

Response to Comments

Respuesta a comentarios

KAREN FROST-ARNOLD*

I am grateful to the authors in this symposium for such thoughtful and thought-provoking analyses of *Who Should We Be Online?* Each piece has helped me better situate the book and clarify my own ideas about what I tried to achieve in this project. Beatriz Jordá's piece usefully frames the book as wrestling with optimism and pessimism about the internet as a tool to combat ignorance. On the one hand, the internet is a space of frenetic epistemic activity—so much of what we learn today is discovered via the internet. And, as Jordá points out, the internet can connect us to people very different than us. This provides opportunities for learning about social justice, unlearning our own biases and stereotypes, and developing virtuous epistemic habits of listening to people who experience oppression. Thus, there are many epistemic benefits from our online lives. On the other hand, the internet is rife with disinformation, hoaxes, and bad actors. I appreciate that Jordá recognizes both my theoretical and practical goals in the book. Theoretically, the book develops a socially situated epistemology with a set of interdisciplinary tools for analyzing the promises and perils of the internet. My goal is not to address every question about internet epistemology, but rather to demonstrate the value of these particular theoretical tools by applying them to several online personas (moderators, imposters, tricksters, fakers, and lurkers). As Jordá notes, the book also tries to help each of us navigate the challenging online landscape. We can think about how our online activities promote objectivity, truth, and epistemic justice, and we can try to cultivate habits for unlearning our biases and ignorance. Since the publication of *Who Should I Be Online?*, I have been particularly excited to see developing work by sev-

* Professor of Philosophy, Hobart & William Smith Colleges and Visiting Associate Professor, the African Centre for Epistemology and Philosophy of Science, University of Johannesburg. Her main lines of research are ethics, philosophy of science, feminist epistemology and social epistemology. Recent publications: Frost-Arnold, K. (2021). "The Epistemic Dangers of Context Collapse Online." In *Applied Epistemology*, (ed.) J. Lackey. New York: Oxford University Press; Frost-Arnold, K. (2020). "Trust and Epistemic Responsibility." In *The Routledge Handbook of Trust*, (ed.) J. Simon. New York: Routledge.
frost-arnold@hws.edu

eral philosophers using the research ethics appendix of the book to work through the tricky problems of how to conduct research about the internet in an ethical way.

Lola Medina Vizuete's piece makes an excellent contribution to the literature on online ignorance by noting a gap in my analysis of miscognition in our collective social memory. In the book, I drew on Charles Mills' concept of white ignorance to show how fake news caused by white racism and white racial domination shapes our understanding of our past (Mills 2007). In "White Ignorance," Mills explains white ignorance as *false belief* or *absence of true belief* caused by white racism. I argued that when search engine archives store racist fake news, they help maintain racist false beliefs, and this is a form of digital white ignorance. Thus, the book explores how the dissemination and storage of false claims online causes ignorance. However, as Medina Vizuete notes, I did not investigate the other type of ignorance suggested by Mills: the absence of true belief. I did not examine whether white racism prevents certain types of knowledge from circulating online or remaining in online archives. Medina Vizuete's piece brilliantly presents several ways that the absence of certain voices online might cause gaps in our collective social memory. Medina Vizuete is certainly right that minoritarian beliefs are less likely to be shared online, and I think she persuasively argues that non-propositional knowledge is also under-represented online. I will present some examples that I think further support and extend her argument.

First, Medina Vizuete argues that small minority groups may have their beliefs proportionately represented. Due to their small number, minoritarian claims may be swamped out by the majoritarian view of the collective social memory. This strikes me as an important mechanism for the creation of gaps in our collective knowledge, and significantly Medina Vizuete points out that this epistemic problem falls outside the scope of epistemic injustice—insofar as the minoritarian groups are not being unjustly discriminated against as knowers. As Medina Vizuete points out, standpoint theory shows that minority groups often have valuable insights that should be attended to by the population at large. It is important that we create online spaces in which these kinds of minoritarian knowledges can be created and heard. For me, the challenge of finding ways for the insights of minority populations to be collectively generated, disseminated, and given uptake motivates the argument for online spaces where these populations can gather to create knowledge and plan for its dissemination. This is a topic which Medina Vizuete and Barbarrusa have skillfully addressed in a paper where they argue for the value of epistemic bubbles to foster the knowledge of patients with rare diseases, such as Cystic Fibrosis (Medina Vizuete and Barbarrusa Forthcoming).

Second, Medina Vizuete points out that lack of resources, education, and access to the internet may prevent some groups from adding their beliefs to the collective social memory. This is an important point. If applied epistemologists hope to provide recommendations to improve the quality of knowledge online, then we ought to consider the need for redistribution of wealth and resources. This is particularly important when we think about calls for epistemic decolonization (cf. Mitova 2020; Tobi 2020). To undo centuries of colonial domination of knowledge production and the suppression of Indigenous knowledge requires not only the removal of colonial influences, but also the "proactive utilisation of the marginalised epistemic resources of the colonised in the advancement of knowledge in various fields" (Mitova 2020, 192). In order for these epistemic resources to be used online, colonized communities need internet access, education, and other material

resources. Thus, global redistribution of wealth and resources has epistemic merit. I think this point is also relevant to Medina Vizuete's concerns about the absence of some types of non-propositional knowledge online. While Medina Vizuete is right that knowledge-how can be resistant to testimonial transmission, it can be learned through digital visual media, such as YouTube videos (Nagel 2017). However, if minority communities lack access to digital technologies to produce and post such videos, some of their knowledge-how may not be well represented online.

Third, Medina Vizuete draws our attention to groups who choose not to participate in online conversations and thus deprive the community of their knowledge. This deprivation can also skew our collective social memory. As Medina Vizuete notes, there are many reasons why people choose not to share their knowledge online. One reason that more social epistemologists should consider is privacy. Fear of online harassment, doxing, having one's information sold to data brokers, and dislike of surveillance by advertisers and others online are all pressing privacy considerations that push people to stay off social media (Citron 2022). When applied epistemologists consider social, legal, and political suggestions to increase the diversity of knowledge online, greater protections for privacy ought to be part of the conversation, in order to encourage a wider array of minoritized groups to share their knowledge online.

Gonzalo Velasco Arias' contribution raises important questions about epistemic virtue and individualism. Velasco Arias focuses on my use of situated virtue epistemology and systems-oriented social epistemology. Bringing these two approaches together asks us to identify how the social-epistemic structure of the internet shapes (and is shaped by) users' virtues and vices. Velasco Arias finds this approach useful for avoiding both doxastic voluntarism and a kind of naïve faith in individual responsibility. This is compatible with one of the goals of the book: to argue that the way to epistemically improve the internet is not simply to expect individuals to choose to do better and act more responsibly. Structural changes are necessary, and structural changes can shape who we become as online knowers. Nonetheless, Velasco Arias finds a lingering individualism in my account of epistemic virtue, and I find his analysis very helpful. His argument begins by drawing our attention to two features of online testimony that pose evaluative challenges for users. First, people can pretend to be someone they are not online (Frost-Arnold 2014; 2023). Second, the internet can give users an illusion of understanding that encourages vicious epistemic arrogance (Levy 2019; de Ridder 2022). Velasco Arias argues that these two problems make the internet a hostile epistemic environment for agents. Thus, we should be skeptical about my argument in chapter 5 that privileged people can gain knowledge from marginalized people's online testimony. I argued there that virtuous lurkers can responsibly gain knowledge by listening to marginalized people, but Velasco Arias is pessimistic about this possibility. He asks for clarification about my notion of the virtuous lurker. Are individual agents virtuous lurkers because they possess other virtuous character traits, or does the structure of the internet shape lurkers to develop the relevant virtues? He worries that if my account rests on the former claim, then my account devolves into a voluntaristic responsibilism. But if my account takes the latter route, then it renders epistemic agents irrational by assuming they take unjustified epistemic risks by believing marginalized agents online.

This is a very insightful argument, which helped me clarify my own thoughts about nature of epistemic virtue in the book. The book recognizes the importance of structural changes to improve epistemic virtue, but it also identifies a role for individual epistemic responsibility. Thus, the chapter on lurking argues that privileged people need to choose to be more careful in spaces where marginalized people gather. Velasco Aris may view these parts of the chapter as unacceptably responsibilist. I do not eschew all forms of responsibilism. And while I reject a naïve faith that individual responsibility can solve all our epistemic problems online, I do think agents can work to develop virtuous habits. Next, I want to focus on his comments on epistemic risk, pessimism, and deference.

First, let's begin with the idea that the internet is a hostile epistemic space. I do not think philosophers are in a position to make this kind of claim today. It is certainly true that we have many theoretical tools to identify epistemic challenges of the internet. Velasco Arias successfully presents two such problems (imposters and the illusion of understanding), and my book acknowledges others. However, the internet also provides many epistemic benefits. Thus, to know whether the internet is a hostile space, we need to know whether the benefits outweigh the costs. And to answer that question, we need to both specify what we take to be the measure of epistemic value (something philosophers can do) and also to have data about what is happening online (something we need social scientists and others to do). For example, suppose we agree with my account that truth is epistemically valuable and that internet imposters often spread falsehoods. Then, to know whether the internet is an epistemically hostile space, we need to know many imposters there are, how influential are they, how many false beliefs do they propagate, and how much do they undermine trust? I do not think we have adequate data to answer these questions. Additionally, I am not sure it is the best approach to ask questions about how hostile or risky the internet is as an overall epistemic environment. The internet is an incredibly diverse set of overlapping epistemic spaces. Wikipedia has grown to become a relatively reliable epistemic tool (Frost-Arnold 2018), but Truth Social is a toxic epistemic space. And individual platforms can change over time. For example, Elon Musk has removed many of the guardrails on Twitter (now X). At this point in time, I do not think philosophers have developed the collaborations with researchers in other fields (or have developed the abilities ourselves) to answer the empirical questions we would need to make assessments about whether a particular epistemic space at a particular point in time is hostile. My book aims to provide situated theoretical tools that would help us formulate questions and develop a better understanding of the kinds of collaborations necessary to answer them in the future. I think Velasco Arias' argument here relies on an empirical claim that believing marginalized people's testimony via lurking requires excessive epistemic risk, and I am not convinced that this is true.

Finally, I found Velasco Arias' comments on deference and the need for self-confidence to be relationally grounded very interesting. However, I am not sure that analyzing the kind of epistemic acts that lurkers engage in as deference is the most helpful framing. In the chapter on lurking, the epistemic agents who I argued would benefit most from lurking are agents already beginning to unlearn their socially situated ignorance. These are privileged people who realize that they have some privilege and that they may be unaware of its scope. They suspect that they may have some false beliefs as a result of their social location, but they do not know how many or what they are. For people at this point in their journey in

unlearning their ignorance, lurking in spaces for marginalized people may be useful, but I do not think they are simply deferring to the judgment of marginalized people, in Velasco Arias' sense. Velasco Arias adopts a broad definition of deference: "A defers to B on the question whether p if A believes that p (or not-p) merely because B believes that p (or not-p)" (Brinkmann 2022, 267). For him, deference is a radical act of trust in the speaker's trustworthiness. This is not what I think is going on at this stage in the process of unlearning one's ignorance. The privileged person does not simply defer to a marginalized testifier. This is less of a matter of replacing previous beliefs with new beliefs based on testimony, and more of a matter of raising questions and doubts and reasons to look for further evidence for one's beliefs. Part of what is going on in this process is helpfully illuminated by Karen Jones' account of a metastance of distrust in one's patterns of distrust (Jones 2002). Jones argues that in the process of unlearning their ignorance, epistemically responsible agents can cultivate habits of reflecting on their patterns of distrust (Jones 2002, 166). Suppose that I, as a white person, notice that I have a pattern of distrusting people of color, and this distrust is best explained by stereotypes and prejudices. Then I ought to adopt a metastance of distrust in my distrust. In other words, I ought to distrust my tendency to distrust people of color. Instead of arguing that I ought to just defer to their judgment about certain questions, Jones argues that this metastance of distrust will push me to look for more evidence. The more I distrust my own distrust, the less weight my judgment that the speaker is untrustworthy has and the more corroborating evidence I ought to seek of the agent's trustworthiness (Jones 2002, 164–65). Thus, by developing habits of reflecting on their prejudices, epistemically responsible agents can adopt attitudes towards their own trust that put them in epistemically better positions to judge when their rejection of testimony requires more evidence.

Lurking in marginalized communities provides opportunities to find such evidence. Lurking is not simply reading one tweet by a marginalized person and deferring to their judgment. Instead, it is a persistent practice of listening. And listening to many voices of a community with a metastance of distrust towards one's own potentially biased habits can put one in a position to gather more evidence that over time will lead to a more virtuous testimonial sensibility, in Fricker's sense (Fricker 2007). Again, how epistemically useful any particular community or space is should be determined case by case. The goal of the book was not to argue that believing the testimony of marginalized people online is always warranted, but instead to draw our attention (as both individual users making decisions about where to lurk online and also as philosophers evaluating the merits of different acts of lurking) to the relevant features of the online space, structures, and habits of members of the community.

Antonio Gaitán Torres' contribution helpfully examines the notion of an epistemic community. He asks for clarification about what distinguishes epistemically good identity-based epistemic communities from bad ones. He claims that "favoring members of our group, limiting contact with members of other groups, openly discriminating against them when distributing resources, time and attention" are "paradigmatic examples" of epistemically harmful group behavior (Gaitán Torres 2024). Noting that I argue for the epistemic benefits of online communities in which members of marginalized groups have space to talk to each other, he asks what distinguishes such groups from epistemically toxic communities, such as anti-vaxxer communities or groups targeting marginalized people for online harassment

(e.g., white users who report people of color discussing racism for being anti-white). These are useful questions, and I appreciate the opportunity to expand on these points.

First, it is crucial to recognize that much social epistemological work on epistemic groups has failed to take the socially situated approach that I pursue in the book. Mainstream social epistemology has taken the generic individual knower stripped of their social location as the primary entity of analysis, thereby failing to recognize that knowers have differing degrees of power and privilege. Similarly, when the field has turned its attention to groups as the primary entity of analysis, it has also failed to recognize the diversity of groups and the epistemic significance of how much power groups of knowers have. As decades of feminist epistemology and epistemology of race have shown, when this diversity is erased, the paradigmatic generic individual knower (or paradigmatic generic group of knowers) tends to resemble privileged individuals (or groups). For example, we end up with accounts of knowers that resemble white men, and accounts of groups of knowers that resemble groups of white men. For this reason, I think we should be very careful when we draw conclusions from what have been taken to be “paradigmatic examples” of epistemic risks of groups. These are likely to be risks of epistemically harmful behavior by privileged groups, but not necessarily epistemically harmful behavior by groups of oppressed people. In fact, a socially situated epistemology can recognize what is often obscured by mainstream epistemology—that behaviors such as favoring members of our group or limiting contact with members of other groups are actually epistemically beneficial for some groups but not others. Why is this? Two sets of insights from the epistemologies of ignorance and standpoint epistemology provide answers.

First, oppressed groups live in a world that systematically denies, undermines, and erases their knowledge. The epistemologies of ignorance literature has uncovered many mechanisms by which this occurs, including white ignorance, testimonial injustice, willful hermeneutical ignorance, gaslighting, testimonial smothering, to name just a few (Mills 2007; Fricker 2007; Dotson 2011; Pohlhaus Jr. 2012; McKinnon 2017). As my book argues, these ignorance-producing practices are often operative online. Additionally, marginalized people are disproportionately targeted for online harassment and abuse. This systemic epistemic oppression marks the first key difference between the “good” epistemic community Gaitán Torres imagines (the blog for Black women discussing their experiences with racism) and his “bad” and “ugly” communities (the anti-vaxxer mom blog and the group organizing to attack oppressed groups). The former group experiences pervasive attacks on its ability to produce and share knowledge, while the latter two do not. For this reason, it is epistemically beneficial for marginalized people to engage in behaviors that mainstream epistemology has improperly labelled as paradigmatically harmful. It is epistemically helpful to withdraw from people who systematically gaslight us, interrupt us when we speak to tell us that we are not competent, who share stereotypes that we are dishonest, or whose ignorance is simply a waste of energy to constantly correct. The “bad” and “ugly” communities do not face this pervasive hostility and do not need to engage in these epistemic practices of withdrawal in order to be able to produce and share their knowledge.

The second difference between marginalized communities and the other communities Gaitán Torres considers is that standpoint theory shows that marginalized communities are more likely to produce reliable knowledge than communities about many topics. The

feminist veritism that I argue for in the book takes truth as a valuable epistemic goal. Thus, the fact that marginalized communities are likely to produce truths marks another important difference between the “good” epistemic communities and the “bad” and “ugly” communities. The insights of standpoint theory do not predict that the anti-vaxxer mom group or the group targeting oppressed people will be epistemically reliable.

Finally, I want to say why I disagree with the two criteria Gaitán Torres proposes for distinguishing between good epistemic communities and harmful ones. First, he proposes that fostering negative attitudes and emotions towards other groups is always epistemically harmful. For Gaitán Torres, the reason why negative attitudes are harmful is that they hinder interaction and debate between groups. However as I have argued, interaction and debate between groups is not always epistemically helpful. For privileged people who are often subject to socially situated ignorance, interaction with people different than them can help them unlearn their ignorance. But for marginalized people who are already more likely to have knowledge about the privileged worldview, interaction often systematically undermines and silences this knowledge. Gaitán Torres does not say exactly what he takes to be negative attitudes and emotions, but I think that many examples of what might be taken as such can be epistemically helpful. For example, believing that white supremacists hold a false ideology is a true belief, and anger at white supremacy can motivate people of color to express their knowledge of racist wrongs and also defend themselves from racist attacks (cf. Cherry 2021). Second, Gaitán Torres suggests that good epistemic communities of marginalized communities must “limit themselves to the articulations of experiences of oppression, avoiding wider political issues or debates” (Gaitán Torres 2024). I do not think this is a tenable distinction because it does not recognize the pervasiveness of oppression or how many features of reality are shaped by oppression. Much of what mainstream epistemology takes to be political issues or debates unrelated to oppression are in fact deeply shaped by oppression. For example, in the United States, it is hard to find any wider political issue or debate that is not shaped by the legacies of white supremacy. Gaitán Torres’ proposal fails to recognize the scope of the epistemic value of marginalized communities. Conversations between people facing oppression are a critical step in the consciousness-raising process that is central to the production of knowledge recognized by standpoint theory. This process involves articulating one’s experiences with oppression, sharing them with others who have similar experiences, and jointly critically reflecting on what these experiences teach us about ourselves, others, political systems, the nature of knowledge production, and many other features of reality relevant to other topics. Thus, to require that marginalized people only articulate their experiences of oppression with each other but refrain from discussing other controversial topics would not only hamper their ability to gain critical insights about many issues, but would also deprive the wider community of these insights.

References

- Brinkmann, Matthias. 2022. “In Defence of Non-Ideal Political Deference.” *Episteme* 19 (2): 264–85. <https://doi.org/10.1017/epi.2020.26>.
- Cherry, Myisha V. 2021. *The Case for Rage: Why Anger Is Essential to Anti-Racist Struggle*. New York: Oxford University Press.

- Citron, Danielle Keats. 2022. *The Fight for Privacy: Protecting Dignity, Identity, and Love in the Digital Age*. New York: W.W. Norton & Company, Inc.
- 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>.
- Fricker, Miranda. 2007. *Epistemic Injustice: Power and Ethics in Knowing*. New York: Oxford University Press.
- Frost-Arnold, Karen. 2014. "Imposters, Tricksters, and Trustworthiness as an Epistemic Virtue." *Hypatia* 29 (4): 790–807. <https://doi.org/10.1111/hypa.12107>.
- 2018. "Wikipedia." In *Routledge Handbook of Applied Epistemology*, edited by David Coady and James Chase, 28–40. New York: Routledge.
- 2023. *Who Should We Be Online? A Social Epistemology for the Internet*. New York, NY: Oxford University Press.
- Gaitán Torres, Antonio. 2024. "Epistemic Communities and Trust in Digital Contexts." *Daimon: The International Journal of Philosophy*.
- Jones, Karen. 2002. "The Politics of Credibility." In *A Mind of One's Own: Feminist Essays on Reason and Objectivity*, edited by Louise Antony and Charlotte Witt, 154–76. Boulder, CO: Westview Press.
- Levy, Neil. 2019. "Due Deference to Denialism: Explaining Ordinary People's Rejection of Established Scientific Findings." *Synthese* 196 (1): 313–27. <https://doi.org/10.1007/s11229-017-1477-x>.
- McKinnon, Rachel. 2017. "Allies Behaving Badly: Gaslighting as Epistemic Injustice." In *The Routledge Handbook of Epistemic Injustice*, edited by Ian James Kidd, José Medina, and Gaile Pohlhaus Jr., 167–74. New York: Routledge.
- Medina Vizuete, Lola, and Daniel Barbarrusa. Forthcoming. "Am I Still Young at 20? Online Bubbles for Epistemic Activism." *Topoi*.
- 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>.
- Nagel, Jennifer. 2017. "Comments on The Internet of Us: Knowing More and Understanding Less in the Age of Big Data." Presented at the APA Pacific Division Meeting, Seattle, Washington.
- Pohlhaus Jr., Gaile. 2012. "Relational Knowing and Epistemic Injustice: Toward a Theory of Willful Hermeneutical Ignorance." *Hypatia* 27 (4): 715–35. <https://doi.org/10.1111/j.1527-2001.2011.01222.x>.
- Ridder, Jeroen de. 2022. "Online Illusions of Understanding." *Social Epistemology*, 1–16. <https://doi.org/10.1080/02691728.2022.2151331>.
- Tobi, Abraham T. 2020. "Towards A Plausible Account of Epistemic Decolonisation." *Philosophical Papers* 49 (2): 253–78. <https://doi.org/10.1080/05568641.2020.1779602>.