

ARTICLE

Zetetic intransigence and democratic participation

Hrishikesh Joshi

Department of Philosophy, University of Arizona, Tucson, AZ, USA
Email: hrishikeshjoshi@arizona.edu

(Received 24 November 2023; revised 23 July 2024; accepted 15 August 2024)

Abstract

A pervasive feature of democracy is disagreement. In general, when we encounter disagreement from someone who is at least more reliable than chance, this puts some pressure on us to moderate our beliefs. But this raises the specter of asymmetric compliance – it's not obvious what to do when we moderate our beliefs but the other party refuses to do so. Whereas an elegant solution is available when it comes to how we can *respond* to our higher-order evidence while still preserving democratic fairness, I argue that no such solution is forthcoming when we consider *zetetic intransigence* – that is, asymmetric compliance with regard to how we inquire, or gather evidence. The upshot is that democratic participation can involve messy trade-offs with no entirely satisfactory resolution.

Keywords: Political disagreement; ethics of voting; inquiry; democracy; epistemic bubbles

1. Democracy and disagreement

Democratic participation brings with it a range of epistemic challenges. Among these is the pervasiveness of disagreement. How should we respond when people disagree with us? In general, the fact that others disagree with us puts *some* pressure on our own beliefs.¹ If I come to some judgment p about a particular issue, based on my own analysis, and then learn that you – someone who I antecedently believe is no less careful, informed, or capable than me – come to the opposing judgment, $\sim p$, it behooves me to become less confident about p than I was prior to learning about the disagreement.

Now, strict epistemic peerhood is rare (King 2012). Importantly, in the literature on peer disagreement, epistemic peers are stipulated to have the same evidence and be just as good at assessing that evidence. It is presumably rare for us to encounter such individuals, especially given that people have diverse sets of evidence, expertise, and so on. You and I need *not* be epistemic peers though, for some moderation, post-disagreement, to be called for. So long as I think you're more reliable than chance about matters like p , there is *some* normative pressure on me to moderate.

¹See Christensen (2007) and Elga (2007) for classic discussions of the idea that disagreeing peers ought to update their credences significantly toward each other's position. On Kelly's (2010) total evidence view also, some moderation is called for because disagreement with others is part of one's total evidence.

© The Author(s), 2025. Published by Cambridge University Press. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (<https://creativecommons.org/licenses/by/4.0/>), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.

Thus, for example, suppose you correctly believe that you are 90% reliable about some subject matter, while I am only 60% reliable. In most of the cases where we disagree, I will be wrong, and you will be right. Nonetheless, given that I'm better than chance, upon noticing this disagreement, it behooves you to reduce your credence from 90% – in this case, you should become 86% confident that you are right.² If I am as good as chance, your credence need not budge, and if I'm worse than chance, your credence should go up.

Now, it is reasonable to think that many of those who disagree with us (though, of course, not all) within the democratic setting are at least more reliable than chance when it comes to the sorts of matters that form the basis for disagreement. This need not mean that we think the other side is equally reliable – the preceding is consistent with thinking that we are much more reliable than they are. However, holding that the opponents are *worse than chance* would have some counterintuitive implications. For one, it would mean that on some general political matter on which we have some antecedent view p , if we then observe widespread agreement, this should make us *less* confident about p . In other words, suppose some portion of the population is anti-reliable with respect to p . Now, we notice first that our group believes p . But, on this issue, it turns out that the opposing group *also* believes p – they agree with our group. (For instance, say everyone agrees that the airport parking lot needs an expansion.) If they're really anti-reliable, this means that upon learning the latter fact, viz., that they agree with us, we should reduce our confidence in p .

That said, I don't want to claim that it's never the case that one's political opponents are less reliable than chance on policy matters.³ However, for present purposes, let us stipulate that neither side is less reliable than chance because it is here that a particular philosophically interesting puzzle arises.⁴

Granting this assumption, we face some pressure to moderate our stances on issues where we observe disagreement. That is, even if my analysis of my first-order evidence supports believing p with a particular level of confidence, upon learning about disagreement from you, it behooves me to reduce that confidence. Disagreement provides higher-order evidence in this way.

Importantly, this need not mean that I must split the difference, as far as credences go, or suspend judgment if we hold opposing beliefs initially. Rather, the core problem addressed in this paper gets going on a much weaker and more plausible assumption – namely that there is normative pressure to moderate *at least somewhat* when faced with disagreement from others who are more reliable than chance on matters like the one in question. Even on total evidence views of disagreement (Kelly 2010), higher-order evidence does count as part of one's total evidence – thus producing some pressure to reduce confidence.

²See White (2009) and Worsnip (2023) for more detailed discussion on this point.

³Indeed, some authors argue that voters are systematically biased in favor of counterproductive economic policies – in part because economic analysis is counterintuitive. See, for example, Caplan (2008). Others have noted that many voters have low political knowledge and even the relatively knowledgeable tend to vote based on party loyalty rather than issues (Achen and Bartels 2016). However, the puzzle of this paper gets going even if you think there is one domain such that one's opponents are not anti-reliable. That is enough to raise the problem of asymmetric compliance with epistemically conciliatory norms.

⁴Strictly speaking, what matters for how we should update our beliefs is whether the other side is less reliable than chance *from one's epistemic perspective*. Suppose two epistemic agents A and B are both less reliable than chance with respect to some domain. It may nevertheless be the case that from A's perspective, B is more reliable than chance. That is, given the evidence A has, they ought to believe that B is better than chance. For example, suppose A and B are medieval doctors trying to cure a patient from the plague. Even if they are both less reliable than chance when it comes to finding the cure, it can be true that from each doctor's perspective, they are justified in thinking they're both more reliable than chance. Disagreement, then, should lead them to moderate somewhat.

In a recent article, Alex Worsnip (2023) builds on the above assumptions to characterize a normative dilemma for the ethics and epistemology of political participation. The dilemma is this: on the one hand, we ought to take appropriate account of the higher-order evidence provided by disagreement. On the other hand, if our moderation is not reciprocated by the other side, it risks a political shift toward that side for the very reason that they are uncompromising. This outcome seems unfair in an important sense – for Worsnip, it violates an ideal he calls “representative fairness.” Worsnip’s way out of this dilemma is to propose the notion of a “personal take,” which he suggests ought to guide political participation. The basic idea is that people ought to *participate politically* according to their first-order evidence (i.e., bracketing their higher-order evidence). At the same time, they ought to *believe* by incorporating their higher-order evidence.

In this paper, I argue that even if this move is successful when it concerns how we ought to *respond* to evidence, the problem reappears when we consider how we ought to *inquire*. This latter problem, I argue, is more significant. That is because, empirical evidence suggests that a pervasive feature of modern partisanship is that partisans consume news, opinion, analyses, and so on, from different sources. Thus, partisans differ in terms of what evidence they have access to. This, I suggest, gives rise to a zetetic requirement to gather more evidence, for instance, by diversifying one’s information sources. However, there is now an analogous venue for asymmetric compliance: *zetetic intransigence*. That is, we can imagine a situation where our side tries to understand why our opponents believe what they do, but our opponents refuse to reciprocate. Unlike the former case, the personal take solution offers no way out of the analogous dilemma here. If Worsnip’s development of the original dilemma is on the right track, then democratic participation, especially in a context where the different sides have access to different sets of information, brings with it messy trade-offs that admit no easy solution.

2. The problem of asymmetric compliance

Suppose that two parties disagree about some issue. This provides each party with some higher-order evidence. However, suppose one party is dogmatically unresponsive to the higher-order evidence, while the other conciliates. The intransigent party’s dogmatism, irrational though it may be, can be rewarded by the principle of conciliation, under certain assumptions (Aikin et al. 2010). In the case of democratic participation, the problem is not merely epistemic. The intransigent can, owing precisely to their intransigence, secure outcomes that favor their position in a way that seems intuitively unfair.

Thus suppose, to use Worsnip’s (2023) terminology, there are two groups, the Twisters and the Stickers. The Twisters respond to their higher-order evidence appropriately and moderate their confidence when they observe disagreement with the Stickers. The Stickers, however, are intransigent – they stick to their own takes based on their first-order evidence; they are epistemically insensitive to disagreement with the Twisters.

This puts the Twisters in a dilemma. On the one hand, they can participate politically (to save words, I will talk primarily about voting in what follows, but participation can also include other activities such as campaigning, donating, assembling, etc.) in accordance with their new, moderated beliefs. However, this result seems importantly *unfair*: if the Twisters act in this way, the political center of gravity will shift toward the positions of the Stickers. In some sense, the Stickers are being rewarded for their epistemic irresponsibility.

On the other hand, perhaps the Twisters can take note of the Stickers’ intransigence and vote according to their original position on the issues. However, this seems non-ideal in a different way. Here, the Twisters are being *insincere* in their actions. Rather than voting their conscience, they’re voting strategically in order to secure an

outcome they desire. In general, insincere voting or campaigning can reflect kind of cynical, Machiavellian stance that seems at odds with core democratic values and principles.⁵ Of course, in light of this problem, some might reject the sincerity norm as an important ideal for democratic participation. Alternatively, they might think that the sincerity norm does not apply to cases like these. However, something like the sincerity norm is foundational for influential accounts of democratic legitimacy and authority, in particular, the public reason and deliberative conceptions (Schwartzman 2011). Insofar as we can offer a solution to the problem of asymmetric compliance without jettisoning sincerity as an ideal then, we should.

Worsnip offers an elegant way out of this problem, one that does not sacrifice the sincerity norm. The idea is that the Twisters ought to vote according to their “personal take” on the issues. Someone’s personal take is simply their judgment about what’s supported by their pre-disagreement first-order evidence. A personal take in this sense comes apart from belief. Insofar as the Twisters take the higher-order evidence from the Stickers seriously, they *do* moderate their beliefs in the Stickers’ direction. Nonetheless, Worsnip’s idea is that we can usefully think of a separate doxastic attitude, the personal take, which is someone’s own assessment of what their first-order evidence supports.

In voting according to their personal takes, Worsnip argues that the Twisters need not be acting cynically or strategically violating a principle of representative fairness. Rather, an appropriate principle of representative fairness, Worsnip suggests, holds that each person’s current *personal take* ought to have equal weight in determining the outcome of the democratic process.

In addition to what Worsnip says in his paper, such a principle seems independently plausible. Democratic decision-making ideally seeks to incorporate a range of different perspectives, which respond to knowledge that is dispersed across society (cf. Hayek 1945; Anderson 2006) and, furthermore, processed through a diversity of lenses that are, in turn, shaped by individuals’ differing social and economic conditions (cf. Lepoutre 2020). A system governed by a principle according to which people vote in response to moderation via higher-order evidence would problematically *lose* important information in some sense. It would yield homogeneity where there is underlying diversity, and would likely not be appropriately responsive to its citizens’ interests. Homogeneity of this kind is also likely to undermine convergence toward the truth, which, some argue, is often aided by robust disagreement and division of epistemic labor (Mercier and Sperber 2017; Hallsson and Kappel 2020).

Before proceeding, I want to say a bit more to motivate why asymmetric compliance with certain epistemic norms seems unfair from the standpoint of democratic legitimacy.⁶ Worsnip himself gives more or less an intuitive gloss given the stylized case, but we might wonder in what sense the asymmetry is unfair given the two main ways of construing democratic legitimacy, namely the instrumentalist conception and the proceduralist conception.⁷ Take first the instrumentalist conception. In this view, democratic procedures are legitimate to the extent they track a procedure independent truth of the matter.⁸ Here,

⁵Something like the norm that participants in the democratic process ought to be sincere is defended by a range of philosophers, including Landemore (2013), Maskivker (2019), Rawls (1993), Ridge (2021), and Schwartzman (2011).

⁶Thanks to an anonymous referee for pressing me to say more about this.

⁷See Peter (2008) for a helpful overview.

⁸Thus, suppose that justice requires a more or less egalitarian distribution of resources. Then, democratic procedures would be legitimate insofar as they yield outcomes that approximate an egalitarian distribution (Arneson 2003; Wall 2007). It is also possible to view democracy’s value as instrumental but then reject its desirability for the reason that there may be other systems which purportedly better track the truth, as in Brennan (2016).

the problem with asymmetric compliance is not so much “fairness” but loss of truth-tracking for the reasons mentioned in the preceding paragraph. We can expect a situation where one side conciliates and votes based on that conciliation, while the other side refuses to budge, to be suboptimal in terms of truth-tracking – in part because it would be systematically biased in a way that precludes an adequate incorporation of information dispersed across society. Even if the problem here is not fairness *per se*, there is still an important way in which the situation is problematic.

The fairness problem gets a more solid grip on proceduralist conceptions of legitimacy. On these views, fair procedures are *constitutive* of, not merely instrumental to, democratic legitimacy.⁹ Of course, there is a sense in which the procedure here is minimally fair – each person gets to vote and gets one vote (and we can also suppose there isn’t voter suppression and the like). However, there is an important sense in which asymmetric compliance with genuine epistemic norms renders the situation epistemically unfair. Fabienne Peter (2008, 51) explains that a purely proceduralist account “defends a set of practices in which epistemic agents critically engage with each other under conditions of transparency and reciprocity . . . legitimate decisions are those which are the result of deliberation under conditions of political and epistemic fairness.” The reciprocity that Peter emphasizes here is missing under conditions of asymmetric compliance with epistemic norms.¹⁰

In what follows, I argue that while Worsnip’s personal take solution offers a plausible way out of the dilemma of asymmetric compliance in the case as described earlier, the dilemma nonetheless resurfaces when we move from the question of how we ought to *respond* to evidence to that of how we ought to *gather* evidence. Moreover, this is significant because it’s plausible that many of the disagreement-generating epistemic pathologies we observe in the modern democratic context – epistemic bubbles, echo-chambers, partisan news – arise because of how we gather evidence (Flores and Woodard 2023).

3. The zetetic aspect of asymmetric compliance

While the primary focus in epistemology has for a long time been how one ought to respond to the evidence one already possesses, recently several philosophers have begun to pay attention to how one ought to inquire, or gather more evidence, given one’s starting point (Friedman 2020; Thorstad 2021; Hughes 2023). There is a dispute here about whether norms on evidence gathering could be genuinely *epistemic* or whether they are practical.¹¹ Regardless, poor evidence-gathering can lead us to become epistemically impoverished, in ways that can be blameworthy (Boult 2021).

Where we end up epistemically – what beliefs we have, the sorts of evidence available to us, our perspective on the world – depends to a large extent on how we inquire, not just how we respond to the evidence we already possess. Depending on the particulars, someone who inquires badly, or refuses to inquire further when they should, might end

⁹There are also views that combine the two grounds for legitimacy – see Christiano (2004) and Estlund (2008).

¹⁰In a similar vein, Peter defends a conciliatory norm for agents faced with peer disagreement which drives its normative force from *accountability* to epistemic peers. “[T]o be a responsible epistemic agent,” Peter (2013, 1263) says, “is to be accountable to one’s peers in a deliberative process.” Even when faced with people who are not peers though, presumably such a norm would support minimal conciliationism so long as those who disagree are better than chance.

¹¹See, for example, Kelly (2003) and Feldman (2000) for a defense of the latter view. See Flores and Woodard (2023) for a defense of genuinely epistemic norms on evidence gathering.

up in a worse epistemic position than someone who gathers much of the relevant evidence but fails to draw the right conclusions, even if the former responds to their evidence flawlessly. Someone who has their “head in the sand” within some domain is likely to have an epistemically impoverished worldview with respect to that domain.

A particular zetetic activity that has gotten some recent attention is double-checking whether p . In some cases where we encounter opposing higher-order evidence, the right response is to double-check (Palmira 2024). A good example might be disagreement involving arithmetic – as when two people add up the dinner total and get different figures. Double-checking is *inquiring* – it is not a matter of simply adjusting one’s beliefs given the evidence one already has. And double-checking whether p can be rational even if one knows that p (Woodard 2024).

In the case of political matters, the relevant bases of evidence are likely to be large, and the interconnections between them complex (this is especially true of geo-political and economic subjects, for example). Here, rather than simply double-checking, a reasonable response to disagreement from people we consider careful and intelligent is to inquire into what their evidence might be. Perhaps they’re seeing something we don’t, and thus it behooves us to read cogent defenses of the other position, listen to well-informed defenders of that position, etc. All this involves gathering new evidence, not merely responding to higher-order disagreement-based evidence.

However, this opens up another avenue for asymmetric compliance. Thus, imagine the following situation. Suppose the Stickers and Twisters initially favor extreme policies E_1 and E_2 , respectively. In large part, this has to do with the types of news and analyses they consume – i.e. they operate on partial, biased subsets of evidence provided to them by different “vendors” at the marketplace of rationalizations (Williams 2023). In reality, however, the total available evidence favors a moderate policy, M .

Now, the Twisters observe the disagreement with the Stickers and think to themselves that perhaps there’s something they’ve been missing that the Stickers are seeing. In good faith, they try to sample evidence from outside their epistemic bubble, so to speak.¹² As they gather evidence from a range of sources, with their different emphases, etc., they come to have the attitude that is supported by the total available evidence, namely, that policy M is the thing to go for. Thus, they perform their epistemic due diligence that goes beyond simply responding to the evidence they already have – they also fulfill their zetetic duties. The Stickers, on the other hand, stay within their epistemic bubble and do not seek out evidence from the outside. What’s more, relative to their first-order evidence, the extreme position E_1 is indeed best supported.

However, this creates a problem with regard to representative fairness. Because the Stickers do not perform their epistemic due diligence, they are left supporting E_1 . And because the Twisters do take on and discharge their zetetic duties, they are left supporting M . The political center of gravity, then, moves toward the Stickers, who are in a sense rewarded for their intransigence.

Unlike the pure disagreement case, where the chief question is how to *respond* to one’s evidence, here there are bound to be uncomfortable trade-offs. Specifically, appealing to one’s personal take, i.e. a person’s own analysis of her first-order evidence, does not help. For, by stipulation in the case above, the personal take that the Twisters now have, given their expanded set of evidence, is that policy M is the way to go. The Twisters are then left with a dilemma. Either they can support policy M , in which case the center of gravity will move toward the Stickers’ position. Or they can strategically vote by continuing to support E_2 , which they now rightly believe is not a good policy.

¹²See Nguyen (2020) for a characterization of “epistemic bubbles,” and Sheeks (2023) for an argument that such bubbles are by nature epistemically problematic.

Of course, this above case is quite stylized, and real-life instances are bound to be less clean cut. First off, there will be a lot of internal diversity among disagreeing political groups, whose members may not all share the exact same position. In addition, some within a particular group may take on their zetetic duties while others may not. What's more, intransigence of the sort discussed above comes as a matter of degree.

Second, I do not mean to suggest that moderate or centrist positions are always, or even more often than not, likely to be correct. Neither do I want to assume that each side's epistemic bubble is always, or even usually, equally impoverished – some bubbles can be worse than others, for myriad reasons. Rather, the weaker assumption on which the problem of zetetic intransigence is significant is the following: Suppose one is in a particular epistemic bubble. It is unlikely that after getting outside of that bubble, by gaining further evidence, one's new set of total evidence will fit precisely with the original doxastic attitudes (on various policy matters, say) one had prior to further inquiry. New relevant evidence often calls for modified attitudes – perhaps increase/decrease of credence, suspension of belief, acknowledgment of risk, and so on. The new attitudes called for need not always move one toward the center, but it will be unlikely that staying put exactly where one started from is going to be appropriate.

To give a simple example, suppose someone starts off by reading only newspapers, magazines, and so on which are in favor of trade protectionism, as a means to protect domestic jobs. It's unlikely that adding to this evidence, by considering the arguments of economists, policy experts, etc., who defend easing restrictions (by pointing out comparative advantage, job creation in other industries, consumer surplus, etc.), calls for *no* change in one's original position. At the very least, the new evidence will point to various trade-offs, potential risks, etc., that would make a difference on the margins as to which policies to support, who to vote for, and so on. In addition, this may modify what it's rational to believe on related economic policy issues such as minimum wages or universal basic income. Cases where the process of getting out of one's epistemic bubble merits absolutely no subsequent moves on attitudes toward policy-relevant questions are likely to be rare.

For my purposes, it is not too important how these epistemic bubbles are maintained, and some ways of maintaining such bubbles may be more pernicious than others. What is important is that there is an "information cocoon" (Sunstein 2017) of sorts, where certain types of evidence are systematically underrepresented or excluded altogether. There may be several causal mechanisms by which the bubble is sustained, some of which might be harder to mitigate than others. Some epistemic bubbles might be sustained by an "echo chamber" wherein important relevant voices are actively discredited (Nguyen 2020). Regardless of the causal mechanism underlying the maintenance of an epistemic bubble, however, the key point for present purposes is that remaining in a bubble violates an important zetetic requirement, namely, to inquire into what countervailing evidence there might be.

Now, it might be thought that if the Stickers really are intransigent in the way sketched above, perhaps the Twisters can disregard them because they are not epistemically virtuous. Certainly, they are not epistemic peers in one sense, given that peers are usually stipulated to be equally careful and diligent regarding the relevant epistemic matters. But even if it's true that the two sides are not equally epistemically virtuous, it doesn't follow that their disagreement places no pressure on the Twisters. For one, as outlined earlier, even if A is more reliable than B within some domain, A still faces some epistemic pressure to moderate so long as B is more reliable than chance within that domain. The fact that B is not appropriately sensitive to higher-order evidence need not mean that she is less reliable than chance.

For instance, in evaluating which football team is likely to win the game, a range of different pieces of evidence is relevant. Some of these are first-order: they concern the team track records, the relative skills of the players, the effectiveness of the coaches, etc. There is also relevant higher-order evidence – in particular, assessments by others, particularly experts, about which team will likely win. Even if you were to totally ignore the second type of evidence (for example, by not watching SportsCenter or listening to commentary), it does not follow that you cannot be more reliable than chance.

What is more, when it comes to the *zetetic* significance of disagreement, matters are even more clear cut. The fact that someone dogmatically refuses to collect further evidence need not mean that they do not now possess evidence that might be relevant to some question. Suppose, for example, that Smith has seen Jones store the murder weapon in his study. Smith now is convinced that Jones committed the crime and refuses to collect further evidence, wrongly assuming that it will be misleading. This does not mean that the detective ought not to inquire as to why Smith believes Jones did it – for Smith possesses evidence, namely that Jones was in possession of the weapon, that is relevant to the case. And this may be so even if the total evidence that the detective has, *ex ante* or even *ex post*, suggests that Jones did not do it.

A somewhat different way of putting the problem is the following. Perhaps zetetic intransigence of the sort I've been discussing counts in favor of thinking that the other side is unreliable. Indeed, depending on how intransigent they are, they might plausibly be less reliable than chance.¹³ To some extent, of course, we are all biased inquirers – we have a well-known tendency for “myside” or “confirmation” bias (Stanovich 2021). This by itself need not mean we're less reliable than chance in general. However, we can certainly imagine inquirers whose zetetic dispositions are so biased that it renders them anti-reliable.

This brings out an important difference between the *zetetic* import of disagreement and the *evidential* import of disagreement. There are two distinct questions here that we could be interested in. The first, which has received much attention in epistemology, is: how should *the fact that someone disagrees with you* about *p* affect your doxastic attitude toward *p*? A second question, though, which is relevant presently is: how should the fact that someone disagrees with you affect *your future inquiries* with respect to *p*? Answers to the first question depend on how reliable the other person is. If the other person is more reliable than chance, you ought to reduce confidence in *p* (assuming you antecedently believe *p* or think it's more likely true than not, and your opponent disagrees).

However, in the zetetic case, we need not assume the other person is more reliable than chance. Suppose A is reasonably good at forming beliefs based on his total evidence. However, he is biased when it comes to gathering new evidence – let's suppose he has a strong tendency to search for evidence that confirms his beliefs and avoid looking for evidence that disconfirms his beliefs. Let's suppose this bias is so strong that it renders A less reliable than chance with respect to certain domains. Now, suppose B finds herself disagreeing with A about some matter. Does this give B reason to inquire into what A's evidence might be? It can, depending on how significant the proposition is. Given that we have limited time and mental resources, we cannot inquire into everything. But some propositions are significant enough to our epistemic lives such that we ought to inquire about them. Presumably, claims about important political or policy matters can fit this bill. The important point for present purposes is that B can have reason to inquire into A's evidence *even if* A is less reliable than chance.

¹³Thanks to an anonymous referee for raising this worry.

To illustrate, suppose that because A is in an epistemic bubble, he only receives evidence that supports p . On this basis, he forms the belief that p , thus properly responding to his first-order evidence. Nonetheless, because of biased sampling of evidence, A forms his beliefs in an unreliable way. B believes $\sim p$ and notices that A disagrees. Here, it still behooves B to inquire into what evidence A has. Even if this evidence does not fully defeat $\sim p$, it is relevant to the question of whether p . For instance, it can lead B to a more nuanced view about p or affect how she thinks it is connected to other issues, and so on. The point is that even a biased subset of evidence can contain good evidence.

Now, supposing that there is this zetetic norm to get outside one's epistemic bubble, it might still be thought that zetetic intransigence does not pose a genuine problem of fairness. Representative fairness, it may be argued, is about giving equal voice, in some sense, to each citizen's viewpoint. Of course, we would depart from this ideal of representative fairness if we gave some people's votes more weight as compared to others. But this is not what's going on. The Twisters expand their evidence base after noticing their disagreement with the Stickers and consequently come to support M. The Stickers, on the other hand, do not sample further evidence. But so long as each Twister and Sticker gets an equal voice in the system – in other words, each person gets one vote – there is no problem of fairness. Personal takes are all that matter here, and as far as representative fairness is concerned, it is unimportant *how* those takes are arrived at.

There are two things to say in response here. The first is that we plainly do care about how other democratic participants arrive at their personal takes. Otherwise, it is puzzling why we would care, for instance, about imposing reasonable limits on campaign finance (Dawood 2015), or the ways in which algorithms might influence voters' behaviors (Christiano 2022). We care about these things because the processes by which voters acquire information are important for fairness.

Second, consider what raises the problem of unfairness in the original case of asymmetric compliance with (non-zetetic) norms of responding to higher-order evidence. Worsnip (2023, 347) puts it in the following way: there seems to be "something unfair about a process that allows the political center of gravity to systematically shift toward the Stickers' views, simply because the Twisters (justifiably!) conciliate and the Stickers (unjustifiably!) don't." Part of what's going on here is that there is a genuine epistemic norm – namely, to moderate our beliefs in response to higher-order evidence – that is being followed by one party and not by the other. Furthermore, it is precisely this fact of non-mutual compliance with the norm that leads to a situation favoring the Stickers' views, whereby the political center of gravity shifts toward them. However, the situation of asymmetric compliance in the zetetic case is structurally analogous. There is a genuine epistemic norm requiring the collection of further evidence that is being followed by one party and not the other. And again, it is non-mutual compliance with this norm that is leading to the center of gravity shift toward the intransigent party.

Another way to put this is that it takes epistemic work to come to a reasonable assessment of the issues at stake in democratic participation. For one, citizens must take adequate account of one another's interests – and this takes some epistemic work, especially in a pluralistic society.¹⁴ It takes work to figure out what others' interests are, and this requires some form of dialog with those whose interests differ from one's own. Furthermore, I think it's plausible that part of the epistemic work required also involves broadening one's evidence base so that it's adequately representative. One fails in this

¹⁴The ideal of fair-mindedness in this sense is a core assumption for many democratic theorists; see for example Barry (1995) and Goodin (2003) for further discussion.

duty insofar as one remains in an epistemic bubble – and more so the more restrictive that epistemic bubble is.

That said, there is a familiar worry about whether the normative demands of democratic participation ought not to be too demanding. In this vein, Julia Maskivker (2019) defends a duty to vote with care, which requires fair-mindedness as well as meeting a minimal epistemic threshold pertaining to the policies under question and their potential consequences. Because the threshold is minimal, the duty to vote with care is not too demanding. However, it is plausible that this minimal threshold involves the good faith effort to sample evidence from a suitable range of sources.

To summarize, there is a structural symmetry between the case of asymmetric compliance with norms pertaining to higher-order evidence, as developed in detail in Worsnip (2023), and asymmetric compliance with plausible zetetic norms. At the very least, I take this structural symmetry to have a burden-shifting import. Supposing someone is convinced of the problem of representative unfairness in the original case from Worsnip, it seems the burden of proof is on them to say why this structurally analogous case does not raise the unfairness worry.

4. Significance

The zetetic aspect of our epistemic lives is significant insofar as what drives disagreement is exposure to distinct sets of evidence. In classic discussions of peer disagreement, the peers are stipulated to have the same evidence. However, this condition is rarely, if ever, met in the case of politically relevant information – for, each of us has a distinct set of relevant evidence constituted by our media/social media exposure, experiences, education, etc. Now, of course, it might be true that politically like-minded individuals have similar (even if technically distinct) sets of evidence, but it's plausible that especially when it comes to individuals on the opposite sides of the spectrum, the sets of evidence they have are substantially different, even though there may be some overlap.

Indeed, much recent empirical work points to differences in exposure to news, especially as a result of social media algorithms, as a driver of political polarization (Sunstein 2017; Levy 2021; Peterson, Goel, and Iyengar 2021). The proper epistemic response to recognizing one's selective exposure in this way goes beyond merely updating one's doxastic attitudes owing to higher-order evidence. Rather, the proper response involves, in many cases, an attempt to get beyond one's epistemic bubble, perhaps by diversifying one's sources (Worsnip 2019; Sheeks 2023).

What this suggests is that democratic fairness involves citizens making good faith attempts to listen to a diversity of voices, beyond simply those who are like-minded. At the very least, if the proportion of individuals within different political camps who fulfill such zetetic duties varies significantly, then there arises an element of democratic unfairness.

This possibility raises tensions for the epistemic dimensions of democratic participation that do not admit of easy resolution. When faced with asymmetric compliance with respect to how people collect evidence, there are three options, none of which are entirely satisfactory. The Twisters, in the stylized case above, can either: i) stay within their epistemic bubble and refuse to collect outside evidence; ii) collect outside evidence but still vote in ways that would be recommended by their original set of evidence, i.e., the initial epistemic bubble; or iii) monitor compliance, or lack thereof, from the Stickers, and depending on levels of compliance, vote strategically to attempt to secure outcomes that most closely match what is supported by the evidence that the Twisters currently have.

Option i) is problematic because it calls for epistemic irrationality. If there are genuine epistemic norms on evidence gathering, then calling for a stop to such gathering is asking people to be epistemically irrational. Now, one possibility here is that if a relevant group happens to be in a “good” epistemic bubble, perhaps refusal to gather outside evidence is epistemically permissible, or even creditworthy. Perhaps any outside evidence would be misleading.¹⁵ However, as Sheeks (2023) has recently pointed out, in order to be justified in believing that one’s epistemic bubble is good in some objective sense, one has to exit the bubble. That is, the judgment that an epistemic bubble is truth-conducive can only be rationally made from outside of it. Hence refusing to collect evidence from the outside, as the Stickers do, is epistemically irrational.

Now, it might be argued that epistemic rationality is not all that matters. Sometimes we should sacrifice epistemic rationality, all things considered, to secure certain practical goals. For example, even if an underdog sports team has a low objective chance of winning, perhaps they should, all things considered, convince themselves that they have a good chance – this belief itself, epistemically irrational though it may be, can raise their objective chance of winning (cf. James 1896). There are two issues with this strategy in the political case, however. First, at the very least, this reveals a non-ideal tradeoff between epistemic and practical rationality within democratic politics. Second, and more importantly, it’s not obvious we can reliably make such judgments from within our bubbles, especially insofar as we are susceptible to motivated partisan reasoning.¹⁶ And even if, unlike the Twisters, the Stickers are behaving epistemically irresponsibly, they may be closer to the truth on some particular issue of partisan disagreement, given the total available evidence. How are the Twisters to rule out this possibility in an epistemically responsible way without collecting outside evidence? Unlike the sports team case, furthermore, ideally what matters when it comes to important policy issues is not simply *winning* but securing just outcomes, hopefully by getting things right. From a proceduralist standpoint, part of what matters here is mutual reciprocation on interpersonal deliberative norms. Epistemic rationality is less dispensable in achieving such goals.

Option ii) is also non-ideal for several reasons. First, it prescribes voting in a way that goes against our total evidence, and what’s more, even against our personal takes – i.e., our total first-order evidence, bracketing the higher-order, disagreement-based considerations. Second, evidence-gathering is a dynamic process, and rational agents continually update their behavior as new evidence comes in. Therefore, in general, it will be highly irrational to have one’s behavior determined by an “original” set of evidence from some previous time. To make things stark, suppose the Twisters notice asymmetric compliance in the year 2000. Now, more than two decades later, how should they participate politically? It would seem absurd to say they should do so in ways that they would have given their original set of evidence from the year 2000, given all that has come to light since then. Furthermore, it will be hard to determine where the appropriate origin point is, considering that disagreement is part and parcel of democratic politics. As Worsnip (2023, 348) puts it, “We form our initial beliefs about many political issues already aware of significant disagreement about the relevant subject matter. As such, for many real-world issues, there simply is no “original” belief that one had prior to encountering disagreement.”

¹⁵Relatedly, Fantl (2018) has argued that sometimes one ought not to be open-minded, as this can be a way of losing knowledge.

¹⁶The literature on this topic is vast, but for a helpful overview, see Van Bevel and Pereira (2018). For a general characterization of motivated reasoning, see Kunda (1990).

Option iii) raises the worry of violating the sincerity norm in democratic politics. Ideally, citizens are to vote and participate in ways that reflect how they see things in terms of what would adequately represent their legitimate interests and their conceptions of fairness. Someone might argue, however, that it's precisely to secure fairness that the Twisters ought to vote strategically. Even bracketing this problem, however, there is a significant remaining worry. Continually monitoring asymmetric compliance, especially when it comes to zetetic norms, is bound to be very difficult if not practically impossible. What is needed to vote strategically in a fair way is not just the knowledge *that* there is asymmetric compliance when it comes to gathering evidence, but the more precise, issue-by-issue nature and extent of the asymmetries. In other words, what is needed is knowledge of the total politically relevant evidence the other side has. This is bound to be challenging, even to a first approximation, for most of us.

This worry is compounded by new survey data showing significant gaps between what partisans think the other side believes versus what the other side actually believes (Yudkin, Hawkins, and Dixon 2019). What is especially noteworthy is that this gap pertains specifically to views about out-party policy positions, which are presumably much more tractable to determine than out-party total *evidence* relevant to such policy positions.

Notwithstanding the above problems, depending on the specifics, one of the above three options might well be the way to go, all things considered. For instance, perhaps, depending on the circumstances, the all-things-considered best option for the Twisters might be to vote strategically. That is, the balance of practical reasons may point to them doing so. Nonetheless, it seems that there is an important tradeoff here so that there is an unresolvable tension between epistemic rationality, political sincerity, and political fairness.

5. Conclusion

Asymmetric compliance with epistemic norms raises relatively underexplored problems for democratic participation. Disagreement with our co-citizens places at least some pressure on us to moderate our beliefs. However, if they do not reciprocate, we seem to be left with two options, both not entirely satisfactory – either we refuse to update in ways that are epistemically called for, or we attempt to vote strategically. If asymmetric compliance concerns simply how the relevant parties *respond* to higher-order evidence, a plausible solution to this dilemma is available – namely, to participate in ways reflecting one's "personal take." On the other hand, if there is asymmetric compliance when it comes to norms of inquiry or evidence gathering, I argue there is no entirely satisfactory way to resolve the underlying tensions. This is significant because, plausibly, much democratic disagreement is driven by how we seek out new evidence, and not simply how we respond to evidence we already have. Democracy thus presents us with a range of messy trade-offs that cross-cut epistemic and practical considerations, and which do not admit of easy resolution.¹⁷

References

Achen C.H., & Bartels L.M. (2016). *Democracy for Realists: Why Elections Do Not Produce Responsive Government*. Princeton, NJ: Princeton University Press.

¹⁷I would like to thank an anonymous referee at *Episteme* for helpful suggestions. For discussion, thanks to Thomas Christiano, Michael Hannon, Elise Woodard, and an audience at the PPE Society in New Orleans.

- Aikin S., Harbour M., Neufeld J., & Talisse R.B.** (2010). 'Epistemic Abstainers, Epistemic Martyrs, and Epistemic Converts.' *Logos & Episteme* 1(2), 211–19.
- Anderson E.** (2006). 'The Epistemology of Democracy.' *Episteme* 3(1–2), 8–22. <https://doi.org/10.3366/epi.2006.3.1-2.8>
- Arneson R.** (2003). 'Debate: Defending the Purely Instrumental Account of Democratic Legitimacy.' *Journal of Political Philosophy* 11(1), 122–32.
- Barry B.** (1995). *Justice as Impartiality*. Oxford: Oxford University Press.
- Boult C.** (2021). 'There Is a Distinctively Epistemic Kind of Blame.' *Philosophy and Phenomenological Research* 103(3), 518–34.
- Brennan J.** (2016). *Against Democracy*. Princeton, NJ: Princeton University Press.
- Caplan B.** (2008). *The Myth of the Rational Voter*. Princeton, NJ: Princeton University Press.
- Christensen D.** (2007). 'Epistemology of Disagreement: The Good News.' *Philosophical Review* 116(2), 187–217.
- Christiano T.** (2004). 'The Authority of Democracy.' *Journal of Political Philosophy* 12(3), 266–90.
- Christiano T.** (2022). 'Algorithms, Manipulation, and Democracy.' *Canadian Journal of Philosophy* 52(1), 109–24.
- Dawood Y.** (2015). 'Campaign Finance and American Democracy.' *Annual Review of Political Science* 18, 329–48.
- Elga A.** (2007). 'Reflection and Disagreement.' *Noûs* 41(3), 478–502.
- Estlund D.** (2008). *Democratic Authority*. Princeton, NJ: Princeton University Press.
- Fantl J.** (2018). *The Limitations of the Open Mind*. Oxford, New York: Oxford University Press.
- Feldman R.** (2000). 'The Ethics of Belief.' *Philosophy and Phenomenological Research* LX(3), 667–95.
- Flores C., & Woodard E.** (2023). 'Epistemic Norms on Evidence Gathering.' *Philosophical Studies* 180, 2547–2571. <https://doi.org/10.1007/s11098-023-01978-8>
- Friedman J.** (2020). 'The Epistemic and the Zetetic.' *Philosophical Review* 129(4), 501–36.
- Goodin R.** (2003). *Reflective Democracy*. Oxford: Oxford University Press.
- Hallsson B., & Kappel K.** (2020). 'Disagreement and the Division of Epistemic Labor.' *Synthese* 197, 2823–47. <https://doi.org/10.1007/s11229-018-1788-6>
- Hayek F.A.** (1945). 'The Use of Knowledge in Society.' *The American Economic Review* 35(4), 519–30.
- Hughes N.** (2023). 'Epistemic Feedback Loops (or: How Not to Get Evidence).' *Philosophy and Phenomenological Research* 106, 368–93.
- James W.** (1896). *The Will to Believe*. New York: Longmans, Green, and Co.
- Kelly T.** (2003). 'Epistemic Rationality as Instrumental Rationality: A Critique.' *Philosophy and Phenomenological Research* LXVI(3), 612–40.
- Kelly T.** (2010). 'Peer Disagreement and Higher-Order Evidence.' In R. Feldman and T. Warfield (eds), *Disagreement*, pp. 111–174. New York: Oxford University Press.
- King N.L.** (2012). 'Disagreement: What's the Problem? Or A Good Peer Is Hard to Find.' *Philosophy and Phenomenological Research* 85(2), 249–72.
- Kunda Z.** (1990). 'The Case for Motivated Reasoning.' *Psychological Bulletin* 108(3), 480–98.
- Landmore H.** (2013). *Democratic Reason: Politics, Collective Intelligence, and the Rule of the Many*. Princeton, NJ: Princeton University Press.
- Lepoutre M.** (2020). 'Democratic Group Cognition.' *Philosophy and Public Affairs* 48(1), 40–78.
- Levy R.** (2021). 'Social Media, News Consumption, and Polarization: Evidence from a Field Experiment.' *American Economic Review* 111(3), 831–70.
- Maskivker J.** (2019). *The Duty to Vote*. New York: Oxford University Press.
- Mercier H., & Sperber D.** (2017). *The Enigma of Reason*. Cambridge, MA: Harvard University Press.
- Nguyen C.T.** (2020). 'Echo Chambers and Epistemic Bubbles.' *Episteme* 17(2), 141–61. <https://doi.org/10.1017/epi.2018.32>
- Palmira M.** (2024). 'Higher-Order Evidence and the Duty to Double-Check.' *Noûs*. 58(3), 799–824. <https://doi.org/10.1111/nous.12479>
- Peter F.** (2008). 'Pure Epistemic Proceduralism.' *Episteme* 5(1), 33–55.
- Peter F.** (2013). 'The Procedural Epistemic Value of Deliberation.' *Synthese* 190, 1253–66.
- Peterson E., Goel S., & Iyengar S.** (2021). 'Partisan Selective Exposure in Online News Consumption: Evidence from the 2016 Presidential Campaign.' *Political Science Research and Methods* 9(2), 242–58.
- Rawls J.** (1993). *Political Liberalism*. New York: Columbia University Press.
- Ridge M.** (2021). 'Voting for Less than the Best.' *Journal of Political Philosophy* 29(3), 404–26.
- Schwartzman M.** (2011). 'The Sincerity of Public Reason.' *Journal of Political Philosophy* 19(4), 375–98.

- Sheeks M.** (2023). 'The Myth of the Good Epistemic Bubble.' *Episteme*. **20**(3), 685–700. <https://doi.org/10.1017/epi.2022.52>
- Stanovich K.E.** (2021). *The Bias That Divides Us*. Cambridge, MA: MIT Press.
- Sunstein C.** (2017). *#Republic*. Princeton, NJ: Princeton University Press.
- Thorstad D.** (2021). 'Inquiry and the Epistemic.' *Philosophical Studies* **178**(9), 2913–28.
- Van Bevel J., & Pereira A.** (2018). 'The Partisan Brain: An Identity-Based Model of Political Belief.' *Trends in Cognitive Science* **22**, 213–224. <https://doi.org/10.1016/j.tics.2018.01.004>
- Wall S.** (2007). 'Democracy and Equality.' *The Philosophical Quarterly* **57**, 416–38.
- White R.** (2009). 'On Treating Oneself and Others as Thermometers.' *Episteme* **6**(3), 233–50.
- Williams D.** (2023). 'The Marketplace of Rationalizations.' *Economics & Philosophy* **39**(1), 99–123. <https://doi.org/10.1017/S0266267121000389>
- Woodard E.** (2024). 'Why Double-Check?' *Episteme*. **21**(2), 644–667. <https://doi.org/10.1017/epi.2022.22>
- Worsnip A.** (2019). 'The Obligation to Diversify One's Sources: Against Epistemic Partisanship in the Consumption of News Media.' In C. Fox and J. Saunders (eds), *Media Ethics: Free Speech and the Requirements of Democracy*, pp. 240–264. London: Routledge.
- Worsnip A.** (2023). 'Compromising with the Uncompromising: Political Disagreement under Asymmetric Compliance.' *Journal of Political Philosophy* **31**(3), 337–57. <https://doi.org/10.1111/jopp.12288>
- Yudkin D.A., Hawkins S., & Dixon T.** (2019). 'The Perception Gap: How False Impressions Are Pulling Americans Apart.' More in Common. <https://perceptiongap.us/media/zaslaroc/perception-gap-report-1-0-3.pdf>

Hrishikesh Joshi is an assistant professor of philosophy at the University of Arizona. He works in social epistemology, political philosophy, and philosophy, politics, and economics. He also has interests in ethical theory and metaethics. He is currently working on a book on politics and public discourse.