bateson’s popularization of the play frame. See his *Steps to an Ecology of Mind: Collected Essays in Anthropology, Psychiatry, Evolution, and Epistemology* (Chicago: University of Chicago Press, 1972). We’ve previously applied this concept to online jokes in Whitney Phillips and Ryan M. Milner, “Putting Jokes in Scare Quotes: How Sparkle Hair Scully Reveals What’s New and What’s Not about Humor Online,” *Flow Journal*, October 2, 2017, <https://www.flowjournal.org/2017/10/putting-jokes-in-scare-quotes>.
- . The COVID-19 crisis has moved the needle slightly, at least when it comes to false and misleading medical information; following the World Health Organization’s designation of COVID-19 as a global pandemic on March 11, 2020, platforms like Facebook and Twitter began aggressively moderating clear-cut lies about the virus and other exploitative advertising content. Some of these efforts have been downright communitarian. But they’re also very limited, as they’ve allowed—for example—QAnon and other Deep State theories to persist on-site, even as they’ve taken down—for example—false coronavirus cures. Given their track record on the subject, how long it took them to intervene (COVID-19 pollution flourished for months between the initial outbreak in Wuhan and the WHO declaration), as well as the sites’ binary cleaving between health information and political information (as if those two things don’t feed into each other), it seems unlikely that the decisive measures taken by the platforms will extend to political speech as well. The fact that the sites are employing any communitarian thought at all, however, provides some hope, and a foothold for further, more sweeping, intervention.
- . Patrick Davison, email interview with the authors, July 16, 2019.
- . Addressing in 2014 the growing tendency to conflate trolling with violent antagonism, ROFLCon cofounder Christina Xu made an uncanny prediction about equating deliberate harms with online play. “I feel there are a lot of dangerous things happening around that conflation, brewing under the surface, and they will bite us in the ass in five years or something,” she explained. Bite us in the ass they did indeed. See Tim Hwang and Christina Xu, “Lurk More’: An Interview with the Founders of ROFLCon,” *Journal of Visual Culture* 13, no. 3 (2014): 376–387.
- . For their case supporting strategic silence in newsrooms, see danah boyd and Joan Donovan, “The Case for Quarantining Extremist Ideas,” *Guardian*, June 1, 2018,

<https://www.theguardian.com/commentisfree/2018/jun/01/extremist-ideas-media-coverage-kkk>.

- . J. Nathan Matias, "Preventing Harassment and Increasing Group Participation through Social Norms in 2,190 Online Science Discussions," *PNAS* 116, no. 20 (2019): 9785–9789, <https://doi.org/10.1073/pnas.1813486116>.
0. Debra Walker King, *African Americans and the Culture of Pain* (Charlottesville: University of Virginia Press, 2008).
1. Stephan Lewandowsky, Ullrich Ecker, Colleen M. Seifert, Norbert Schwarz, and John Cook, "Misinformation and Its Correction: Continued Influence and Successful Debiasing," *Psychological Science in the Public Interest* 13, no. 3 (2012): 106–131; Brendan Nyhan and Jason Reifler, "Displacing Misinformation about Events: An Experimental Test of Causal Corrections," *Journal of Experimental Political Science* 2, no. 1 (2015): 81–93.
2. Lewandowsky et al., "Misinformation and Its Correction."
3. George Lipsitz, *The Possessive Investment in Whiteness: How White People Profit from Identity Politics*, 2nd ed. (Philadelphia, PA: Temple University Press, 2006).
4. Robin Wall Kimmerer, *Braiding Sweetgrass: Indigenous Wisdom, Scientific Knowledge, and the Teachings of Plants* (Minneapolis: Milkweed Editions, 2013), 16.