

Discursive optimism defended

Article

Published Version

Creative Commons: Attribution 4.0 (CC-BY)

Open Access

Lepoutre, M. ORCID: <https://orcid.org/0000-0001-7573-8585>
(2023) Discursive optimism defended. *Politics, Philosophy and Economics*, 22 (3). pp. 357-374. ISSN 1741-3060 doi:
<https://doi.org/10.1177/1470594X231179665> Available at
<https://centaur.reading.ac.uk/112036/>

It is advisable to refer to the publisher's version if you intend to cite from the work. See [Guidance on citing](#).

To link to this article DOI: <http://dx.doi.org/10.1177/1470594X231179665>

Publisher: Sage

All outputs in CentAUR are protected by Intellectual Property Rights law, including copyright law. Copyright and IPR is retained by the creators or other copyright holders. Terms and conditions for use of this material are defined in the [End User Agreement](#).

www.reading.ac.uk/centaur

CentAUR

Central Archive at the University of Reading

Reading's research outputs online

Discursive optimism defended

Maxime Lepoutre 

University of Reading (Politics and International Relations), UK

Politics, Philosophy & Economics

1–18

© The Author(s) 2023



Article reuse guidelines:

sagepub.com/journals-permissions

DOI: 10.1177/1470594X231179665

journals.sagepub.com/home/ppe



Abstract

This article defends the democratic ideal of inclusive public discourse, as articulated in *Democratic Speech in Divided Times*, against the critiques offered by Billingham, Fraser, and Hannon. Specifically, it considers and responds to three core challenges. The first challenge argues, notably, that the “shared reasons” constraint should either apply everywhere or not at all, and that, if this constraint is to apply in divided circumstances, its justificatory constituency must be idealized. The second challenge contends that the resistance of hate speech and misinformation to counterspeech cannot adequately be explained by considerations of salience, and therefore cannot adequately be countered (as I suggest) by “positive” forms of counterspeech. Finally, the last challenge objects that the ideal of inclusive public discourse I defend remains, as pessimists allege, excessively idealistic.

Keywords

democracy, deliberation, public reason, anger, hate speech, misinformation, counterspeech, ignorance

Introduction

I could not hope for better critics than Billingham, Fraser, and Hannon, whose replies are unfailingly clear-minded, generous, and astute—and, though these replies are framed constructively, they raise important challenges for the ideal of inclusive public discourse defended in *Democratic Speech in Divided Times*.

Corresponding author:

Maxime Lepoutre, University of Reading (Politics and International Relations), Reading, Berkshire, UK.

Email: m.c.lepoutre@reading.ac.uk

In what follows, I aim to address the main challenges articulated by these critics. These challenges pertain to three different parts of my argument, which I will therefore revisit in turn: first, my defence of the “shared reasons” constraint on public reason-giving; second, my support for “more speech,” or “counterspeech,” as a viable instrument for responding to harmful speech (e.g., hate speech and misinformation); and, finally, my contention that the optimistic account of inclusive public speech outlined in the book is realistic, in the sense that it is appropriately sensitive to the real-world problems posed by distrust, ignorance, and fragmentation.

Sharing reasons

Billingham focuses first and foremost on my defence of the shared reasons constraint (SRC). Though *prima facie* appealing, this constraint risks excluding appeals to too many reasons when applied in divided settings. To offset this exclusiveness, the book argues that we should a) confine the SRC to the most formal and empowered arenas of public discourse (e.g., the latter stages of legislative debate; the deliberation of Supreme Court justices), and b) welcome emotionally charged narratives, including narratives of anger, in sites where the SRC doesn’t apply. At the same time, I suggest that c) we should refrain, as far as possible, from idealizing the SRC’s constituency to mitigate its exclusionary implications. Billingham exposes problems with each of these contentions.

The scope dilemma

According to Billingham (2023: 4–7), the idea of confining the SRC to the most formal arenas of public discourse encounters a dilemma. If we accept my argument for imposing this constraint in formal discursive arenas, then we are committed, for reasons of consistency, to imposing it in all stages of the discursive process (including in informal arenas populated by ordinary citizens). So *either* we should radically expand the arenas in which the SRC applies, *or* we should reject my argument for the SRC, and do away with this constraint altogether. Either way, the middle course I try to steer is untenable.

Against this challenge, I wish to suggest that it is possible to hold on to my argument for the SRC without being committed to applying it in all discursive arenas. My argument for the SRC appeals to considerations of domination. When I’m given a reason I share as a justification for an exercise of political power, this puts me in a better position to assess and contest that justification—and so, it puts me in a better position to exert control over this power, and thereby make it less dominating (Lepoutre, 2021: 21). Billingham’s key move is to point out that political power isn’t only exercised in the most formal political arenas. After all, ordinary citizens “exercise a modicum of political power when they cast their votes” (2023: 6). So shouldn’t they, too, offer shared reasons before voting, to make it easier for others to assess and contest their decision? Wouldn’t failure to do so expose other citizens to domination?

Billingham is right that political power isn’t confined to formal arenas. Consequently, the worry that being offered non-shared reasons might expose one to domination does apply, as he presses, outside of formal arenas. But we can accept this without accepting the further claim that the SRC should therefore apply “to all stages of deliberation” (2023: 7).

There are two reasons for this. The first is that domination comes in degrees. I understand domination as subjection to unchecked or uncontrolled power.¹ How much one is dominated thus depends on at least two factors: the extent to which the power one is subjected to is *uncontrolled*; and the *magnitude* of this power. Now, the power exercised by most ordinary citizens (e.g., when they vote) is typically vastly smaller in magnitude than the power exercised by legislators and Supreme Court justices. So, although there *is* a domination-related reason to apply the SRC to ordinary citizens in informal arenas, this reason is usually far less weighty.

The second point relates to the potential epistemic cost associated with the SRC. Applying this constraint only to formal arenas leaves ample room for epistemically valuable non-shared perspectives to be expressed in other arenas. But—as Billingham (2023: 13–14) helpfully notes—applying the SRC to all arenas (or to far more arenas) eliminates (or substantially reduces) this room, and therefore comes at an important epistemic cost.

Put together, these two points suggest that we can consistently accept my argument for the SRC without being committed to applying it everywhere. The domination-related reasons for applying the SRC beyond the most formal arenas are usually comparatively weak, while the epistemic reasons against doing so are comparatively strong. In many cases, therefore, the latter reasons plausibly outweigh the former.

Having said this, Billingham’s argument *does* show that determining where exactly the SRC applies is a more nuanced question than I originally recognized. Most individual voters exercise fairly little political power. But some people outside formal arenas (e.g., extremely wealthy celebrities with millions of followers, like Elon Musk) do wield significant political power. What the first horn of Billingham’s dilemma suggests, and I did not previously acknowledge, is that such powerful individuals may be subject to the SRC.² When exactly this is the case is impossible to say in the abstract. Yet, in line with the foregoing analysis, it will depend on how much power they wield, and on the extent to which their adherence to the SRC would prevent epistemically important and marginalized perspectives from being publicized.

One might still worry about the second horn of Billingham’s dilemma. Everything I’ve said so far assumes that being offered non-shared reasons in support of an exercise of power really can reduce one’s ability to contest that power—and so, really can increase the extent to which one is dominated. But maybe we should reject this claim, and jettison the SRC altogether. In support of this approach, Billingham suggests that the connection between non-shared reasons and domination is doubtful. Take, for example, “welfare policies justified by God’s special concern for the poor.” For Billingham, even if some citizens reject the reason this justification appeals to, it’s “not at all obvious that this leaves them dominated, or unable to hold the exercise of political power to account by contesting the laws” (2023: 7).

To hit home, however, this objection would need to be refined in two ways. First, recall that domination comes in degrees. Accordingly, my claim is not that someone who is offered non-shared reasons is wholly unable to contest the exercise of power that those reasons are meant to support. Rather, it’s that, comparatively speaking, they are *less* able to do so, other things being equal, than someone who is offered shared reasons. Second, it matters *who* offers the justification. As mentioned earlier, I’m focusing primarily on justifications offered by the (typically powerful) speakers in formal

arenas. So let's suppose that the non-shared consideration in question (e.g., God's special concern for the poor) is offered by the leader of a parliamentary majority during a legislative debate.

With these clarifications in place, I believe Billingham's example supports, rather than undermines, the connection between non-shared reasons and domination. There *is* something intuitively problematic—from the standpoint of contestation—about this justification, compared with a justification, by the same speaker, that deploys a shared reason (e.g., the value of freedom). As an authoritative speaker, the majority leader can contribute importantly to setting the terms of the debate. By offering a justification in terms of non-shared reasons, she, therefore, puts those who attach no normative force to the consideration that *God has special concern for the poor* in a challenging discursive position. If they attempt to respond to this justification on its own terms, this is likely to distort, or at least distract from, the considerations they find genuinely relevant. And if, alternatively, they ignore this justification, they leave unanswered what may now be the most publicly salient justification for this policy. This doesn't mean that they are unable to contest the policy at all.³ But they nonetheless seem less well-placed to contest the policy than if a shared reason had been offered. This suggests that offering a non-shared reason *does* come at a domination-related cost.

Emotional speech and emotional excesses

The second mechanism I propose for offsetting the SRC's exclusiveness appeals to the importance of narrative, and emotionally charged narrative in particular. Indeed, the epistemic value of emotion (including negative emotions like anger) is a major theme of the book. I therefore suggest that emotionally charged narratives have a key role to play, in sites where the SRC doesn't apply, in publicizing novel and marginalized perspectives. Yet Billingham worries that this proposal might be overinclusive. Even *if fitting* emotions are epistemically valuable, welcoming emotional appeals could also open the door for misdirected or unfitting emotional appeals (e.g., anger inappropriately directed at refugees). Hence, "the novel considerations added into the pool of shared reasons via emotional narrative might suffer from significant epistemic defects" (2023: 12).⁴

This is certainly a risk. But it's a risk that, on balance, is often worth taking. The first thing to emphasize is that the epistemic function performed by narratives expressing fitting emotions is critically important in divided social settings. Divided groups are often ignorant of one another's policy-relevant experiences (e.g., experiences of injustice or oppression), and may even lack the conceptual resources needed to fully grasp those experiences. In such contexts, emotional narratives have a crucial—and, the book argues, irreplaceable—role in conveying these important yet misunderstood experiences (Lepoutre, 2021: 59–69, 72–73).

Of course, one might nonetheless worry that the epistemic costs of misguided narratives can outweigh these weighty epistemic benefits. But—and this is the second point—we can mitigate these epistemic costs via counterspeech. As I discuss in Chapters 3 and 4, and will return to shortly, suitably designed counterspeech can alleviate the nefarious effects (including the nefarious epistemic effects) of hateful or otherwise misguided public interventions.

Finally, even though counterspeech cannot wholly block the epistemic costs of misguided emotional appeals, alternative discursive strategies are liable to engender similar costs. Norms of dispassionate argument, for example, can also be abused. Public speakers can (and frequently do) offer fallacious inferences under the guise of rational argument. And these fallacious inferences are often highly effective. Billingham, who anticipates this last response, counters that “fallacious inferences do not count as *properly* fulfilling the public justification requirement. Even if actual citizens are persuaded by them, we should not see them as truly providing shared reasons” (2023: 12, emphasis added). Yet I don’t find this reply comforting—at least not if we are interested, as this book is, in what norms should govern public discourse. The problem with a norm stipulating that people should only offer *proper* inferences is that, in context, which inferences count as proper is often precisely what is in dispute. So even if norms of dispassionate argument call for proper inferences, those norms remain likely to open the door to fallacious arguments—and with them, epistemically defective considerations. My point here is twofold: first, abandoning emotionally charged narrative in favour of dispassionate argument wouldn’t necessarily reduce epistemic risk; and second, just as we don’t usually take these epistemic risks to warrant shunning argumentative speech, so too we shouldn’t shun emotional narrative because it incurs such risks.

Idealization

Yet one might worry that the two strategies outlined above don’t do enough to alleviate the SRC’s exclusiveness. As Billingham notes, if “shared” means shared by all actual citizens, it’s likely that no reasons are shared in highly divided societies—and so, that no reasons will be available in spheres where the SRC applies. What’s more, Billingham continues, emotionally charged narrative seems unlikely to succeed in making new considerations shared by absolutely everyone. To overcome this problem, he suggests, we need to idealize the SRC’s constituency (2023: 7–10). On this approach, a shared reason isn’t shared by all actual citizens, but rather by all *idealized* citizens, where idealized citizens satisfy certain moral or epistemic conditions. I’ll refer to actual citizens who satisfy these idealizing conditions as “reasonable citizens.” Restricting the SRC’s constituency helps, because some reasons *are* shared by reasonable citizens. But this seems like a problem for me, because *Democratic Speech in Divided Times* expresses a deep reluctance to idealize.

I concede that a wholly non-idealized SRC is unviable, for exactly the reasons Billingham highlights. This is something I do sometimes gesture at in the book. I note, for instance, that idealization might sometimes be necessary, but “that we should only accept this solution as a last resort” (2021: 28, 30). Thus, I regard idealization, not as something that is ruled out entirely, but rather as something to be avoided “as far as possible” (2021: 45). Billingham is nevertheless right to press me on this point, first, because this is an important concession; and second, because the book is too unclear—and indeed, ambivalent—on this point.

But even with this concession, the point I want to insist on is that we should accept idealization only with great reluctance. That’s because idealization comes at a meaningful moral cost, notably because it exposes the unreasonable to domination. This cost should

be at the forefront of our thinking when we decide whether, and how much, to idealize. Specifically, it suggests that we should only idealize to the smallest extent needed to produce an adequate pool of shared considerations. It's precisely to ensure that we keep idealization as limited as possible that I explore, and give greater emphasis to, two alternative strategies for offsetting the SRC's exclusiveness. If, by contrast, we *don't* bear this cost in mind, we risk defining reasonableness in an excessively narrow way, which excludes large swathes of actual people. Nor is this a hypothetical worry: many theories of public reason do adopt quite a narrow account of reasonableness.⁵ When that happens, compliance with the SRC means that public discourse is framed in terms acceptable only to a relatively small set of "reasonable" people. Since I aim to defend inclusive democratic speech, this is an upshot I am keen to avoid.

Billingham, however, raises two doubts about the claim that idealization comes at a domination-related cost. First, he suggests that excluding the unreasonable dominates them only if we accept a non-moralized account of domination (such that whether one is dominated by some exercise of power doesn't depend on whether or not the exercise of power is morally permissible) (2023: 8).⁶ Second, Billingham questions whether, even on a non-moralized account of non-domination, the unreasonable are exposed to domination when offered "shared" reasons they don't recognize (2023: 8–9). After all, if the SRC only applies to formal arenas, aren't unreasonable citizens still free to contest such reasons, in whatever terms they please, outside these arenas?

Billingham is right that I rely on a non-moralized account of domination. Indeed, following Pettit (2008: 177) and Lovett's (2022: 55) influential accounts, I take it that uncontrolled power is *pro tanto* problematic *even when* this exercise of power is otherwise morally justifiable (and, relatedly, even when those who object to it do so on unreasonable grounds).⁷ It's beyond the scope of my discussion here to reiterate Pettit or Lovett's defence of this view, or to offer a full-fledged defence of my own. But, to appreciate the non-moralized account's intuitive appeal, consider, briefly, two scenarios (which are adapted from a case Billingham offers).⁸ In S1, I stop A from murdering B by unilaterally (i.e., without any checks or controls on my interference) coercing B. In S2, I stop A from murdering B by coercing B. But my interference, though its consequences are identical to those in S1, is constrained by a democratically enacted law, whose enactment A had a meaningful opportunity to check and contest.⁹ Even if interference is overall permissible in both cases, I take it there is *something* morally preferable about S2. This moral difference is readily explained by a non-moralized account of domination, but not a moralized account. For, on a non-moralized account, A is exposed to greater domination than B.

As for Billingham's second concern, I agree that the fact that the SRC is restricted in scope reduces the extent to which the unreasonable would be dominated as a result of idealization. But I doubt that this *eliminates* their exposure to domination. Where the idealized SRC still applies, it means that unreasonable citizens (in formal arenas, these might be unreasonable political representatives) receive, and are expected to respond in terms of, considerations whose normative force they reject. Insofar as this is the case, there is still a cost in terms of contestatory power. Although this cost doesn't suffice to rule idealization out entirely, as Billingham persuasively argues, it does yield a significant reason to keep idealization to a minimum.

Counterspeech

Hate speech and misinformation are widespread in divided democracies. A central theme of *Democratic Speech in Divided Times* is that we can and should use sophisticated forms of counterspeech to counteract these harmful forms of speech. Fraser is broadly sympathetic to this proposal. But she argues that my salience-based diagnosis of why hate speech and misinformation tend to resist counterspeech is flawed—and that, consequently, so too are some of my recommendations regarding the form counterspeech should take.¹⁰

Countering hate speech with “more speech”

The book argues, building on an analysis originally proposed by Robert Simpson (2013), that counterspeech directed at hate speech often fails for reasons of salience. The worry, in short, is that counterspeech risks reinforcing the salience of the hateful associations or stereotypes it is responding to. To alleviate this worry, I recommend “positive” forms of counterspeech, which are less about reiterating and explicitly rejecting hateful claims (this is “negative” counterspeech) and more about affirming an alternative vision of the world that is inconsistent with these hateful claims.

Fraser finds this analysis unsatisfactory. It’s not clear, she suggests, what is wrong with making a mere *association* (e.g., between Xs and parasites) salient, because associations are semantically transparent (2023: 7–8). And even if we focus on hateful *stereotypes* (e.g., the stereotype *that Xs are parasites*), my analysis faces, among other things, an extensional problem: some forms of counterspeech (notably, sarcastic counterspeech) make hateful stereotypes more salient; yet Fraser thinks that these salience-raising responses seem adequate (2023: 9).

Fraser therefore offers an alternative diagnosis of why hate speech often resists counterspeech, drawing on the “question under discussion” (QUD) framework. On this analysis, hate speech is problematic partly because it changes the QUD (e.g., to something like “Are Xs parasites?”) such that the status of targets of hate speech becomes “under discussion,” or “at issue,” in public debate. And the problem with counterspeech (including to some extent positive counterspeech) is that, even when it challenges hate speech, it continues to treat a problematic question as the QUD.

Fraser’s analysis is elegant and illuminating. Furthermore, it captures much of the spirit of my original argument (while lending it greater precision). When articulating what seems problematic about salience, I notably suggest that “when hateful propositions are salient in the public sphere, this means that the public conversation is partly *about* these hateful propositions.” Consequently, targets’ “standing is precisely *what is at issue in public debate*.” (2021: 99, emphasis added; see also 115). This language fits naturally with the QUD diagnosis. So I happily accept Fraser’s diagnosis, along with the recommendation she extracts from it: namely, that (positive) counterspeech “will more effectively undo the harms of hate speech if it occurs as not-at-issue content” (2023: 11).

But I’m not convinced that this means we should reject the salience analysis of why hate speech often resists counterspeech. The first reason for this is that the QUD analysis can plausibly be understood as an *interpretation* of the salience analysis, rather than an

alternative to it. The salience explanation is compatible with different accounts of *what* is being made salient (an association, a stereotype, etc.). What Fraser's discussion of the QUD framework shows is that we should (also) be concerned about the salience of *questions*. What seems harmful, in the cases Fraser cites, is that the question under discussion has been changed in a problematic direction (e.g., "Are Xs parasites?") *combined* with the fact that this change is salient to targets of hate speech. Indeed, it's unclear why the change in QUD would threaten targets' sense of their own standing, if this change weren't visible to them. Accordingly, Caponetto and Cepollaro (2022: 13) explicitly cite the QUD framework as a way of interpreting the salience diagnosis.

The second (and more substantive) reason why we shouldn't reject the salience analysis is that salience seems to matter over and above considerations relating to QUD. Indeed, making a hateful stereotype salient can be harmful *even* when doing so doesn't shift the QUD in a problematic direction.

To see this, consider first that there is a tight connection between the question under discussion, and whether a proposition is *at issue* in the conversation. This is why, as mentioned above, Fraser recommends that, to avoid shifting the QUD in a problematic direction, we should present counterspeech as content that is "not at issue." To illustrate this, Fraser contrasts two ways of responding to a claim like "Xs are parasites":

When we assert "Xs help support the NHS [the UK's National Health Service]", the claim about Xs is at-issue content. By contrast, when we say things like "The NHS, which heavily depends on Xs, is a treasured institution", it is not. The former risks embedding questions about Xs as "under discussion" in a way that the latter does not (2023: 11–12, underscore added)

But now consider the following instance of counterspeech: "Xs (who, incidentally, are not parasites) tend to live in Y part of town." There is something intuitively problematic about this counterspeech. Yet this isn't readily explained by the QUD framework. Here, consistent with Fraser's recommendation, the counterspeech claim *that Xs are not parasites* is part of the utterance's not-at-issue content.¹¹ The at-issue content is the claim *that Xs tend to live in Y part of town*. Hence, the QUD encouraged by this utterance is not "Are Xs parasites?" but rather "Where do Xs tend to live?" And this last QUD does not seem threatening to the dignity of Xs.

So what goes wrong with the above counterspeech? At least part of the answer, I suggest, is that it makes a hateful stereotype (i.e., *Xs are parasites*) salient.¹² This salience, in turn, is likely to be painful to Xs. It confronts them with a humiliating representation of themselves. Moreover, even if the stereotype is not "up for discussion" here and now, its salience is likely to remind them that some people do believe it. Thus—and this is the important point—the salience of a hateful stereotype seems potentially harmful even when it is not "at issue," and thus not the object of the QUD.

I want to return, finally, to the case of sarcastic counterspeech that reiterates (and so, raises the salience of) a hateful stereotype in order to ridicule it. Fraser (2023: 10–11) initially mentions sarcastic counterspeech to motivate the thought that counterspeech can be adequate *despite* making hateful stereotypes salient. But this case is also instructive for a different reason—namely, that, in practice, sarcastic counterspeech often backfires. The

ongoing controversies surrounding the French satirical magazine *Charlie Hebdo* exemplify this. Defenders of *Charlie Hebdo* often claim that its writers are sarcastically printing Islamophobic stereotypes to *ridicule* these stereotypes. Yet the magazine is nonetheless regularly accused of harming those targeted by these stereotypes through this practice (see, e.g., Fisher (2015) and Noorlander (2015)).

The foregoing discussion helps us understand this ambivalence. On the one hand, as Fraser persuasively argues, sarcastic counterspeech can (when its aims are properly understood) change the QUD in a felicitous direction (e.g., from “Are Xs parasites?” to “How could anyone be so foolish as to believe these stereotypes?”). At the same time, however, it raises the salience of vilifying stereotypes in a way liable to cause pain or distress (for the reasons outlined above) to those whom the stereotypes target. On closer inspection, then, the case of sarcastic counterspeech shows, not that salience doesn’t matter, but rather how two important dimensions of salience can come apart: the salience of *problematic questions*; and the salience of *problematic stereotypes*.

Countering misinformation with “more speech”

I also appeal to salience to explain why false beliefs induced by misinformation frequently persist after they’ve been corrected. I suggest, in short, that counterspeech can reinforce the salience of misinformation; that the more salient misinformation is, the more familiar (or “fluent”) it’s likely to feel; and that, as cognitive scientists have shown, we often find “fluent” claims more believable. Here too, I therefore recommend avoiding negative counterspeech, which reiterates and explicitly rejects the misinformation (“p is false!”), and deploying positive counterspeech instead. In the context of misinformation, positive counterspeech involves affirming an accurate representation of the issue at hand, which does not reiterate the misinformation, but entails or implies its falsity.

Fraser’s first worry with this salience-based analysis is that it risks unwarrantedly portraying people as irrational—in particular, due to their reliance on seemingly irrelevant considerations of familiarity or fluency (2023: 15). Yet this initial charge is too quick. Gigerenzer and Sturm (2012), for example, have argued that relying on considerations of fluency can be epistemically useful, and often constitutes an efficient way of arriving at accurate judgments.¹³ Of course, fluency heuristics *can* sometimes lead to false beliefs (notably, as I emphasize, in contexts where misinformation feels familiar). But so too can paradigmatically rational modes of reasoning. For instance, logical argument might consistently lead to false conclusions if, because one finds oneself in a bad epistemic environment, one starts from false premises. So even if fluency can lead us astray, there doesn’t seem to be anything inherently irrational about relying on this heuristic.

Fraser’s second worry is that my salience-based analysis seems to overlook alternative rationalistic explanations for the stickiness of misinformation. In particular, Fraser rightly points out that it can be perfectly rational to maintain a false belief, even after it has been corrected, if doing so is needed to preserve the coherence of one’s “cognitive eco-system.” “If someone has a false belief that *p* that plays an important role in their cognitive eco-system,” Fraser contends, “you can’t always make it rational to give up on that

belief just by telling them p is false [...] You will have to provide them with some alternative that can play that role” (2023: 17).

Why is this coherence-based explanation a problem for me? One reason might be that it makes my salience-based explanation redundant. Fraser occasionally hints at this worry when she refers to the two accounts as “competing” explanations (2023: 18). But I suggest that we should rather see them as *complementary* explanations, both of which have a part (among other explanations still)¹⁴ in explaining why misinformation often resists counterspeech. Indeed, there’s evidence that fluency affects people’s willingness to accept misinformation independent of considerations relating to cognitive coherence.¹⁵ So, although Fraser persuasively demonstrates that my salience-based explanation doesn’t exhaust the reasons for misinformation’s stickiness, this doesn’t show that we should abandon it.

But perhaps the problem is that, even if the two explanations are complementary, they yield different practical recommendations. The upshot of Fraser’s diagnosis is that, when correcting misinformation, we must do so in a way that “repairs” people’s disrupted cognitive eco-systems, so as to restore their coherence. But “positive counterspeech,” Fraser suggests, “provides no more reparative tools than negative counterspeech” (2023: 17). For example, to counter a vaccine conspiracy theory, positive counterspeech might involve telling someone *that vaccines are safe*. However, this might jar with that person’s pre-existing beliefs about how vaccines were made, and the fact that members of their network told them otherwise.

It’s true that we can engage in positive counterspeech without repairing people’s disrupted cognitive eco-systems. But this isn’t a problem unless engaging in positive counterspeech (as the salience explanation recommends) is *inconsistent* with engaging in reparative counterspeech (as the coherence explanation recommends). I see no reason to think they are. Fraser’s diagnosis doesn’t show that we shouldn’t affirm accurate representations that entail or imply (without explicitly reiterating) the falsity of misinformation (e.g., “vaccines are safe”). It just shows that we shouldn’t stop there. To repair people’s cognitive eco-systems, we should also explain, for instance, why this claim is correct (particularly if people have false and worrying pre-existing beliefs about how vaccines were made) and why members of their networks may have told them differently. This seems wholly consistent with positive counterspeech. Indeed, and as I have discussed elsewhere (Lepoutre, 2022: 8), we can explain why vaccines are safe, and why people have been fed falsehoods, without reiterating those falsehoods.¹⁶ On balance, then, Fraser’s appeal to reparative counterspeech doesn’t seem like an objection to positive counterspeech. Rather, it helps us discriminate between better and worse forms of positive counterspeech—where better forms make a concerted effort to tell a positive story that leaves people with a coherent cognitive eco-system.

However, this response leads straight to Fraser’s final concern. If the stickiness of misinformation has multiple interacting causes, which yield different (albeit compatible) recommendations for counterspeech, designing and deploying successful forms of counterspeech on a large scale seems extremely challenging. “And so,” Fraser concludes, “pessimism about counterspeech may be forced on us after all” (2023: 18). I agree that designing and implementing sophisticated forms of counterspeech is an arduous task. But it doesn’t follow that we should give up on it: first, because (as Chapters 3 and 4

also suggest) alternatives to counterspeech are often similarly arduous to implement; second, because even if eliminating speech-based harms through counterspeech is unrealistic, considerably improving on current counterspeech seems feasible; and finally, because the problems posed by hate speech and misinformation are immensely weighty, and therefore worth our time and effort, imperfect though the result may be.

Between pessimism and idealism

Part II of the book assesses the realism of the ideal of inclusive public discourse articulated in Part I. It does so, specifically, by considering whether this ideal can withstand challenges generated by political ignorance, mutual dislike, and social and spatial fragmentation. Hannon, as well as Billingham, raise three main concerns about my arguments in this Part: namely, that some of my arguments are actually irrelevant to inclusive public discourse; that some of my arguments are too idealistic, because they presuppose conditions that are substantially better than real-world conditions; and finally, that some of my optimistic responses are insufficiently supported by the evidence.

The problem of relevance

One might worry that political ignorance makes people incapable of engaging competently in public debate (see, e.g., Somin, 2016: ch. 2). To address this worry, I appeal to the possibility of using information shortcuts (I highlight, notably, the use of social group identity as a heuristic for determining whom or what to believe). For Hannon, however, defending information shortcuts isn't the same thing as defending inclusive public discourse. Accordingly, my defence of (some) information shortcuts might seem *irrelevant*: even if information shortcuts help people circumvent their political ignorance, this doesn't show that inclusive public discourse is fine. In fact, Hannon adds, this defence might actually backfire: it implies that inclusive public discourse is unnecessary, since people can get the guidance they need simply by deferring to their (group-based) representatives or opinion leaders (2023: 5–7).

My first response is that information shortcuts typically operate via public discourse. Information shortcuts often take the form of testimony (e.g., advice and recommendations of representatives and opinion leaders from one's group). And the practice of giving and receiving testimony is part and parcel of democratic public discourse. So, showing that information shortcuts enable people to circumvent political ignorance *doesn't* show that such discourse is unnecessary. On the contrary, it's relevant to demonstrating the value of democratic public discourse in conditions of political ignorance.

Hannon adeptly anticipates this reply, however, and objects that this picture of public discourse seems a far cry from the ideal of *inclusive* public discourse I meant to defend (2023: 6). After all, the practice of deferring to others' testimony is wholly consistent with the "elite" theories of democracy I reject in Chapter 8.

This is true, and Hannon is therefore absolutely right to put pressure on this first response. But I am by no means suggesting that deferring to testimony is *all* that ordinary people should do. Much of my discussion of political ignorance, in the book, focuses on (group-related) information shortcuts, and how they can constitute a useful discursive

strategy for circumventing political ignorance. Yet this focus sometimes overshadows a further and more basic point, which is that many of the forms of speech explored in Part I simply *aren't* undermined by existing evidence of political ignorance.

For example, Part I emphasizes the importance of emotionally charged narratives aimed at disclosing politically relevant experiences (e.g., experiences of oppression). Now, political ignorance tests don't normally assess such experiential knowledge. So existing evidence of political ignorance doesn't show that people cannot fruitfully share their experiences, and thus perform the crucial epistemic function outlined earlier.¹⁷

A similar observation applies to counterspeech directed at hate speech. Many of the falsehoods promulgated by hate speech (e.g., representations of people as subhuman) are not the sorts of claims that political ignorance tests normally ask about—and so, they aren't the sorts of propositions we have evidence of most people accepting. Moreover, not all hate speech operates by saying things that are false. Hate speech can reject people's basic standing by saying things that are true ("We want Xs out" spoken by the leader of a white supremacist group) or things that aren't truth-apt ("Get out of this country!"). In this latter case, since the point of counterspeech isn't to correct a falsehood, one's ability to respond to hate speech doesn't clearly depend on one's political knowledge or ignorance. Thus, one might be ignorant about the different branches of US government, or about which state Ted Cruz hails from, and still perfectly capable of reassuring targets of hate speech ("I stand by you.").

My response to the problem of political ignorance is therefore twofold: first, democratic public discourse supplies tools (notably, information shortcuts) for circumventing political ignorance; and second, even when ordinary people are politically ignorant in the ways evidenced by public opinion research, they often remain capable of performing many of the discursive roles outlined in Part I. This last point marks the distinction between elite theories of public discourse—in which elites speak and the public simply listen—and inclusive public discourse—where people from different walks of life participate robustly in communication on matters of public concern.

The problem of idealism

Hannon also worries that the optimistic story I tell about inclusive public discourse depends on conditions that are substantially better than actual conditions. If so, I might be guilty of out-of-touch idealism despite my intended aim of providing a realistic account of inclusive public discourse.

This problem is most salient in the context of my discussion of group cognition. The book argues that group identity's well-documented influence on political judgment (group cognition) can be epistemically valuable. And it suggests that this can be true *even* when group identity influences judgments regarding politically relevant technical matters (technical group cognition). But this defence of technical group cognition might seem idealistic. As Hannon notes, the conditions (which I outline in the book) under which technical group cognition is likely to be epistemically reliable are quite restrictive, and unlikely to be met in practice. Hence, in practice, technical group cognition is likely to lead people to behave dogmatically, by being "too resistant to some technical evidence and insufficiently resistant to others" (2023: 14).

Now, the book does have a response to this worry, which is that dogmatic manifestations of (technical) group cognition are not necessarily bad. Dogmatism, I suggest, can play a valuable epistemic function when considered from a broader systemic perspective. But this response *itself* presupposes better-than-actual conditions. For example, it requires that information circulate readily between different dogmatic groups, which seem unlikely given the extent of intergroup social and spatial fragmentation in divided democracies (Hannon, 2023: 15–16). Nor does the problem end there. My response to the problem of (lack of) goodwill also partly depends on information circulating between groups, which—to reiterate—is made unlikely by intergroup fragmentation. Now, as Hannon acknowledges, I am aware of this, and therefore spend a full chapter arguing for integrative policies aimed at tackling intergroup fragmentation (Hannon, 2023: 16–17). But this doesn't change the fact that my defence of inclusive public discourse seems to depend on conditions that aren't *currently* in place.

The first thing to say here is that, in adopting a realistic approach, my stated aim is not to say that there are no problems with existing public discourse. It's rather to develop an ideal of inclusive public discourse that is relevant to, and useful in, non-ideal conditions (Lepoutre, 2021: 3). This clarification matters, because it's acceptable for there to be some distance between an ideal and reality. Ideals, after all, are meant to guide action, and guiding action wouldn't be needed if everything were already as it should be. So it's not enough for critics to show that my optimistic account depends on better-than-actual conditions. They need to show that these conditions are too demanding, such that the ideal is *too* distant from reality.

Is the ideal I offer too distant from reality? There are several reasons to think it isn't. For one thing, the normative ideal I propose is compatible with (and often presupposes) the existence of many characteristic features of real-world divided democracies—including expressions of anger, hate speech, misinformation, significant intergroup trust, and group cognition. And although the ideal does require greater-than-actual social and spatial integration, I contend (in my discussion of “networked integration”) that it's nonetheless compatible with a fair amount of intergroup fragmentation (2021: 202–06).

A second consideration relates to action-guidance. If an ideal is too distant from reality, it risks losing its capacity to guide action in non-ideal conditions. This is because it might be exceedingly difficult to know what we should do to get from where we are now, to where we should be. But the ideal I've offered does provide guidance here and now. Some of this guidance relates, of course, to integration. As the book stresses, the fact that some of my optimistic arguments (e.g., about dogmatic group cognition) depend on greater-than-actual integration indicates that we should adopt extensive policies aimed at reducing online and offline fragmentation (2021: ch. 7). But the ideal also provides guidance regarding discursive strategy. Many of the discursive tools I highlight retain some of their value even prior to implementing extensive integrative policies. Take, for example, my defence of positive counterspeech. Existing social and spatial fragmentation may slow down the spread of counterspeech across groups. But this doesn't mean that we shouldn't currently deploy counterspeech against hate speech and misinformation. Nor does it mean that we shouldn't currently prefer, for the reasons the book outlines, positive forms of counterspeech to negative ones. Likewise, with my defence of emotional narrative. Integration would help these narratives circulate faster and more effectively. But this doesn't mean that their epistemic value is *nonexistent* in the meantime.

A final reason for thinking that the ideal I offer isn't too far from reality is comparative. How far is *too far* depends to some degree on what alternatives are available. If alternative political ideals offered a normatively appealing picture that somehow didn't rely on substantially better-than-actual conditions, we might still conclude that, as a comparative matter, the ideal I've offered is excessively idealistic. But I argue, in the book's final chapter, that this isn't the case. Conditions of ignorance, distrust, and fragmentation are *also* problematic for the main competing ideals (e.g., elite democracy, epistocracy, and lottocracy). If I am right about this, then their desirability does seem to depend on significantly better-than-actual conditions.

The upshot is this. Even though the optimistic ideal I offer does, as Hannon persuasively demonstrates, presuppose some better-than-actual conditions (notably, relating to fragmentation), we have good reason to resist the conclusion that it is therefore insufficiently realistic. This is because it nevertheless retains many characteristic features of non-ideal democracies; because it seems capable of providing guidance here and now; and because, despite appearances to the contrary, it isn't clearly more idealistic than its principal competitors.

The problem of evidence

But even if the optimistic story I tell about inclusive public discourse isn't too distant from reality, you might still worry that this story—or at least some part of it—is insufficiently supported by the evidence. Both Billingham (2023) and Hannon (2023) raise versions of this worry. Billingham, in particular, expresses concerns about my use of the “systemic” approach to democratic public discourse. At various points in the book, I suggest that some features of public discourse are less problematic than they initially appear when considered from a systemic perspective. For example, as discussed above, in response to concerns about dogmatic group cognition, I suggest that, when information circulates readily between groups across the broader discursive system, dogmatism can actually play a positive epistemic function—notably by facilitating the emergence of new and varied evidence. Likewise, in response to the worry that expressing anger can provoke distrust of the angry speaker, I observe that, even if this is true, it can also help other (non-angry) speakers in the discursive system *gain* trust, for example by helping them appear comparatively more approachable or moderate.

The problem, Billingham notes, is that this all seems too speculative. For each of these “optimistic stories” about the discursive system, he suggests, we could tell an equally plausible “pessimistic story” about what might happen in these circumstances. For instance, perhaps group-based dogma is so inflexible that, even if information circulates readily between groups, it *won't* spur the generation of new and varied evidence. And maybe expressing anger in support of a particular cause will lead non-angry speakers who support the same cause to be treated as guilty by association—and thus, to *lose*, rather than to gain, trust (Billingham, 2023: 14–18; see, relatedly, Hannon, 2023: 14–18).

It's true that existing empirical evidence does not speak decisively in favour of the optimistic stories Billingham highlights (a point I acknowledge in the book). But articulating these stories nonetheless helps achieve two important goals. First and foremost, doing so performs a defensive function, by helping to block common pessimistic

inferences. In each of the cases Billingham cites, I'm responding to an objection that starts from a seemingly troubling observation (e.g., anger's apparent erosion of trust; group cognition's tendency to be dogmatic), and, from this observation, draws a pessimistic conclusion about democratic public discourse. In this dialectical context, the purpose of offering an optimistic story is to show that these pessimistic inferences are too quick. The initial observation can be reconciled with a plausible optimistic hypothesis. And so, from the initial observation, the pessimistic conclusion does not immediately follow. Even if group cognition can be dogmatic, for instance, this doesn't necessarily mean that it's epistemically bad for democratic public discourse. And even if anger can breed distrust locally, it doesn't follow that it's all-things-considered bad for trust.¹⁸

This defensive function matters. It may not (as Billingham correctly observes) allow us to *vindicate* the optimistic hypotheses I offer. But it nonetheless shows that the evidence we have *also* does not vindicate the pessimistic claim I'm responding to. In a context where many are inclined to despair about the ideal of inclusive public discourse, this achievement—staving off pessimism—constitutes an important first step. It takes the ideal of inclusive public discourse away from the realm of utopia or wishful thinking, and back into the realm of things we can rationally and realistically hope for.

The second goal follows from the first. The optimistic stories Billingham points to don't simply show that existing evidence is insufficient to vindicate pessimistic claims about inclusive public discourse. In addition, they also help *direct* the search for empirical evidence. In other words, they aim to provide us with empirically and philosophically informed hypotheses (about anger's effect on trust, about the epistemic function of dogma) that provide guidance for ongoing empirical research into the effects of different discursive mechanisms. As Billingham recognizes, I take this function—of hypothesis generation—to be a key function of political philosophy (2023: 17–18). And it's one that I therefore return to time and again in *Democratic Speech in Divided Times*.

If what I'm offering are informed hypotheses, rather than sure-fire solutions, does this mean that they can't or shouldn't guide action? I don't think so. Uncertainty is an inevitable feature of action under non-ideal conditions. We face uncertainty when we tackle climate change, when we implement nationwide lockdowns, when we decide whether to take sides in a militarized conflict, and so on. How much uncertainty we can bear depends on the costs of error. We'd better have excellent evidence for our hypothesis if, say, the cost of error is a nuclear war. In the cases, I'm concerned with, however, not only are the consequences of error obviously not as extreme, but there can also be important risks involved in *not* acting. For example, refraining from expressing anger, because of what might turn out to be ill-founded fears about its impact on trust, may mean foregoing a critical tool for publicizing overlooked injustices. On balance, therefore, one might rationally decide that it's worth taking the risk. Of course, not all cases will be like this. But the main point remains: the fact that something is a working hypothesis doesn't necessarily mean that we aren't warranted in acting on it.

All of this does mean, admittedly, that there is something contingent about my defence of inclusive public speech. Billingham is keenly aware of this: "Lepoutre's arguments," he concludes, "often leave me with a nagging doubt: do things really work like that? Will we really see this positive effect?" (2023: 18) I don't think I can dispel these lingering doubts. In fact, I don't think any normative account that is seriously committed to

looking at democracies as they are, in the real world, with its complexity and contingency, can do that. But the account I offer in *Democratic Speech in Divided Times* nonetheless *can*, I hope to have shown, stave off rising pessimism about inclusive public discourse. It *can* give us rational grounds for hope. And it *can*, despite uncertainty, provide guidance for both evidence-gathering and action.

Acknowledgements

For helpful comments and discussion, I am deeply grateful to Samuel Bagg, Jacob Barrett, Paul Billingham, Bianca Cepollaro, Joanna Demaree-Cotton, Jamie Draper, Rachel Fraser, Michael Hannon, Jeff Howard, Henrik Kugelberg, Ruairi Maguire, Zeynep Pamuk, and Jennifer Saul.

Declaration of conflicting interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: The author grateful acknowledges support from a Leverhulme Research Fellowship (RF-2022-001).

ORCID iD

Maxime Lepoutre  <https://orcid.org/0000-0001-7573-8585>

Notes

1. Here, I am following Pettit (2008) and Lovett (2022).
2. For a recent argument to this effect, see Kugelberg (2021b, 472).
3. Kugelberg (2021a, 13) notes, in a similar vein, that citizens who are offered non-public reasons are “*disempowered* by being forced to debate and engage with others on unfamiliar ground” (emphasis added).
4. I also discuss this worry in Lepoutre (2021: 46–47, 69–71).
5. For a defence of this point, see Enoch (2015: 121–22). Some recent accounts, such as Kugelberg’s, are much more sensitive to this problem. But even Kugelberg’s proposed “moderate idealization” ends up being quite demanding: idealized citizens, for instance, don’t make faulty inferences (2021a, 9). See Billingham and Taylor (2022) for an up-to-date overview.
6. See List and Valentini (2016: 1046) on the distinction between moralized and non-moralized freedom.
7. List and Valentini (2016) offer an alternative account of non-moralized freedom (freedom as independence) which they deem superior to Pettit’s. My argument doesn’t hinge on whether their critique is successful: arguments analogous to my own could be developed using the freedom-as-independence framework.

8. Billingham's (2023: 8) version involves a murderer who is stopped *as a result of a democratically enacted law*. But this case is less clean than we might like, since, even on a non-moralized account of domination, the democratic process constitutes a form of control on the interference.
9. For similar intuitive appeals, see List and Valentini (2016: 1058).
10. Fraser also raises doubts regarding my characterisation of hate speech and (correspondingly) what makes it harmful. I focus on the dignitarian harm in hate speech. For Fraser, this focus is problematic, first, because it seems insufficiently sensitive to hate speech's role in maintaining hierarchy; and second, because depending on how we understand dignity, the concept seems either too narrow or too broad (2023: 5–6). These concerns are slightly too quick. Following Waldron, I understand dignity as a social status: the status of being a member of society in good standing (2012: 59–60). To say that hate speech assaults its targets' assurance of their dignity, then, is to say that it assaults their assurance of their standing in society. This seems entirely consonant with attending to hate speech's role in maintaining hierarchy (which I take to be a matter of relative standing). Fraser objects that this understanding of "dignity" risks being too broad, as "most people do not have a 'basic entitlement' to be regarded as members of most societies" (2023: 7). I agree with the claim that people don't have a basic entitlement to be regarded as members of most societies. But this is consistent with thinking that most people have a basic entitlement to be regarded as members of *some* societies (perhaps because their livelihood, or ability to live a life free of domination, depends on it), and that hate speech, when it is harmful, is harmful because it assaults its targets' standing in those societies.
11. Following Fraser, and standard practice, I use an apposite relative clause to present *Xs are not parasites* as not-at-issue content. Consistent with a common test for identifying not-at-issue content, it doesn't seem possible to directly deny this proposition (e.g., by simply responding "No" or "That's not true" to the above utterance). See Tonhauser (2012, §2.1) for discussion.
12. Fraser expresses doubts about the claim that saying not-p (where p is a proposition) can reinforce p's salience (2023: 8). But in her subsequent discussion of the "loose label" theory, she acknowledges that claims that we initially encode as not-p can subsequently lose their negation "tags." This is one reason why saying not-p could increase p's salience.
13. Fraser (2023: 15n12) briefly acknowledges this possibility.
14. Two others, which the book discusses, are distrust and group-based motivated reasoning (2021: 120).
15. See, e.g., Berinsky's (2017: 257) finding that people made to repeat misinformation, thus reinforcing its fluency, are less likely to reject it. Lewandowsky et al. (2012) also sees them as complementary explanations.
16. Fraser's (2023: 18) parenthetical example also suggests a way of doing this.
17. Hannon (2023: 7) justifiably worries that elite manipulation can warp people's interpretation of their own experiences. But this is not a decisive point against the value of emotionally charged narrative. For even when elites successfully corrupt people's *interpretation* of their own experiences, this doesn't necessarily mean that people can no longer accurately relay those experiences. For example (to use one of Hannon's cases), someone who experiences rural hardship might be misled into interpreting this hardship as the result of an unfair distribution of tax money. But they can still relay valuable experiences of rural hardship. For a demonstration of this point, see Cramer (2016).
18. Billingham (2023: 19) acknowledges this towards the end of his commentary.

References

- Berinsky A (2017) Rumors and health care reform. *British Journal of Political Science* 47: 241–262.
- Billingham P (2023) Sharing reasons and emotions in a non-ideal discursive system. *Politics, Philosophy & Economics*, Online First: 1–21.
- Billingham P and Taylor A (2022) A framework for analyzing public reason theories. *European Journal of Political Theory* 21: 671–691.
- Caponetto L and Cepollaro B (2022) Bending as counterspeech. *Ethical Theory and Moral Practice*, Online First: 1–17.
- Cramer K (2016) *The Politics of Resentment*. Chicago: Chicago University Press.
- Enoch D (2015) Against public reason. In: Sobel D, Vallentyne P, and Wall S (eds) *Oxford Studies in Political Philosophy*. Oxford: Oxford University Press, Vol. 1, 112–140.
- Fisher M (2015) What everyone gets wrong about Charlie Hebdo and racism. *Vox*. Available at: <https://www.vox.com/2015/1/12/7518349/charlie-hebdo-racist>.
- Fraser R (2023) How to talk back. *Politics, Philosophy & Economics*, Online First: 1–21.
- Gigerenzer G and Sturm T (2012) How (far) can rationality be naturalized? *Synthese* 187: 243–268.
- Hannon M (2023) Public discourse and its problems. *Politics, Philosophy & Economics*, Online First: 1–21.
- Kugelberg H (2021a) Civic equality as a democratic basis for public reason. *Critical Review of International Social and Political Philosophy*, Online First: 1–23.
- Kugelberg H (2021b) Public justification versus public deliberation: the case for reconciliation. *Canadian Journal of Philosophy* 51: 468–473.
- Lepoutre M (2021) *Democratic Speech in Divided Times*. Oxford: Oxford University Press.
- Lepoutre M (2022) Narrative counterspeech. *Political Studies*, Online First: 1–20.
- Lewandowsky S, Ecker UKH, Seifert CM, et al. (2012) Misinformation and its correction. *Psychological Science in the Public Interest* 13(3): 106–131.
- List C and Valentini L (2016) Freedom as independence. *Ethics* 126: 1043–1074.
- Lovett F (2022) *The Well-Ordered Republic*. Oxford: Oxford University Press.
- Noorlander P (2015) *When satire incites hatred: Charlie Hebdo and the freedom of expression debate*. EUI Centre for Media Pluralism and Media Freedom. Available at: <https://cmpf.eui.eu/when-satire-incites-hatred-charlie-hebdo-and-the-freedom-of-expression-debate/>.
- Pettit P (2008) Republicanism: Three Axioms, Four Theorems. In: Laborde C and Maynor J (eds) *Republicanism and Political Theory*. Malden: Polity Press, 102–132.
- Simpson R (2013) Un-Ringing the bell. *Australasian Journal of Philosophy* 91(3): 555–575.
- Somin I (2016) *Democracy and Political Ignorance*. Redwood: Stanford University Press.
- Tonhauser J (2012) Diagnosing (not-) at-issue content. *Proceedings of Semantics of Under-Represented Languages in the Americas* 6: 239–254.
- Waldron J (2012) *The Harm in Hate Speech*. Cambridge: Harvard University Press.

Author biography

Maxime Lepoutre is a Lecturer in Political Theory at the University of Reading. His research focuses primarily on the norms that should govern democratic public discourse in divided settings.