

Beyond Corporate Social Media Platforms: The Epistemic Promises and Perils of Alternative Social Media

Karen Frost-Arnold

This is a preprint, please cite final version published in *Topoi*

DOI: <https://doi.org/10.1007/s11245-024-10102-2>

Abstract

In recent years, we have witnessed increased interest in alternatives to the dominant corporate social media sites, such as Facebook, Twitter (now X), and TikTok. Tired of disinformation, harassment, privacy violations, and the general degradation of platforms, users and technologists have looked for non-corporate alternatives. Not-for-profit social media platforms emerging from free/libre and open-source (FLOSS) communities based on non-centralized infrastructure have emerged as promising alternatives. For applied epistemology of the internet, these alternative social media platforms present an opportunity to study different ways of producing knowledge together online. This paper evaluates the epistemic potential for such alternative, non-corporate social media. I present an epistemological framework for analyzing the epistemic promises and perils of alternative social media. Then I apply this framework to the case of Mastodon, a federated, open-source microblogging platform. Mastodon's structure and culture of openness present opportunities to avoid many of the epistemic perils of biased and untrustworthy large corporate platforms. However, Mastodon's risks include techno-elitism, white ignorance, and isolated, epistemically toxic communities.

Introduction

In recent years, we have witnessed increased interest in alternatives to the dominant corporate social media sites, such as Facebook, Twitter (now X), and TikTok. Tired of disinformation, harassment, privacy violations, and the general degradation of platforms (Doctorow 2023), users and technologists have looked for non-corporate alternatives. Elon Musk’s disastrous takeover of Twitter caused many users to flee the site, and some of them joined non-corporate social media spaces, such as Mastodon. There are many ways to evaluate the success of these non-corporate alternatives. Popular media have tended to focus on the number of people using these platforms, leading to an ongoing hype cycle of “boom and bust” narratives about the latest cool social media site, followed by accounts of its supposed demise (Zulli, Liu, and Gehl 2020). In contrast, internet studies scholars have examined alternative social media’s potential to show us different ways of living together online. Non-corporate social media platforms, no matter how large, can teach us about different economies, governance structures, and aesthetics that are driven by goals other than profit (Caelin 2022; Dunbar-Hester 2024; Gehl 2015; Gehl and Zulli 2022; Mannell and Smith 2022; Raman et al. 2019; Rozenshtein 2022; Stevenson and Valente Pinto 2024; Zuckerman 2020; Zuckerman and Rajendra-Nicolucci 2021; Zulli, Liu, and Gehl 2020). This paper evaluates the epistemic potential for alternative social media. For applied epistemology, which has increasingly turned its gaze on social epistemology of the internet, alternative social media present an opportunity to study different ways of producing knowledge together online.

Section 1 introduces alternative social media platforms, with an eye to details that are particularly epistemically relevant. Section 2 outlines the epistemological framework that I will use throughout the paper. Section 3 applies this framework to analyze the epistemic promises

and perils of alternative social media, focusing on a case study of a federated alternative social media platform: Mastodon.

1. What are alternative social media platforms?

Robert Gehl coined the term “alternative social media” (ASM) to describe non-corporate social media projects that are explicitly designed to avoid the problems of corporate social media (CSM), such as Facebook, Twitter, and TikTok (Gehl 2015). ASM are diverse and constantly changing, with different projects being created and falling into disuse. While there is no concrete set of features that all ASM share, there are some patterns that tend to reoccur in ASM.

First, rejecting the idea that social media platforms should be primarily a profit-making enterprise, ASM eschew advertising. By disallowing advertising, ASM resist the notion that people with money to pay for advertising can have more of a voice in social media spaces (Gehl 2015, 5). This also enables ASM administrators to focus on the needs of users rather than the demands of advertisers and shareholders (ibid.). The lack of advertising and profit-motive shapes many features of ASM. For example, without advertising, ASM lack an incentive to push users to spend time on platforms by seeing engaging and paid-for content boosted by an algorithmically-generated feed. Mastodon, for example, shows users posts in a chronological feed as posts appear in the order they are posted. This means that what users see is less susceptible to manipulation by platform providers and advertisers.

A second common feature of ASM is that many ASM use non-centralized infrastructure, meaning that the data is not stored in central servers, as is the data in CSM. Some ASM, such as Scuttlebutt, use decentralized peer-to-peer technology to create a “network of devices that share

the responsibility for data storage” (Mannell and Smith 2022, 5). In the Scuttlebutt network, the data for your user profile is stored on your device, and also backed up on the devices of your friends using Scuttlebutt. Other ASM, such as Mastodon, have a federated infrastructure in which an individual or group of admins for a Mastodon community (called an ‘instance’) host the data for the community on an independent server. Each community instance can choose to connect with other instances hosted on other independent servers. This form of non-centralized social media creates a federated network, which is “a special type of network topology where ‘each resource provider maintains local autonomy, and the ability to set policy for use of its own resources’ while allowing members to connect across the network (Berman et al. 2014, 17)” (Gehl and Zulli 2022, 3).¹ This relative independence and autonomy balanced with connection to others shapes many features of Mastodon’s federated social media. For example, each Mastodon instance can create its own moderation policies for speech in the community. However, to be listed on joinmastodon.org, the primary entry point for new users looking to find a Mastodon instance, each instance must commit to the Mastodon server covenant (Gehl and Zulli 2022, 12), which requires (1) “active moderation against racism, sexism, homophobia, and transphobia,” (2) daily backups of data, (3) “at least one person with emergency access to the server infrastructure,” and (4) “commitment to give users at least 3 months of advance warning in case of shutting down instance” (“Mastodon Server Covenant for Joinmastodon.Org,” n.d.). Gehl and

¹ In a move that has been controversial among Mastodon users, Threads, Meta’s CSM microblogging platform, has joined Mastodon’s federated network creating some communication between Mastodon instances and users on Threads (Forristal 2024). Despite these connections, Threads remains a thoroughly corporate social media platform with centralized moderation and profit-driven incentives. Bluesky is another CSM platform that is slightly closer to the ASM model in that it runs on an open-source protocol and is working to develop a federated network of servers. However, Bluesky is a corporate platform with profit-driven incentives that currently has centralized moderation on its one server (Silberling, Stringer, and Corral 2024).

Zulli call this governance model “covenantal federalism,” describing it as a system “where small units consent to band together while abiding by a shared ethical code” (Gehl and Zulli 2022, 2).

A third feature of ASM is their grounding in free/libre and open-source software (FLOSS) communities. FLOSS communities promote users’ freedom to “run, copy, distribute, study, change and improve the software” (GNU Project, n.d.). CSM, such as Twitter, depend on proprietary software which is inaccessible to users, and which users cannot modify for their own needs. In contrast, developers behind ASM create and share tools that can be studied and changed by anyone to create a variety of social media platforms with different affordances. This ASM-FLOSS culture promotes what Gehl calls “network and code pedagogies,” which teach users technical skills (Gehl 2015, 7). Gehl quotes Matt Lee from GNU social (an earlier micro-blogging ASM site) saying, “it is vitally important to me that anyone can set up a GNU social server on virtually any web hosting. I also want to make it as easy as possible to set up and install. To that end, I will personally help anyone who wants to get set up” (qtd. in Gehl 2015, 7). Thus, as several internet studies scholars have argued, ASM provide a more fully participatory model of social media than CSM, one where users not only participate by creating content, but also participate in the design and maintenance of the infrastructure and education of the community (Gehl 2015; Gehl and Zulli 2022; Mannell and Smith 2022). Each of these three common features of ASM (no advertising, non-centralization, and FLOSS grounding) will be relevant to evaluating the epistemic value of ASM. The next section provides a normative framework for making that evaluation.

2. A framework for evaluating epistemic promises and perils

This paper draws on Deborah Johnson’s promises-and-perils approach to evaluating emerging technologies (Johnson 2021) and Karen Frost-Arnold’s approach to social epistemology for the internet (Frost-Arnold 2023). Johnson provides a framework for ethical evaluation of the plausible outcomes of new technology, and Frost-Arnold presents a way to adapt this framework for epistemological evaluation of technology. I will explain each and outline how they will be combined in this paper.

Johnson begins with a sociotechnical systems perspective according to which technologies are “ensembles of artifacts together with social practices, social relationships, institutional arrangements, and forms of knowledge” (Johnson 2021, 648). On this approach, the material objects of technology should be analyzed as shaped by (and shaping) the social arrangements in which they are (or will be) embedded. This means that in order to assess what effect a piece of technology will have on human flourishing, we need to take into account not only its design and how its parts work together to produce an outcome, but also the social arrangements that determine how it will be used and what meanings it will have. For example, to understand the potential impact of autonomous cars, we need to not only understand how well the steering mechanisms and GPS navigations will function, we also need to look at the legal regulatory environment in which they will be used, what the public level of trust is, and how these cars will change our own understanding of how we move through space (Johnson 2021, 657–58).

One common way of evaluating the outcomes of technology is cost-benefit analysis, which Johnson argues has several limitations (Johnson 2021, 651–55). The type of cost-benefit analysis (CBA) that Johnson critiques provides a consequentialist analysis of the predicted consequences of a piece of technology by measuring the value of each outcome in monetary

terms. If the monetary benefits to society of a piece of technology outweigh the monetary costs, then the technology is considered beneficial. Johnson argues that this type of CBA analysis of emerging technologies tends to overestimate benefits, underestimate costs, face difficulty with nonquantifiable costs and benefits, fail to identify unforeseen consequences, and ignore considerations of justice. In contrast, Johnson's approach to promises and perils analysis (PPA) recognizes multiple values, instead of simply monetary value. PPA attempts to identify the plausible outcomes of a piece of technology and evaluates not only how they will contribute to human happiness, freedom, and privacy, but also to what extent the benefits and costs are distributed justly, fairly, and equally (Johnson 2021, 653–55). On this approach, to evaluate an emerging technology, we need to ask many questions about what is likely to happen when this artifact is used and given meaning in a particular social context, and we can evaluate these outcomes using several different ethical values.

This paper adapts PPA to social epistemology for the internet by drawing on Frost-Arnold's account of several different epistemological frameworks that draw our attention to different relevant epistemological features of technology (Frost-Arnold 2023). Frost-Arnold shows how (1) feminist accounts of objectivity, (2) veritism, (3) epistemologies of ignorance, (4) virtue epistemology, and (5) epistemic injustice (what Frost-Arnold calls the FOVIVI frameworks) can provide mutual reinforcing evaluative frameworks for epistemological analysis of the internet.² Feminist accounts of objectivity show the epistemic importance of diverse networks of trust that highlight the voices of the oppressed for weeding out bias online. Veritism highlights the value of truth and evaluates sociotechnical systems by how well they help agents

² For a different approach to a similar set of epistemological tools for social epistemology of the internet, see (Habgood-Coote 2024).

obtain true beliefs and avoid false beliefs. Epistemologies of ignorance reveal the mechanisms by which ignorance is constructed and maintained on the internet. Virtue epistemology explains how individual agents can cultivate virtuous habits conducive to the dissemination of knowledge online. The epistemic injustice literature identifies ways that knowers are disrespected and knowledge is erased due to individual prejudice of users and systemic oppression in online platforms.

Combining PPA with the FOVIVI frameworks provides a robust set of theoretical tools to evaluate the epistemic merits of emerging internet technologies. As with PPA, we can ask what the plausible outcomes of implementing the technologies in particular social environments are likely to be. Then, we can use the FOVIVI frameworks to ask which of these outcomes are likely to be epistemically promising or epistemically perilous. Just as PPA combines both consequentialist considerations of outcomes with deontological considerations of justice, a FOVIVI-based PPA analysis will look at the outcomes of internet technologies in terms of producing true beliefs and weeding out false beliefs, ignorance, and bias, while also paying attention to considerations of epistemic justice. A focus on epistemic virtue not only adds a focus on epistemic responsibility and virtue, but also illuminates the role that social structures play in shaping who we are online. One addition this paper makes to the FOVIVI frameworks is by including the acquisition of knowledge-how as an epistemic benefit. Several critics of Frost-Arnold's FOVIVI approach have argued that veritism's focus on true belief as the primary epistemic value is too narrow for social epistemology of the internet (Medina Vizueté Forthcoming; citation blinded for review). They argue that online practices should also be evaluated according to how well they produce other epistemic consequences for users, such as understanding and knowledge-how, that are not clearly cases of users simply having true beliefs.

I find these arguments convincing, and this paper will investigate to what extent ASM provide users with opportunities to gain a variety of epistemic benefits in addition to true belief. Given this framework, the rest of the paper will identify several epistemic promises and perils of alternative social media.

3. The epistemic promises and perils of alternative social media

3.1. Epistemic promises

3.1.1. Moderation, epistemic justice, and trust

ASM provide opportunities for epistemic communities to set their own moderation policies, which can promote epistemic justice and epistemically beneficial trust. To make this argument, it will be helpful to first present some of the problems with moderation on CSM. Some of the central problems with CSM are that the corporate policies for what can and cannot be said on the platform (and who can and cannot speak) are made with limited input from the users, lack transparency or meaningful opportunity for appeal, and are often biased against marginalized people (Barnes 2022; Frost-Arnold 2023; Kaye 2019; Roberts 2019; “Santa Clara Principles on Transparency and Accountability in Content Moderation” 2018; Stewart, Cichocki, and McLeod 2022; York 2021). There are many reasons why CSM tend to exhibit these moderation failings, including profit-based incentives to protect speech by dominant groups, profit-based incentives to perform moderation cheaply (by insufficiently supporting human moderators and relying on biased algorithmic moderation), lack of diversity in tech company management, and a top-down imposition of moderation rules on users. As I will argue, the three features of ASM described above provide an opportunity for epistemic communities to create

their own moderation policies that best fit their needs, rather than the needs of advertisers and management.

To illustrate these problems with CSM further, let us focus on problems with racist moderation. CSM have a persistent problem with both allowing racist hate speech on platforms and also punishing people of color who speak out against racism. For example, Frost-Arnold documents how Facebook moderation policies encode an interpretive framework aimed at protecting white people from co-called ‘reverse racism’ (Frost-Arnold 2023). In 2017, leaked slides from a training presentation for Facebook moderators showed that Facebook took white men to be a protected category, while bizarrely denying this status to other groups, such as Black children (Angwin and Grassegger 2017). These kinds of policies cause people of color who speak out against racism to be punished for what are viewed as attacks on white users. Not only have Facebook’s human moderators been involved in racist moderation, but when Facebook used algorithmic moderation, its algorithms were set up to be “race-blind” in ways that “resulted in the company being more vigilant about removing slurs lobbed against White users while flagging and deleting innocuous posts by people of color on the platform” (Dwoskin, Tiku, and Kelly 2020). A civil rights audit of Facebook recognized that “too many mistakes were being made removing content that was actually condemning hate speech” (“Facebook’s Civil Rights Audit - Final Report” 2020, 45). These biased moderation practices perpetuate epistemic injustice against marginalized people, as their voices are discounted, misinterpreted, and silenced (Frost-Arnold 2023; Stewart, Cichocki, and McLeod 2022).

Eugenia Siapera argues that the imposition of such racist moderation policies on the world can be understood as a kind of colonialism (Siapera 2022). Siapera draws on Quijano’s conception of the “coloniality of power” as a type of power manifested when colonial powers

impose their racial classifications on the whole world (Quijano 2007). She notes that CSM, such as Facebook, create racial classifications (such as making white men a protected category) and policies (such as imposing penalties on users who criticize white racism) which they then impose on the whole world. She argues that “In dismissing the views and knowledge of targeted communities, and in excluding them from participating in the definition of their own oppression, content moderation policies continue and reinforce the colonial tactic of dismissing the knowledge of colonized communities” (Siapera 2022, 62). CSM, with their moderation policies created by corporate management located in the Global North, impose their own understanding of racism and hate speech on the rest of world who use their platform. In the face of this colonialism, it is helpful to turn to decolonial theory to see alternative ways of living. The following explanation of Franz Fanon’s work provides a helpful entry point to seeing the value of ASM: “For Fanon, struggles for decolonization are first and foremost about self-ownership... In his eyes, self-ownership is a precondition, a necessary step towards the *creation of new forms of life* that could genuinely be characterized as fully human” (Mbembe 2015, 12).

ASM provide the potential for self-ownership of social media spaces in ways that can avoid the problems with CSM’s approach to moderation. When people can create their own social media communities, set their own moderation policies, and host their own data, they are not subject to corporate policies imposed from corporate powers. On Mastodon, special-interest and identity-based communities have created their own instances with specific ground rules to support the kinds of discussion they want and to prevent the kinds of attacks to which members of their group are particularly vulnerable (e.g., policies against deadnaming in LGBTQ spaces). This freedom to create their own moderation policies can lead to more just policies not only in spaces specifically designed for marginalized people, but also in other groups. For example, I

have an account on a Mastodon instance operated by an academic professional organization that has very detailed moderation policies designed to protect the speech of marginalized people in the space. Thus, one way in which ASM are epistemically promising is that they provide opportunities for people to create communities that are designed and governed to avoid epistemic injustice.³

This self-determination also provides the potential for decolonization of our epistemic infrastructure. Unlike CSM, the policies governing speech, data collection, and data storage in ASM are not imposed on local communities by corporate agents primarily located in the Global North (Kwet 2018; 2021). Thus, when ASM are used by colonized people in the Global South, they can decolonize their online epistemic communities. Veli Mitova characterizes epistemic decolonization as having both negative and positive programs. Negatively, epistemic decolonization involves the “elimination of unreflective western influences on our knowledge supplies and production” (Mitova 2020, 2). Positively, epistemic decolonization includes “proactive utilisation of the marginalised epistemic resources of the colonised in the advancement of knowledge in various fields” (ibid.). Epistemic decolonization is often analyzed in terms of removing the Western concepts and ways of knowing that have been imposed on colonized people and replacing them with Indigenous concepts and worldviews. But colonialism also damages the epistemic infrastructure that facilitates the production and dissemination of knowledge. For example, colonial domination of academic infrastructure has lasting impacts, such as the curricular entrenchment of “philosophy” departments that only teach Western

³ In fact, Mastodon was created by Eugen Rochko in order to provide an alternative to the hate speech and white supremacy he saw on Twitter (Zulli, Liu, and Gehl 2020). That said, in sections 3.2.1 and 3.2.3, I will discuss how one of the perils of ASM is that they also provide opportunities for people to create communities hostile to marginalized people

philosophy, or the dominance of academic conferences in Global North countries with restrictive visa requirements that often limit participation by colleagues from the Global South. Similarly, digital colonialism shapes our online lives in a myriad of ways, such as the imposition of moderation policies on the world by Global North companies (cf. Kwet 2018; 2021).⁴ Claiming control of their own social media infrastructure and creating policies that best fit their own ways of sharing knowledge is one way that colonized people can take steps to decolonize their online epistemic infrastructure.

This self-ownership can also promote epistemically beneficial trust. Frost-Arnold argues that trust in the “practices, institutions, and social structures that create avenues for critical discourse” is essential for objectivity and truth acquisition (Frost-Arnold 2023, 76). If people lose faith in the avenues for discussion, then they are likely to stop participating in those spaces. For veritists, this loss of trust is clearly epistemically damaging because the community loses the true claims that would be shared and also loses the critiques that would help others identify false beliefs. Feminist accounts of objectivity also recognize the damage of distrust in public avenues for discussion, especially when marginalized people lose trust, because when people withdraw from discussion, they stop pointing out the biases present in those spaces and discussions. The degradation of discourse on Twitter illustrates this epistemic demise of social media spaces when trust is lost. Following Musk’s acquisition of Twitter, hate speech increased (Elliott 2023), and usage of the app has decreased relative to other social media apps (Hern 2024).

When ASM are built by users who can participate in the creation and maintenance of their own spaces, they gain reasons to trust the space. Even if not every user of a Mastodon

⁴ For a thorough analysis of other forms of digital colonialism see (Kwet 2019).

instance is involved in the original creation or the day-to-day maintenance of the space, there are other features of federated social media that can ground trust: accountability and transparency. Gehl and Zulli argue that users' ability to migrate their data from one Mastodon instance to another federated instance is a key mechanism for accountability that grounds users' trust in the platform (Gehl and Zulli 2022). Mastodon is different than CSM, such as Twitter, in that if a user decides to leave a particular instance on Mastodon and move to another, they can take their account with them; whereas on Twitter, if you leave, there is no way to take your account elsewhere. This means that administrators of Mastodon instances have a strong incentive to take steps to maintain the trust of their users: "[J]ust as users must consent to abide by an instance's code of conduct, so must an administrator foster trust to avoid users withdrawing their consent to be governed and transferring it to another instance whose administrator might be a better steward of their participation and data" (Gehl and Zulli 2022, 8–9). One might argue that CSM administrators also have incentives to maintain the trust of their users in order to keep eyeballs for advertising revenue and maintain a profit-making platform. However, open-source software developer and advocate Coraline Ehmke argues that this trust in CSM is often coerced:

“Without accountability and transparency, users can only see what providers *say*, not what they actually *do*. This makes an informed assessment of a platform's trustworthiness nearly impossible. Combined with an artificially high cost of opting out of a particular platform, users are routinely coerced into risking their own safety and well-being in exchange for access.” (Ehmke 2023)

Emke calls this phenomenon “trust coercion,” and she argues that the fact that users cannot take their accounts and network of relationships with them when they leave a CSM platform weakens the incentive for CSM administrators to act trustworthily. Additionally, as she mentions in this passage, CSM lack the kind of transparency that open-source software provides to users. The ability to investigate the source code and participate in open-source community discussions

makes ASM much more transparent than CSM. For all these reasons, users often have more reason to trust ASM as public avenues for discussion than do users of CSM, and as I have argued, this trust is epistemically beneficial.

3.1.2. Scale

Federated ASM, such as Mastodon, have promising ways of handling the opportunities and challenges of scale in social media. To see the merits of how Mastodon handles issues of scale, let us first investigate scale in CSM.

On the one hand, the large number of users and the global reach of CSM constitute some of the great benefits of CSM. Social media platforms exhibit network effects because the more people who use the platform, the more value the platform has (Constine 2020). This is true for many non-epistemic values. For example, if I want to reconnect with friends from my childhood, they are more likely to be on a larger social media platform than a smaller one. And if I want to be entertained by funny memes, then a larger social media platform will be more likely to circulate more memes. Networks effects can also increase the epistemic value of social media platforms. Veritism can help us see this. Goldman's veritistic social-epistemology accords more epistemic value to true beliefs that answer questions that interest the inquirer (Goldman 1999, 87–96, 351–52). For example, the question “What is the number of hairs on my cat's head?” is likely to be of much less interest to inquirers than the question “What is the number of wildfires spreading across my region?” Veritistic systems-oriented social epistemology evaluates sociotechnical systems according to how well they help inquirers form true beliefs about questions that interest them. Networks effects are relevant here. A large-scale social media platform can give me access to many users and many conversations from all over the world. I

have a better chance of getting answers to questions that interest me on larger social media platforms than smaller ones.⁵ Not only is it more likely that someone will be posting about something that I am interested in on a larger platform, but I have a variety of interests, and large, global platforms are likely to have a variety of topics discussed. Thus, the global reach and large scale of CSM present epistemic promises.

On the other hand, scale also presents epistemic challenges for CSM. As platforms grow, so do the number of bad actors drawn to a site. This leads to increases in hate speech, harassment, internet imposters, conspiracy theories, and fake news, all of which can be epistemically damaging. To compound these problems, moderation at scale is notoriously difficult for CSM platforms (Frost-Arnold 2023; Gorwa, Binns, and Katzenbach 2020; Roberts 2019). The sheer quantity of problematic content that needs to be removed overwhelms the mechanisms that CSM companies have put in place. For example, 500 hours of video are uploaded to YouTube every minute (Roose and Conger 2019). Relying on users to flag offensive content, hiring tens of thousands of poorly paid human moderators to screen flagged material, and using algorithmic filtering all present epistemic problems of bias, unreliability, and epistemic injustice (Frost-Arnold 2023). Thus, while enabling many users to communicate with each other is one of the epistemic promises of CSM, typical CSM governance models often fail to properly handle the epistemic perils of large-scale communication platforms.

Federated ASM have a different structure that combines the opportunities of both small and large-scale platforms. Mastodon enables users to form their own small interest-based communities. Moderating smaller communities is much easier for administrators, so Mastodon

⁵ One complication here is the problem of ‘epistemic flooding’ identified by Glenn Anderau in which users are exposed to more information than they can responsibly process (Anderau 2023).

instances can avoid some of the problems CSM face with moderation at scale. ASM administrators are dealing with a smaller quantity of content to handle, so they typically do not rely on poorly paid and poorly supported moderation workers, and they do not have to use biased algorithmic moderation. With more time to carefully assess potentially problematic content and to work to instill a healthy culture in the community, administrators for smaller ASM communities can be more reliable moderators. Additionally, since users can choose instances based on their moderation guidelines, users may be more motivated to follow the rules of the community. So Mastodon's open-source infrastructure and culture of self-determination facilitates the creation of smaller communities that are easier to moderate and create a climate in which epistemically valuable discussion can occur. Furthermore, Mastodon's federated structure ensures that users are not stuck in small islands without the benefits of larger-scale communities. For example, while I have an account on a relatively small Mastodon instance run by a professional organization, I am not limited to only seeing posts by members of this community. Since this Mastodon instance is federated with many other Mastodon instances, I can see posts from a variety of other communities. Thus, members of federated Mastodon communities can gain the epistemic benefits of learning from people from a wide variety of backgrounds and expertise. This enables users to gain true beliefs about a variety of questions that interest them. So, federated ASM combine the benefits of well-moderated small communities with the power of a larger network.

One might ask how federated ASM avoid the problems of large-scale CSM, since as one is connected to more users through federation, the chance of encountering bad actors would seemingly increase. Even if I have an account in a well-moderated small instance that protects me from epistemically damaging disinformation and harassment, am I not exposed to these

harms when I read posts on other federated communities? This is where the tool of blocking or “defederating” an instance plays a useful role. Administrators of a Mastodon instance can choose to block posts from other Mastodon instances, so that members of their community are protected from the blocked instance.⁶ Gehl and Zulli provide a case study of this process with respect to Gab, the right-wing extremist social media network (Gehl and Zulli 2022, 11–13). Gab was created in July 2019 using Mastodon’s open-source protocol. As a cesspool of bigotry and harassment, Gab grew as a right-wing social media space and was associated with the January 6 insurrection at the U.S. Capitol. The existence of Gab as part of Mastodon gave the platform much negative press, but, as Gehl and Zulli document, it was less accurately reported how Mastodon administrators took several steps to denounce and isolate Gab from the rest of the network. Gab’s policies failed to follow the Mastodon Server Covenant, so Gab was not listed as a server on joinmastodon.org. And Gab was blocked by servers in the covenantal federation. The result was isolation of Gab and protection from its harms for users in the other federated instances. As one admin quoted by Gehl and Zulli described:

They [Gab] are off on their own little island and hardly anybody actually communicates with them. They showed up, everybody was pretty well prepared to close them [out] as soon as they did...Almost every server on the fediverse implemented moderation features to lock Gab out. (qtd. in Gehl and Zulli 2022, 12)

Thus, as Gehl and Zulli’s useful case study shows, federated ASM have blocking tools to protect users from some of the harms of connections to a larger federated community. This is by no means a perfect system, as I will discuss in section 3.1.3. Nonetheless, the ability to separate one

⁶ Note that this is an additional blocking tool that Mastodon provides on top of the ability to block individual users which Mastodon and many CSM platforms provide. Blocking individual users (and using blocklists) on CSM can be time consuming, confusing, and awkward for individual users. By shifting some of the labor of blocking bad actors from the shoulders of individuals to the responsibility of admins who can defederate, Mastodon can make the process of handling bad actors easier in some ways. That said, defederation cannot solve all problems with bad actors (Captain 2022).

epistemic community from another dysfunctional community can have some epistemic benefits in ASM.

3.1.3. Developing skills

If we return to the third common feature of ASM (i.e., their grounding in FLOSS communities), we find a clear epistemic promise of ASM: the opportunity they provide for users to develop technical skills. Recall Gehl’s argument that ASM emphasize “network and code pedagogies” (Gehl 2015, 7). Open-source communities highly value freedom of knowledge and software that provides people the freedom to investigate, tinker with, and adapt software. With this ethic as a foundation, ASM communities often devote time and resources to teaching new users how to navigate the platform. And the openness of the source code provides people an opportunity to look “beneath the hood” of social media and learn how this works. With their propriety source code and lack of an ethical commitment to pedagogy, CSM lack these kinds of opportunities. Thus, ASM are epistemically promising in their ability to foster technical knowledge-how.

3.2. *Epistemic perils*

3.2.1. Epistemic bubbles and echo chambers

Recent social epistemological work on epistemic bubbles and echo chambers suggests the first peril of ASM: the formation of epistemically toxic communities that are isolated in various

ways.⁷ C. Thi Nguyen defines an epistemic bubble as “A social epistemic structure in which some relevant voices have been excluded by omission” (Nguyen 2020, 142). Online epistemic bubbles are formed in many ways, including users choosing to follow certain people and block other users, search engine personalization, and algorithmic filtering of social media feeds. Nguyen defines an echo chamber as “a social epistemic structure in which other relevant voices have been actively discredited” (ibid.). Online echo chambers have been implicated in epistemically damaging phenomena such as socially constructed ignorance, conspiracy theories, extremism, and polarization (Alfano et al. 2021; Levy 2023; Nguyen 2020; 2021; Santos 2020; Talisse 2021). Epistemic bubbles differ from echo chambers in that in an epistemic bubble a user is at risk of not hearing evidence that challenges their false beliefs, but in an echo chamber the user will hear that evidence, but will discount it due to distrust. The FOVIVI frameworks recognize that both of these outcomes can be epistemically damaging. According to veritism, they can prevent users from forming true beliefs and avoiding false beliefs. Both epistemic bubbles and echo chambers can perpetuate biases by preventing users from having their beliefs subject to transformative criticism; thus feminist accounts of objectivity also raise concerns about these phenomena. Finally, epistemic bubbles and echo chambers can cause epistemic injustice when social media algorithms compound and amplify the discrediting of marginalized voices both online and offline (Stewart, Cichocki, and McLeod 2022).

It is important to recognize that epistemic bubbles and echo chambers are not always epistemically harmful. As we saw in section 3.1.1, protecting marginalized users from harassment by bad actors can play an important epistemic role in social media communities.

⁷ Not all epistemologists find the terms ‘epistemic bubble’ and ‘echo chamber’ philosophically fruitful. For critiques of these concepts, see (Coady 2024; Munroe 2023).

Campbell and Furman argue that this filtering and also active discrediting of hostile voices can sometimes be epistemically beneficial, especially when they provide safe spaces for marginalized people (Campbell 2023; Furman 2023). Frost-Arnold and Habgood-Coote, Ashton, and El Kassar argue that counterpublics of marginalized people who develop their own knowledge to fight oppression can play an epistemically valuable role online (Frost-Arnold 2023; Habgood-Coote, Ashton, and El Kassar 2024). Begby, Fantl, and Lackey argue that filtering information and constructing echo chambers can sometimes protect users from systemic misinformation and fake news (Begby 2022; Fantl 2021; Lackey 2021). Barbarrusa and Medina Vizquete show that epistemic bubbles can help marginalized people develop epistemic resources necessary for fighting hermeneutical marginalization (Barbarrusa and Medina Vizquete This issue).

That said, the example of Gab illustrates a potentially perilous manifestation of these phenomena in ASM. Promoting itself as a haven for right-wing elements that have been deplatformed from CSM, Gab drew a significant number of users into an epistemic bubble and echo chamber. Since the ideology that attracts users to Gab is predicated on a distrust of so-called ‘woke’ culture on CSM, it is ripe for becoming an echo chamber in which outside views are distrusted. Gab displays the features of an epistemic bubble due to not only users’ choices not to follow others who do not share their views, but also the choices of the admins of other Mastodon servers who defederated from Gab, leaving them isolated. While in section 3.1.2, I argued that defederation can have some epistemic benefits by protecting users from bad actors, we are now in a position to see that it can also have epistemic costs through the creation of communities of bad actors who are isolated in epistemically toxic spaces. Whether the epistemic benefits of this feature of ASM outweigh the epistemic costs is beyond the scope of this paper,

since it depends on empirical research on the scope and magnitude of both the benefits and costs.

As Zuckerman and Rajendra-Nicolucci put it:

Deplatforming bad actors from mainstream social media... likely limits the spread of their ideas. However, it could also lead to insular echo chambers, filled with only the most devoted extremists, a dynamic that could lead to further radicalization and extremism... Understanding the trade-offs of deplatforming, and looking for ways to combat extremism on alt-tech platforms will be crucial as the space develops. (Zuckerman and Rajendra-Nicolucci 2021)

3.2.2. Scale, elitism, and dictators

Despite FLOSS communities' commitment to openness and teaching new users technical skills, ASM run the risk of remaining relatively small and geographically narrow social media platforms. In October 2023, Mastodon had 1.8 million monthly active users (Perez 2023), while in the second quarter of 2023 Facebook had roughly 3 billion monthly active users (Dixon 2024). Mastodon is a very small social media platform compared to CSM competitors. Mastodon also has limited global reach. Raman et al. found that between 2017 and 2018, 89.1% of posts were on instances in Japan, the U.S., and France (Raman et al. 2019, 6). There are many more Mastodon servers in the Global North than the Global South. This narrowness of Mastodon threatens to prevent ASM from attaining the epistemic promises discussed in sections 3.1.1 and 3.1.2. The dominance of the Global North in the fediverse undermines the epistemic potential for Mastodon to be a vehicle for epistemic decolonization. If colonized people are not using Mastodon, then it simply will not function as a tool for decolonizing online epistemic infrastructure. Similarly, if Mastodon remains a relatively small social media platform, then it will not provide close to the same epistemic benefits of scale that CSM can offer.

There are many reasons why Mastodon has fewer users than CSM competitors, one of which is particularly epistemically relevant: techno-elitism. Gehl and Zulli identify “techno-elitism” as the main form of exclusion preventing Mastodon’s growth (Gehl and Zulli 2022, 13). There is much a steeper learning curve required to start using Mastodon compared to CSM platforms. It requires users to understand the federated structure of Mastodon before they know that they need to find an open instance.⁸ Then they need to find an instance that meets their goals. And then learning how to set up a new account and follow others is also a relatively complicated affair. If simply starting a Mastodon account is complicated, this pales in comparison to the technical knowledge required to start and host one’s own instance. Thus, while ASM are more open than CSM because they theoretically provide the potential for anyone to host their own social media community, in reality only a narrow group of technically skilled people can do so. As Gehl and Zulli put it, “Mastodon thus remains a niche system for free and open-source software enthusiasts and other more technically minded internet users, and this has brought about the exclusion of the vast majority” (Gehl and Zulli 2022, 13). This exclusion is epistemically significant since it means that the epistemic benefits of ASM are less widely shared than they could be. Additionally, since the population of FLOSS community participants who have the technical skills to fully engage with ASM tends to skew more white and male than the general population, ASM lack diversity (Dunbar-Hester 2020; Gehl and Zulli 2022). I will discuss the whiteness of Mastodon in more detail in section 3.2.3, but here I will note that a lack of diversity in ASM makes these platforms vulnerable to some of the problems of bias and epistemic injustice discussed in section 3.1.1.

⁸ An open instance is one receiving new members. When an instance experiences rapid growth that outstrips the admin’s ability to responsibly manage the community, Mastodon instances can close themselves to new members.

This risk is exacerbated by the tendency of FLOSS communities to move towards hierarchical structures. Dunbar-Hester’s study of FLOSS communities explains how differences in technical expertise and reputational capital can create hierarchies within hacker communities, despite their commitment to democratization (Dunbar-Hester 2020, 51). This pattern has repeated in Mastodon with Eugen Rochko, the creator of Mastodon, acting as what some critics have called a “benevolent dictator for life” (Valens 2019). Valens describes how many queer users initially found Mastodon an attractive space in 2016 after Rochko created it to be an alternative to Twitter by providing a space that banned Nazis (Valens 2019). Many queer coders invested time in the project by developing features for Mastodon, but some were frustrated by Rochko’s failure to listen to some of their suggestions and give them adequate credit for their work. Additional complaints have been raised against Rochko’s moderation of the flagship, largest Mastodon instance, Mastodon.social, with concerns that he does not take necessary steps to protect queer users and people of color on the site (Captain 2022; Valens 2019). For his part, Rochko has called Mastodon “the child of my imagination,” arguing that “[I] created it the way that I wanted to do it” and that “[I] did things the way I wanted them to work” (qtd. in Valens 2019).

Dictator-led ASM communities undermine the epistemic potential of ASM. As feminist accounts of objectivity show, moderation policies created by communities that include the voices of the oppressed are less likely to encode biases and perpetuate epistemic injustices, while policies created by one individual are more liable to bias. Even a benevolent dictator who tries to moderate fairly is likely to have some unconscious biases or socially constructed ignorance due to the limitations of their own social location. Additionally, when a dictator does things the way they want them to work, they are liable to lose the trust of users when users perceive that their

concerns are ignored. This means that dictator-led communities can fail to have all of the epistemic benefits of objectivity and epistemic trust discussed in section 3.1.1.

3.2.3. The whiteness of alternative social media

What has been called “the whiteness of Mastodon” illustrates a third epistemic peril of ASM (Doyle-Burke and Smith 2023; Flowers and Hendrix 2022). Johnathan Flowers has examined how social norms and choices about the design and use of Mastodon’s features have created a space that centers white users and creates a hostile climate for people of color (ibid.). In this section, I draw on Flowers’ analysis to argue that ASM platforms can manifest socially constructed ignorance and epistemic vice.

People of color on Mastodon have raised concerns about racism on the platform and failures to address it for years (Valens 2019). These concerns drew more attention in 2022 after Mastodon received its largest influx of users as people fled Twitter following Musk’s takeover (Captain 2022). One of the much-discussed problems during this time was social norms about the use of content warnings on Mastodon (Flowers and Hendrix 2022). Unlike Twitter, Mastodon’s interface includes a feature that allows users to add a content warning (CW) to a post. The CW is the only thing that other users see about the post unless they click a button to “show more” (or unless they have applied a setting to open all CWs). While the practice of adding some text at the top of a post to describe the nature of the rest of the post is a somewhat common practice in many social media platforms, Mastodon’s design that enables users to hide a post behind a CW was novel to new users migrating to the platform. What was also novel to many new users was the culture of extensive use of CWs on Mastodon. In 2022’s Mastodon

culture, it was already an entrenched (if not uncontested) practice to use CWs for a broad range of topics, including topics that might be traumatic or difficult for some users (e.g. sexual assault, photos with eye contact), spoilers for TV or movies, or tedious and depressing (e.g., U.S. politics) (buckenham 2022; Stefanello 2022). With the influx of new users, older Mastodon users tried to instill norms of widespread CW usage.

There are two problems with how this attempt to instill norms of CW usage encodes a culture of whiteness: the nature of the norms surrounding CWs about racism, and the manner in which white users try to enforce the norms. People of color on Mastodon often reject the idea that their posts about their experiences with racism should be hidden behind a CW (Flowers and Hendrix 2022; Valens 2019). Recall that if a post has a CW attached, then someone scrolling through their feed will not see the details of the post. This means that CWs about racism enable users to avoid encountering details of racism in the lived experiences of people of color. As Flowers argues, this enables white people to avoid being confronted the realities of racism:

[S]ome of the arguments for content warnings, or at least the majority of them, at least in my view, come down to an attempt to preserve a particular white entitlement to comfort or freedom from racialized stress through the ways in which whiteness is held center as an inheritance of the platform. (Flowers and Hendrix 2022)

This centering of white people's desire for comfort is ethically problematic, but it is also epistemically pernicious as a form of white ignorance. Charles Mills defines white ignorance as a systemic non-knowing caused by white racism or white racial domination (Mills 2007, 20). One manifestation of white ignorance is the systemic erasure of the history of white racism and the effects of racism on the lives of people of color. When people of color's testimony about their lived experiences of racism is erased, ignored, or discounted, knowledge is lost (Mills 2007).

Flowers' analysis of the centering of whiteness on Mastodon shows how social norms surrounding CWs on Mastodon can maintain a culture of white ignorance on Mastodon by hiding people of color's testimony from view.

If we look at the ways that these social norms were enforced, we find further epistemic perils of the whiteness of Mastodon. There has been a pattern on Mastodon of people of color being chastised for omitting CWs when they discussed their experiences with racism. In situations in which a person of color posted something about racism without a CW, often another user would appear in their mentions telling them that they should have hidden the post behind a CW. This is another manifestation of white entitlement that is usefully analyzed using Shannon Sullivan's concept of "ontological expansiveness" (Sullivan 2006). Sullivan explains this habit of racial privilege as follows: "As ontologically expansive, white people tend to act and think as if all spaces—whether geographical, psychical, linguistic, economic, spiritual, bodily, or otherwise—are or should be available to them to move in and out of as they wish" (Sullivan 2006, 10). Frost-Arnold applies this concept to the internet by pointing out that many white people feel entitled to enter the online conversations of people of color and hijack spaces in which people of color congregate (Frost-Arnold 2023, 180–83). When white people feel entitled to enter the mentions of people of color to chastise them for failing to use a CW when discussing their experiences with racism, they act in ontologically expansive and epistemically vicious ways. Online ontological expansiveness is an epistemically damaging habit because it can distract people of color from fruitful conversations about racism that could reduce ignorance, correct for racist biases, and share valuable truths about the nature of racism. Furthermore, ontological expansiveness is destructive of trust in the public avenues for discussion that, as we saw in section 3.1.1, is epistemically critical for the functioning of online epistemic

communities. When people of color do not trust fellow users to respect their decisions about when to use CWs, they are likely to feel disrespected as knowers and may choose to curtail their testimony or leave the platform altogether. This not only deprives the community of their insights, but also constitutes a form of epistemic injustice. When people of color truncate their testimony due to lack of trust in their white peers in ASM, this constitutes an example of the epistemic injustice of testimonial smothering (cf. Dotson 2011). Thus, the culture of whiteness on Mastodon generates a variety of epistemic harms.

4. Conclusion

I have argued that ASM platforms, such as Mastodon, present a broad range of epistemic promises and perils. It is outside the scope of the paper to make a determination about whether Mastodon, or ASM in general, are an all things considered improvement over CSM platforms. There are limitations to the Mastodon case study, and other ASM platforms also deserve the attention of epistemologists. One of the themes that emerges from this case study is that, as the sociotechnical approach predicts, both the technical affordances of ASM and the social context shape the epistemic consequences of ASM platforms. For example, we cannot simply disentangle the technical affordance CWs on Mastodon from the culture of whiteness in FLOSS communities. And the epistemic effects of the ways that Mastodon's federated infrastructure balances large networks with small communities cannot be disentangled from the political environment of right-wing extremism. I hope to have shown that ASM present rich opportunities for understanding the complexity of epistemology of social media. As both philosophers and social media users, we would do well to continue to explore new ways to produce knowledge together online.

References

- Alfano, Mark, Amir Ebrahimi Fard, J. Adam Carter, Peter Clutton, and Colin Klein. 2021. “Technologically Scaffolded Atypical Cognition: The Case of YouTube’s Recommender System.” *Synthese* 199 (1): 835–58. <https://doi.org/10.1007/s11229-020-02724-x>.
- Anderau, Glenn. 2023. “Fake News and Epistemic Flooding.” *Synthese* 202 (4): 106. <https://doi.org/10.1007/s11229-023-04336-7>.
- Angwin, Julia, and Hannes Grassegger. 2017. “Facebook’s Secret Censorship Rules Protect White Men from Hate Speech but Not Black Children.” ProPublica. June 28, 2017. <https://www.propublica.org/article/facebook-hate-speech-censorship-internal-documents-algorithms>.
- Barbarrusa, Daniel, and Lola Medina Vizueté. This issue. “Am I Still Young at 20? Online Bubbles for Epistemic Activism.” *Topoi*.
- Barnes, Michael Randall. 2022. “Online Extremism, AI, and (Human) Content Moderation.” *Feminist Philosophy Quarterly* 8 (3/4). <https://doi.org/10.5206/fpq/2022.3/4.14295>.
- Begby, Endre. 2022. “From Belief Polarization to Echo Chambers: A Rationalizing Account.” *Episteme*, June, 1–21. <https://doi.org/10.1017/epi.2022.14>.
- Berman, Mark, Jeffrey S. Chase, Lawrence Landweber, Akihiro Nakao, Max Ott, Dipankar Raychaudhuri, Robert Ricci, and Ivan Seskar. 2014. “GENI: A Federated Testbed for Innovative Network Experiments.” *Computer Networks* 61 (14): 5–23. <https://doi.org/10.1016/j.bjp.2013.12.037>.
- buckenham, v. 2022. “CW: Mastodon | V21.” April 30, 2022. <http://v21.io/blog/uses-of-cws>.
- Caelin, Derek. 2022. “Decentralized Networks vs The Trolls.” In *Fundamental Challenges to Global Peace and Security : The Future of Humanity*, edited by Hoda Mahmoudi, Michael H. Allen, and Kate Seaman, 143–68. Cham: Springer International Publishing. https://doi.org/10.1007/978-3-030-79072-1_8.
- Campbell, Douglas R. 2023. “In Defense of (Some) Online Echo Chambers.” *Ethics and Information Technology* 25 (3): 45. <https://doi.org/10.1007/s10676-023-09715-9>.
- Captain, Sean. 2022. “Can Mastodon Be a Twitter Refuge for Marginalized Groups?” Fast Company. November 30, 2022. <https://www.fastcompany.com/90817452/can-mastodon-be-a-twitter-refuge-for-marginalized-groups>.
- Coady, David. 2024. “Stop Talking about Echo Chambers and Filter Bubbles.” *Educational Theory* 74 (1): 92–107. <https://doi.org/10.1111/edth.12620>.
- Constine, Josh. 2020. “‘Content Network Effect’ Makes TikTok Tough to Copy.” TechCrunch. April 1, 2020. <https://techcrunch.com/2020/04/01/content-network-effect/>.
- Dixon, Stacy Jo. 2024. “Facebook MAU Worldwide 2023.” Statista. May 21, 2024. <https://www.statista.com/statistics/264810/number-of-monthly-active-facebook-users-worldwide/>.
- Doctorow, Cory. 2023. “The ‘Enshittification’ of TikTok.” *Wired*, January 23, 2023. <https://www.wired.com/story/tiktok-platforms-cory-doctorow/>.
- Dotson, Kristie. 2011. “Tracking Epistemic Violence, Tracking Practices of Silencing.” *Hypatia* 26 (2): 236–57. <https://doi.org/10.1111/j.1527-2001.2011.01177.x>.

- Doyle-Burke, Dylan, and Jess Smith. 2023. Twitter vs. Mastodon with Johnathan Flowers. The Radical AI Podcast. <https://www.radicalai.org/twitter-vs-mastodon>.
- Dunbar-Hester, Christina. 2020. *Hacking Diversity: The Politics of Inclusion in Open Technology Cultures*. Princeton Studies in Culture and Technology. Princeton, New Jersey: Princeton University Press.
- . 2024. “Showing Your Ass on Mastodon: Lossy Distribution, Hashtag Activism, and Public Scrutiny on Federated, Feral Social Media.” *First Monday*, March. <https://doi.org/20240313092153000>.
- Dwoskin, Elizabeth, Nitasha Tiku, and Heather Kelly. 2020. “Facebook to Start Policing Anti-Black Hate Speech More Aggressively than Anti-White Comments, Documents Show.” *Washington Post*, December 3, 2020. <https://www.washingtonpost.com/technology/2020/12/03/facebook-hate-speech/>.
- Ehmke, Coraline Ada. 2023. “Dimensions of Digital Coercion.” Organization for Ethical Source. <https://ethicalsource.dev/publications/digital-coercion/>.
- Elliott, Vittoria. 2023. “Twitter Really Is Worse Than Ever.” WIRED. May 3, 2023. <https://www.wired.com/story/twitter-really-is-worse-than-ever/>.
- “Facebook’s Civil Rights Audit - Final Report.” 2020. <https://about.fb.com/wp-content/uploads/2020/07/Civil-Rights-Audit-Final-Report.pdf>.
- Fantl, Jeremy. 2021. “Fake News vs. Echo Chambers.” *Social Epistemology* 35 (6): 645–59. <https://doi.org/10.1080/02691728.2021.1946201>.
- Flowers, Johnathan, and Justin Hendrix. 2022. The Whiteness of Mastodon. Podcast. <https://techpolicy.press/the-whiteness-of-mastodon/>.
- Forristal, Lauren. 2024. “Threads Opens Beta to Let Users Connect Their Accounts to the Fediverse.” TechCrunch. March 21, 2024. <https://techcrunch.com/2024/03/21/threads-opens-beta-to-let-users-connect-their-accounts-to-the-fediverse/>.
- Frost-Arnold, Karen. 2023. *Who Should We Be Online? A Social Epistemology for the Internet*. New York, NY: Oxford University Press.
- Furman, Katherine. 2023. “Epistemic Bunkers.” *Social Epistemology* 37 (2): 197–207. <https://doi.org/10.1080/02691728.2022.2122756>.
- Gehl, Robert W. 2015. “The Case for Alternative Social Media.” *Social Media + Society*. <https://journals.sagepub.com/doi/full/10.1177/2056305115604338>.
- Gehl, Robert W., and Diana Zulli. 2022. “The Digital Covenant: Non-Centralized Platform Governance on the Mastodon Social Network.” *Information, Communication & Society* 0 (0): 1–17. <https://doi.org/10.1080/1369118X.2022.2147400>.
- GNU Project. n.d. “What Is Free Software?” GNU Operating System. Accessed June 17, 2024. <https://www.gnu.org/philosophy/free-sw.html>.
- Goldman, Alvin I. 1999. *Knowledge in a Social World*. New York: Clarendon Press.
- Gorwa, Robert, Reuben Binns, and Christian Katzenbach. 2020. “Algorithmic Content Moderation: Technical and Political Challenges in the Automation of Platform Governance.” *Big Data & Society*, February. <https://doi.org/10.1177/2053951719897945>.
- Habgood-Coote, Joshua. 2024. “Towards a Critical Social Epistemology of Social Media.” In *Oxford Handbook of Social Epistemology*, edited by Jennifer Lackey and Aidan McGlynn. New York: Oxford University Press.
- Habgood-Coote, Joshua, Natalie Ashton, and Nadja El Kassar. 2024. “Receptive Publics.” *Ergo* 11 (5): 113–49.

- Hern, Alex. 2024. "Twitter Usage in US 'Fallen by a Fifth' since Elon Musk's Takeover." *The Guardian*, March 26, 2024. <https://www.theguardian.com/technology/2024/mar/26/twitter-usage-in-us-fallen-by-a-fifth-since-elon-musks-takeover>.
- Johnson, Deborah G. 2021. "Emerging Technology as Promise and Peril." In *The Oxford Handbook of Philosophy of Technology*, edited by Shannon Vallor, 647–62. New York: Oxford University Press. <https://doi.org/10.1093/oxfordhb/9780190851187.013.39>.
- Kaye, David. 2019. *Speech Police: The Global Struggle to Govern the Internet*. New York: Columbia Global Reports.
- Kwet, Michael. 2018. "Break the Hold of Digital Colonialism." *The Mail & Guardian*, June 29, 2018. <https://mg.co.za/article/2018-06-29-00-break-the-hold-of-digital-colonialism/>.
- . 2019. "Digital Colonialism: US Empire and the New Imperialism in the Global South." *Race & Class* 60 (4): 3–26. <https://doi.org/10.1177/0306396818823172>.
- . 2021. "To Fix Facebook, We Need to Socialise the Networks." *The Mail & Guardian*, October 30, 2021. <https://mg.co.za/opinion/2021-10-30-to-fix-facebook-we-need-to-socialise-the-networks/>.
- Lackey, Jennifer. 2021. "Echo Chambers, Fake News, and Social Epistemology." In *The Epistemology of Fake News*, edited by Sven Bernecker, Amy K. Flowerree, and Thomas Grundmann, 206–27. New York: Oxford University Press.
- Levy, Neil. 2023. "Echoes of Covid Misinformation." *Philosophical Psychology* 36 (5): 931–48. <https://doi.org/10.1080/09515089.2021.2009452>.
- Mannell, Kate, and Eden T. Smith. 2022. "Alternative Social Media and the Complexities of a More Participatory Culture: A View From Scuttlebutt." *Social Media + Society* 8 (3): 20563051221122448. <https://doi.org/10.1177/20563051221122448>.
- "Mastodon Server Covenant for Joinmastodon.Org." n.d. Accessed June 14, 2023. <https://joinmastodon.org/covenant>.
- Mbembe, Achille. 2015. "Decolonising Knowledge and the Question of the Archive." <http://wiser.wits.ac.za/system/files/Achille.pdf>.
- Medina Vizuete, Lola. Forthcoming. "What about My True Beliefs? On the Construction of Our Collective Memory Online." *Daimon: The International Journal of Philosophy*.
- Mills, Charles. 2007. "White Ignorance." In *Race and Epistemologies of Ignorance*, edited by Nancy Tuana and Shannon Sullivan, 26–31. Albany, NY: SUNY Press.
- Mitova, Veli. 2020. "Decolonising Knowledge Here and Now." *Philosophical Papers* 49 (2): 191–212. <https://doi.org/10.1080/05568641.2020.1779606>.
- Munroe, Wade. 2023. "Echo Chambers, Polarization, and 'Post-Truth': In Search of a Connection." *Philosophical Psychology* 0 (0): 1–32. <https://doi.org/10.1080/09515089.2023.2174426>.
- Nguyen, C. Thi. 2020. "Echo Chambers and Epistemic Bubbles." *Episteme* 17 (2): 141–61. <https://doi.org/10.1017/epi.2018.32>.
- . 2021. "Was It Polarization or Propaganda?" *Journal of Philosophical Research* 46 (October): 173–91. <https://doi.org/10.5840/jpr20211022183>.
- Perez, Sarah. 2023. "Mastodon Actually Has 407K+ More Monthly Users than It Thought." TechCrunch. October 9, 2023. <https://techcrunch.com/2023/10/09/mastodon-actually-has-407k-more-monthly-users-than-it-thought/>.
- Quijano, Anibal. 2007. "Coloniality and Modernity/Rationality." *Cultural Studies* 21 (2–3): 168–78. <https://doi.org/10.1080/09502380601164353>.

- Raman, Aravindh, Sagar Joglekar, Emiliano De Cristofaro, Nishanth Sastry, and Gareth Tyson. 2019. "Challenges in the Decentralised Web: The Mastodon Case." In *Proceedings of the Internet Measurement Conference*, 217–29. IMC '19. New York, NY, USA: Association for Computing Machinery. <https://doi.org/10.1145/3355369.3355572>.
- Roberts, Sarah T. 2019. *Behind the Screen: Content Moderation in the Shadows of Social Media*. New Haven: Yale University Press.
- Roose, Kevin, and Kate Conger. 2019. "YouTube to Remove Thousands of Videos Pushing Extreme Views." *The New York Times*, June 5, 2019. <https://www.nytimes.com/2019/06/05/business/youtube-remove-extremist-videos.html>.
- Rozenshtein, Alan Z. 2022. "Moderating the Fediverse: Content Moderation on Distributed Social Media." SSRN Scholarly Paper. Rochester, NY. <https://doi.org/10.2139/ssrn.4213674>.
- "Santa Clara Principles on Transparency and Accountability in Content Moderation." 2018. Santa Clara Principles. 2018. <https://santaclaraprinciples.org/>.
- Santos, Breno R. G. 2020. "Echo Chambers, Ignorance and Domination." *Social Epistemology*, November. <https://www.tandfonline.com/doi/abs/10.1080/02691728.2020.1839590>.
- Siapera, Eugenia. 2022. "AI Content Moderation, Racism and (de)Coloniality." *International Journal of Bullying Prevention* 4 (1): 55–65. <https://doi.org/10.1007/s42380-021-00105-7>.
- Silberling, Amanda, Alyssa Stringer, and Cody Corral. 2024. "What Is Bluesky? Everything to Know about the App Trying to Replace Twitter." TechCrunch. March 19, 2024. <https://techcrunch.com/2024/03/19/what-is-bluesky-everything-to-know-about-the-app-trying-to-replace-twitter/>.
- Stefanello, Viola. 2022. "What Is Up With All the Content Warnings on Mastodon?" November 10, 2022. <https://www.dailydot.com/debug/mastodon-content-warnings-twitter/>.
- Stevenson, Michael, and Carolina Valente Pinto. 2024. "Distinction and Alternative Tech: Exploring the Techno-Critical Disposition." *New Media & Society*, March, 14614448241239579. <https://doi.org/10.1177/14614448241239579>.
- Stewart, Heather, Emily Cichocki, and Carolyn McLeod. 2022. "A Perfect Storm for Epistemic Injustice: Algorithmic Targeting and Sorting on Social Media." *Feminist Philosophy Quarterly* 8 (3/4). <https://doi.org/10.5206/fpq/2022.3/4.14291>.
- Sullivan, Shannon. 2006. *Revealing Whiteness: The Unconscious Habits of Racial Privilege*. Bloomington, IN: Indiana University Press.
- Talisse, Robert B. 2021. *Overdoing Democracy: Why We Must Put Politics in Its Place*. New York: Oxford University Press.
- Valens, Ana. 2019. "Mastodon Is Crumbling—and Many Blame Its Creator." The Daily Dot. January 18, 2019. <https://www.dailydot.com/debug/mastodon-fediverse-eugen-rochko/>.
- York, Jillian. 2021. *Silicon Values: The Future of Free Speech under Surveillance Capitalism*. Brooklyn: Verso Books.
- Zuckerman, Ethan. 2020. "The Case for Digital Public Infrastructure." Knight First Amendment Institute at Columbia University. January 17, 2020. <https://knightcolumbia.org/content/the-case-for-digital-public-infrastructure>.
- Zuckerman, Ethan, and Chand Rajendra-Nicolucci. 2021. "Deplatforming Our Way to the Alt-Tech Ecosystem." *Knight First Amendment Institute at Columbia University* (blog). January 11, 2021. <http://knightcolumbia.org/blog/deplatforming-our-way-to-the-alt-tech-ecosystem>.

Zulli, Diana, Miao Liu, and Robert Gehl. 2020. "Rethinking the 'Social' in 'Social Media': Insights into Topology, Abstraction, and Scale on the Mastodon Social Network." *New Media & Society* 22 (7): 1188–1205. <https://doi.org/10.1177/1461444820912533>.