

The Rational Partisan

Brian Hedden

bhedden@mit.edu

Massachusetts Institute of Technology

Abstract

Our politics are increasingly polarised. Polarisation takes many forms. One is increasing clustering, whereby people hold down-the-line liberal or conservative views on a wide range of political issues, even when those issues are orthogonal to each other. Some philosophers think that such clustering is indicative of irrationality, and so if you find yourself in one of several clusters of opinion, you should decrease your confidence that all your political beliefs are true. I argue that the reverse is true, presenting a simple model of belief-formation on which finding yourself in one of several clusters of opinion on orthogonal issues should increase, rather than decrease, your confidence that all your beliefs are true.

Keywords: polarization, clustering, social epistemology, Condorcet Jury Theorem, judgment aggregation

Brian Hedden is a professor of philosophy at the MIT. He works primarily on epistemology, decision theory, and ethics, and has written on collective action problems, legal standards of proof, value aggregation, and algorithmic fairness, among other topics.

1 Polarization and Orthogonality

Political polarization is widely regarded as one of the most pressing issues facing the world today. Polarization also takes many different forms (Bramson et al., 2016). The type of polarization that I will discuss is *clustering*.¹ Clustering is best illustrated by example. Consider the following political questions:

- Should abortion be legal?
- Should there be strict gun control?

¹Political scientists sometimes use the terms ‘ideological consistency’ to refer to such clustering. See Converse (1964) for a classic early study of ideological consistency in the US.

- Should the minimum wage be increased?
- Should more immigrants be allowed into the country?
- Should there be affirmative action?
- Is climate change a major problem?
- Should the government provide universal health care coverage?

Many people hold down-the-line ‘liberal’ beliefs, answering ‘yes’ to all of these questions. And many hold down-the-line ‘conservative’ beliefs, answering ‘no’ to all of them. Moreover, these sets of beliefs are more common than other particular alternative sets of beliefs, corresponding to a mix of yes and no answers to these questions. Note that this is compatible with there being many people who answer with a mix of yeses and noes than with all yeses or all noes. It’s just that there are fewer people who give any *particular* sequence of answers consisting of both yeses and noes than who answer with all yeses or with all noes.²

In this way, there is a ‘liberal’ cluster consisting of those people who answer ‘yes’ to all—or at least most—of these questions, along with a ‘conservative’ cluster consisting of those who answer ‘no’ to all—or at least most—of them.

Michael Huemer (2016) and Hrishikesh Joshi (2020) have recently argued that such clustering is epistemically suspicious and indicates irrationality (or nearby epistemic foibles³) on the part of those in the clusters. They think that if you learn that you fall into one of several clusters, you should increase your confidence that you have been irrational, and you should accordingly increase your confidence that you’ve made an error somewhere, and hence that not all of your political beliefs are true.⁴

This is because the political questions above are largely orthogonal or logically unrelated to each other. Roughly, the idea is that your view on abortion shouldn’t obviously determine, or even influence, your view on the minimum wage, or *vice versa* (and so on for all pairs of issues mentioned above). And

²See <https://www.pewresearch.org/politics/2014/06/12/political-polarization-in-the-american-public/> for a detailed, albeit somewhat dated, analysis of clustering in the American populace.

³Joshi (2020) suggests that being in a cluster is evidence that your beliefs have been subject to irrelevant influences, or that you possess only a biased subset of the possible evidence on political matters.

⁴Of course, this presupposes that political beliefs can be true or false even when they concern moral rather than purely empirical matters. I adopt this ‘moral realist’ presupposition in what follows.

so it's suspicious that we find the clustering that we do. It indicates the operation of groupthink, or social conformity, or some other non-truth-tracking mechanism for forming beliefs.

I will argue that this pessimistic response to clustering is wrong. Learning that you fall into one of several clusters need not give you evidence that you have been irrational. And it should increase, rather than decrease, your confidence that you're right about everything, i.e. that all your political beliefs are true.

2 Orthogonality

The notion of orthogonality (or logical unrelatedness, which I'll understand as equivalent) plays a key role in the arguments of Huemer and Joshi. Now, for our purposes, we don't need a precise definition of orthogonality. A rough and ready, intuitive understanding will suffice. But I want to try to pin it down a bit more. (The impatient reader can skip the next few paragraphs without loss.)

Huemer (2016, 458) writes that 'Two beliefs are "logically unrelated" if neither of them, even if true, would constitute evidence for or against the other.' And Joshi (2020, 37) writes that two issues are orthogonal if and only if 'your position on one of them doesn't commit you to any particular position on the others.' These glosses are quite different, with Huemer's employing the concept of evidential connections and Joshi's employing the concept of rational commitment. And neither gloss is adequate as it stands.

Against Joshi, the notion of commitment is too strong. Consider the issues of whether the government should provide universal health care coverage and whether the government should provide publicly funded education. These issues are not orthogonal, or at least not to the same degree as those mentioned above. It is quite natural for your view on government-funded health care to be closely connected to your view on government-funded education. But your view on government-funded health care doesn't fully *commit* you to a given view on government-funded education. You *could*, for instance, favor government-funded health care while opposing government-funded education. And so Joshi's gloss would wrongly classify these issues as orthogonal.

And against Huemer, we will see later that the very fact of clustering can create evidential connections between beliefs that, intuitively, are orthogonal or unrelated. It gives us a kind of 'higher-order' evidence (as epistemologists call it) suggesting that e.g., a given belief on abortion is true if and only if some

belief on gun control is, due to the fact that many people hold either both beliefs or neither.⁵

We may be able to improve on the glosses from Huemer and Joshi to avoid these problems. The problem for Joshi was that with two non-orthogonal issues, it could be that there is rational pressure to adopt some position on one of them if you adopt a given position on the other, even though the latter position doesn't fully *commit* you to the former. So we could instead just say that two issues are orthogonal just in case adopting a given position on one doesn't yield any rational pressure to adopt a given position on the other.

And the problem for Huemer was that—as we'll see later—clustering itself can provide a kind of higher-order evidence suggesting that certain intuitively orthogonal or logically unrelated beliefs are either all true or all false. But we could instead say that two beliefs are orthogonal or logically unrelated just in case neither one of them, if true, would constitute evidence for or against the other, *setting aside any higher-order evidence about who holds which beliefs*.⁶

Now, we have two different glosses on the notion of orthogonality, one in terms of rational pressure and the other in terms of probability. We might wish to investigate whether the two glosses are equivalent, given certain background assumptions, and if not, which gloss is superior. But again, for our purposes, nothing much hangs on exactly how we formulate the notion of orthogonality, and so I will set this issue to the side.

You might question whether the issues we began with really are orthogonal. You might suggest, for instance, that liberals are united by an underlying belief in big government, while conservatives are united by an underlying belief in small government. But whatever truth there is to the association of liberalism

⁵Higher-order evidence is evidence about how confident you ought to be in a given proposition. One kind of higher-order evidence is evidence about the beliefs of other people, even people who have the very same evidence as you. For classic discussions of higher-order evidence, see Kelly (2005) and Christensen (2007). For a model of clustering in which higher-order evidence plays a role in generating distinct clusters, see Freeborn (2024).

⁶We could make this idea of evidential unconnectedness more precise by cashing it out in terms of probabilistic independence. Two propositions A and B are probabilistically independent of each other just in case $P(A \mid B) = P(A)$ and $P(B \mid A) = P(B)$, i.e. just in case the conditional probability of each one given the other is equal to the unconditional probability of the one. Then, we can say that two beliefs are orthogonal just in case they are probabilistically independent of each other *given only the first-order evidence*, or the evidence bearing directly on the question (e.g., the evidence about the likely effects of an abortion ban, whether fetuses are conscious, etc.), regardless of whether they're probabilistically independent of each other given the higher-order evidence (e.g., the evidence concerning who holds which beliefs). Note that this gloss relies on their being a sharp distinction between first-order and higher-order evidence. See Hedden and Dorst (2022) for an argument that all evidence is higher-order evidence, and hence we cannot partition one's evidence into two disjoint subsets, one containing all the first-order evidence and the other containing all the higher-order evidence.

with big government and of conservatism with small government, it fails to explain why liberals are pro-choice while conservatives are pro-life. After all, the pro-life position is a big government one, insofar as it involves a commitment to intruding into private life by banning and punishing abortion. Alternatively, you might think that liberals are united by secularism while conservatives are united by religiosity. But even if—and it’s a big ‘if’—this could explain why liberals and conservatives adopt opposing views on certain hot-button social issues like abortion and marriage equality, it seems to have little to do with their opposing views on issues like immigration, affirmative action, and climate change. Or you might think that liberals are united by a commitment to the social good and standing up for the little guy, while conservatives lack this commitment. But this doesn’t explain the divide on abortion, where conservatives would argue that their pro-life position involves standing up for the littlest guy of all, the unborn fetus. And it seems flatly irrelevant to the question of whether climate change is a big problem.

There is much more to say here.⁷ But I’m going to concede the orthogonality of these political questions, for the sake of argument. And I refer the interested reader to Joshi (2020, 43-8) and Lewis and Lewis (2022, Ch. 4) for extensive critical surveys of possible underlying principles and character traits explaining the liberal-conservative divide; they find them all wanting.

Given the orthogonality of our political questions, why think that clustering is epistemically suspicious? In the next section, I will explain my opponents’ arguments for this conclusion.

3 Is Clustering Suspicious?

As I said, some philosophers argue that learning that you’re in one of two clusters, e.g., the liberal one or the conservative one, gives you evidence that you’ve been irrational and that not all of your political beliefs are true. Huemer (2016, 458) argues that if there were some people who were just good at getting to the truth, then

we would not expect the existence of an *opposite* cluster of be-

⁷Lakoff (1996) argues that liberals and conservatives employ different guiding metaphors, with conservatives thinking of the government as a ‘strict father’ and liberals thinking of it as a ‘nurturant parent.’ Haidt (2012) argues that liberals and conservatives differ in their views about the importance of different ‘moral foundations’ such as care, fairness, loyalty, respect for authority, sanctity, and liberty. See Joshi (2020, 45-8) for excellent critical discussion of the latter.

liefs. . . Why would there be a significant number of people who tend to embrace the opposite beliefs on all these issues? It is not plausible to suppose that there are some people who are in general drawn toward falsity. Even if there are people who are not very good at getting to the truth (perhaps they are stupid, ignorant, etc.), their beliefs should be, at worst, *unrelated* to the truth; they should not be systematically directed *away* from the truth. Thus, while there could be a “true cluster” of political beliefs, the present consideration strongly suggests that neither the liberal nor the conservative belief-cluster is it. (emphasis in original)

And Joshi (2020, 36) considers an agent who finds herself in a cluster and argues that

If she is justified in consistently adopting the party line, it must be true that her side is reliable on the issues that are the subject of disagreements. It would then follow that the other side is anti-reliable with respect to a host of orthogonal political issues. Yet, it is difficult to find a psychologically plausible explanation for why one side would get things reliably wrong with respect to a wide range of orthogonal issues.

The conclusion of Huemer and Joshi is also espoused by Lewis and Lewis (2022, 93), who say that since ‘there is no essence behind ideology and therefore no essence behind what each party stands for, then it follows that neither party (nor its associated ideology) has a monopoly on truth.’ Here, an ‘essence’ is an underlying principle, commitment to which underlies all of that cluster’s particular political beliefs.

Note that the argument from Lewis and Lewis is different from those of Huemer and Joshi. They make no mention of anti-reliability. Instead, Lewis and Lewis seem to think that the very orthogonality of the issues entails that it’s very unlikely that either cluster has all true beliefs. In the next section, I will show that while this may be true of *individuals*—orthogonality makes it unlikely that any given individual gets everything right—the same is not true of clusters.

What about Huemer and Joshi? They argue that if there are two clusters, and if members of one them are right about everything, then members of the other would have to be anti-reliable (or systematically directed away from the

truth, or in general drawn toward falsity). But it is implausible that there are many such anti-reliable people. So probably, neither cluster is such that its members are right about everything.

But against Huemer and Joshi, supposing that members of one cluster are right about everything does *not* require seeing members of the other cluster as anti-reliable. For the clusters actually agree on a vast number of political issues. For instance, liberals and conservatives alike think that the US should not launch a preemptive nuclear strike against the UK.⁸ They agree that interracial marriage should be legal. They agree that slavery should be prohibited. And so on. It's just that these issues aren't seen as salient political issues, given the agreement of the two main clusters. This is reminiscent of Bertrand Russell's suggestion that the apparent lack of progress in philosophy may be an artefact of issues ceasing to be seen as philosophical issues once they are definitively resolved (Russell, 1912, Ch. 15).⁹ It is worth adding that the two main clusters agree even on some issues that not everyone agrees on. For instance, they agree that the government can justifiably mandate the wearing of seatbelts, or that the government can justifiably require a driver's licenses to operate motor vehicles, even though libertarians tend to disagree.

So the argument from Huemer and Joshi rests on a false premise, namely that members of one cluster being right about everything requires that members of the other be anti-reliable. Nevertheless, that is not the end of the story. For we still need to explain why, if members of one cluster are right about everything, there is nonetheless another *cluster*, rather than just a bunch of people whose errors are randomly distributed.

Here is a natural explanation: Those in the other cluster are following an opinion leader, and that's why they have formed a cluster. But Huemer and Joshi would likely say that this is more grist for their mill. For insofar as one cluster was generated by its members following some opinion leader, so,

⁸Or so I assume (and hope!).

⁹To be fair to Joshi, he does sometimes qualify this claim about anti-reliability, saying that if members of one cluster are right about everything, then those in the other cluster must be anti-reliable 'with respect to political issues on which there is partisan disagreement' (Joshi, 2020, 41). But even with this qualification, I still take issue with his assertion. It follows from the supposition that members of one cluster are right about everything that anyone not in the cluster is wrong about the issues on which they disagree with members of that cluster. This much is trivial. But anti-reliability is a modal notion, and nothing modal follows from this trivial observation. It might be highly contingent which issues are the subject of partisan disagreement. And, even when a person is wrong about a certain issue, it might be highly contingent that they happen to arrive at this wrong view. So even if members of one cluster are right about everything, it doesn't follow that anyone else—even members of some other cluster—are anti-reliable about anything.

probably, was the other cluster generated by the same mechanism. The situation is probably symmetric. Therefore, if you find yourself in a cluster, you should think that you yourself likely formed your own beliefs by following an opinion leader. But—they might argue—following your opinion leader is irrational, insofar as you lack evidence that your opinion leader is more reliable than the other cluster’s opinion leader. And unless your opinion leader is hyper-reliable, which is implausible, your beliefs are unlikely to all be true.

Is that right? Is the situation inherently symmetric, such that insofar as one cluster was generated via deference to an opinion leader, so, probably, was the other? No. In the next section, I’ll present a simple model on which we should expect an asymmetry. In my model, if we find that there are two clusters of opinion, we should think that one of them was generated via deference to some opinion leader, while the other consists of independent-minded people who are right about everything. And so learning that you are in a cluster should increase, rather than decrease, your confidence that all of your beliefs are true.

4 Clustering and Condorcet

Here is a simple model, inspired by that underlying the Condorcet Jury Theorem (Condorcet, 1785/1976). We have a large population of citizens and a set of orthogonal binary political questions. Some citizens are randomly selected to be potential opinion leaders. Everyone else has some probability of following an opinion leader, in which case one randomly chooses one of the pre-selected potential opinion leaders and adopts all of that opinion leader’s answers to our political questions.¹⁰ Whether one follows a leader, and if so which leader one follows, is probabilistically independent of what the right answer is to any given question.

If one doesn’t follow any opinion leader, then one is independent-minded. All independent-minded people are better than random at getting the right answer to each political question. (Because the potential opinion leaders are independent-minded, this means that everyone in the population, including the followers, is better than random with respect to each question.) Their being better than random on each question means that each person’s probability of getting the right answer to any given question is greater than 0.5. And in the simplest version of the model, which I’ll assume going forward, this probabil-

¹⁰Compare the model of opinion leaders and followers in Boland et al. (1989).

ity—which is often called a ‘competence level’—is the same for everyone and for each question.

To capture the idea of independent-mindedness, whether one independent-minded person gets the right answer to a given question is probabilistically independent of whether any other independent-minded person gets the right answer to that question.¹¹ And to capture the idea that the questions are orthogonal, whether a given independent-minded person gets the right answer to one question is probabilistically independent of whether they get the right answer to some other question.

That is the model. Here is the key implication: If we find n clusters of opinion, we should think that $n - 1$ of them are based on people following $n - 1$ different opinion leaders, while the remaining cluster consists of independent-minded people who have clustered around the truth. People in the core of that cluster are right about everything. (Those giving nearby sequences of answers should probably still count as being in that cluster, but they are only right about *nearly* everything.) It then follows that learning that you are in a cluster should increase your confidence that all of your beliefs are true, since anyone with all true beliefs almost certainly is a member of a cluster.¹²

To see why this implication falls out of the model, ignore all of the followers and focus on the independent-minded people. Because each one’s getting the right answer to a given question is probabilistically independent of any other’s doing so, and because each one’s getting the right answer to one question is probabilistically independent of their getting the right answer to any other, we can think of them as tossing coins, one coin toss per political question. And because each is better than random at getting the right answer to any given question, we can think of them as tossing *biased* coins, one toss per political question. In particular, where getting heads on the i -th toss corresponds to getting the right answer to the i -th question, we can think of them as tossing coins which are biased towards heads.

Then, given a suitably large¹³ population, we expect the coin tossers to

¹¹See Goodin and Spiekermann (2018, Ch. 5) for discussion and defense of the independence assumption, which they describe as ‘the most contested and most misunderstood aspect of Condorcet’s jury theorem.’ Among other things, they make clear that independence does not require lack of any interaction.

¹²More precisely, the conditional probability of being in a cluster given that all your beliefs are true is higher than the unconditional probability of being in a cluster, and so, by the symmetry of positive probabilistic relevance, the conditional probability of having all true beliefs given that you’re in a cluster is higher than the unconditional probability of having all true beliefs.

¹³In particular, suppose that the population is large enough for it to be nearly certain that

cluster, and to cluster around the all-heads sequence. That is because the all-heads sequence is more probable for any given person than any alternative sequences, followed in probability by all of the sequences consisting of all heads but one, followed in probability by all of the sequences consisting of all heads but two, and so on. So, the all-heads sequence has the highest expected frequency in the population, followed by the various all-heads-but-one sequences, followed by the various all-heads-but-two sequences, and so on.

This fact, in turn, follows from the nature of the binomial distribution. Where n is the number of tosses and p is the probability of getting heads on a given toss, the probability of any particular sequence consisting of m heads and $n - m$ tails is $p^m \times (1 - p)^{n-m}$. If $p > 0.5$, then $p > (1 - p)$, and so $p^m \times (1 - p)^{n-m}$ is increasing in m . This means that any given sequence with more heads is more probable (and hence has a higher expected frequency) than any given sequence with fewer heads.¹⁴

Returning to the original setting of politics, we see that just as the tossers of heads-biased coins are expected to cluster, and to cluster around the all-heads sequence, so the independent-minded and better than random people are expected to cluster, and to cluster around the truth. Among them, the sequence of answers in which all the answers are true has a higher expected frequency than any other sequence, followed by the various sequences consisting of all true answers but one, followed by the various sequences consisting of all true answers but two, and so on.¹⁵

This means that learning that you are a member of a cluster should increase your confidence that you're right about everything, since anyone who is right about everything almost certainly is a member of a cluster. This is true despite

relative frequencies closely approximate probabilities, so that we can, with near certainty, infer approximate relative frequencies from probabilities.

¹⁴This does not mean that for any person, the expected number of heads is n . In fact, the *expected* proportion of heads in one's sequence is p . For instance, if $p = 0.6$, then the expected percentage of heads in one's sequence is 60%. Still, the all-heads sequence is more probable than any *particular* sequence consisting of around 60% heads. It's just that there are so many more of the latter than the former that it's more probable that you get around 60% heads in some arrangement or other than that you get all heads.

¹⁵This is how things go in the simplest case where the competence level is the same for each person and for each question. But things go similarly even if competence levels differ by person and/or by question, provided that they are always above 0.5. In particular, each person's probability of giving all true answers is greater than their probability of giving any alternative sequence of answers. But now, it needn't be the case that all sequences consisting of all true answers but one are more probable than *all* sequences consisting of all true answers but two, and so on. But still, sequences will generally be probable to the extent that they include more true answers, and so we still expect people to cluster, and to cluster around the truth.

the fact that learning that you’re in one of several clusters should *also* increase your confidence that you’ve just followed an opinion leader.

My main conclusion—that learning that you’re in a cluster should make you more confident that you’re right about everything, even if there are several clusters—doesn’t depend on this model’s specific mechanism—that of following opinion leaders—for generating other clusters. We could replace that mechanism with something else—some kind of groupthink or social conformity, say—that can generate clusters.¹⁶ What is important is just that there be a bunch of independent-minded people who are better than random at getting the right answer to any given question. No matter what else is going on, we expect those people to cluster, and to cluster around the truth, and so learning that you’re in a cluster should increase your confidence that you’re right about everything.

Let me briefly circle back to the arguments from Huemer, Joshi, and Lewis and Lewis discussed in the previous section. Lewis and Lewis think that the very orthogonality of the political questions means that neither cluster is likely to be right about all of them. Our model illustrates that while this may be true of individuals, it is not true of clusters. Take the simple version of the model, where the competence level is the same for each individual and for each question. Then, where p is that competence level and n is the number of questions, the probability of an individual being right about everything is p^n , which will be rather low unless p is large and n is small. (For instance, $0.6^{10} \approx 0.006$.) But conditional on being in a cluster, the probability of an individual being right about everything is much higher. For instance, conditional on being in the ‘core’ of one of two clusters—giving the sequence of answers which is most common among members of that cluster—the probability of being right about everything is around 0.5.

Huemer and Joshi think that if members of (the core of) one cluster are right about everything, then those of the other cluster must be anti-reliable, which is implausible. But in the simple version of the model, where competence levels are the same for everyone, this is clearly not the case. Everyone is reliable, and equally so. For even conditional on being a follower, one has that same competence level, which is greater than 0.5, since all of the opinion leaders whom one might follow have that same competence level, and since whether

¹⁶Weatherall and O’Connor (2020) present a model in which clusters arise as a result of people sharing evidence with their neighbors in an epistemic network, but distrusting evidence shared by those with whom they disagree on many issues. And Freeborn (2024) presents a model in which clusters arise as a result of people’s differing background beliefs affecting how their beliefs evolve in response to new evidence.

one follows (and if so, whom one follows) is probabilistically independent of what the right answer is to any given question.

Let me close this section with one final point. My opponents hold that learning that you are in a cluster should make you more confident that you have been irrational, and also that not all of your beliefs are true. I have rebutted the latter claim, showing that being in a cluster raises, rather than lowers, the probability that you are right about everything. But what of the former claim? Well, I conceded that learning that you're in a cluster should increase your confidence that you've followed an opinion leader. But in my model, it is hard to see why this should increase your confidence that you've been irrational. After all, in the simplest version of the model, everyone—including the potential opinion leaders—has the same competence level, and so you do equally well in expectation by following an opinion leader as by being independent-minded.¹⁷ Moreover, we could tweak the model so that potential opinion leaders have higher competence levels than others, in which case following an opinion leader might even be rationally required. In that case, learning that you're in a cluster should *decrease* your confidence that you've been irrational while still increasing your confidence that you are right about everything.

5 Clarifications

The reader may have some lingering questions about my argument and the model on which it rests. First, does my argument rely on the political questions being binary, i.e. having just two possible answers? No. It still goes through with non-binary questions, provided that people are still more likely to give the right answer than to give any alternative answer. To see this, suppose that the questions have six possible answers. Then, instead of thinking of things in terms of tossing biased coins, we can think of them in terms of rolling loaded dice, where each die is loaded so as to be more likely to land on the side corresponding to getting the right answer than on any other side. For if all dice are loaded in favour of the same side, then when everyone rolls some number of times, the sequence in which the dice always land on that privileged side is more probable, and hence has a higher expected frequency, than any alternative sequence.

¹⁷Perhaps following an opinion leader is irrational for other reasons. Perhaps deferring to moral testimony—testimony about moral, rather than descriptive, matters—is somehow problematic, for it precludes understanding or moral knowledge, or something of that sort. See Sliwa (2012) for critical discussion. But it isn't irrational in virtue of lowering your expected accuracy, for it does no such thing.

Second, does my argument rely on the assumption that there are more independent-minded people than followers? No. This assumption was not made in the previous section, and it is not needed. Of course, the exact shapes and sizes of the various clusters will depend on the percentage of the population who are followers, as well as the number of potential opinion leaders and people's competence levels. But provided that there are a bunch of independent-minded, better than random people, my argument goes through, for we expect those people to cluster, and to cluster around the truth.

Third, won't the opinion leaders themselves wind up close to the truth, leaving the clusters indistinguishable from each other? Well, even if this were true, it would not affect my conclusion, for the truth would still probably lie in a cluster, and so learning that you're in a cluster should still increase your confidence that you have the truth. But in any case, it isn't true. Take again the simple version of the model where the competence level is the same for everyone and for each question. Then, since potential opinion leaders are randomly selected, the expected percentage of questions that an opinion leader gets right is equal to that competence level. For instance, if the competence level is 0.6, then the expected percentage of questions that a randomly selected opinion leader gets right is 60%. And so, in expectation, follower clusters will disagree with the (core of the) truth cluster on fully 40% of the questions.

Fourth, does my model suggest that one of the two major American political parties is right about everything? Not quite, for they are surely not the only clusters we find in the US population. My model does suggest that wherever the truth lies, there will be a cluster of people around it. But the truth cluster needn't be especially big, prominent, or salient.¹⁸ Nevertheless, learning that your beliefs coincide with those of one of the two main parties should still increase your confidence that you're right about everything, unless you antecedently have reason to believe that neither of these two clusters is the truth cluster.¹⁹

Fifth, granting that learning that you're in a cluster should increase your confidence that you're right about all the questions, how should it affect your expected accuracy, i.e. your expectation of the percentage of questions that you

¹⁸If the independent-minded people are few in number, then the truth cluster likely won't be very big, and if they are only slightly better than random, then it likely won't be very prominent.

¹⁹As Daniel Greco (p.c.) suggested, perhaps the ways in which the two parties and their associated clusters have changed their positions over time gives reason to doubt that either of them contains the independent-minded people who have clustered around the truth.

got right? It should also increase it, at least in the simple version of the model where everyone has the same competence level. Here's why. Your unconditional expected accuracy is just your competence level. Your expected accuracy, conditional on being in a follower cluster, is the same, since opinion leaders have the same competence level as everyone else, including you. But your expected accuracy, conditional on being in the cluster consisting of independent-minded people who have clustered around the truth, is of course much higher. But these are (very probably) the only two ways of being in a cluster, and since one increases your expected accuracy while the other leaves it unchanged, overall your expected accuracy is higher conditional on being in a cluster than not.

Sixth, how should learning that you're in a cluster affect your confidence in your answer to any given question? Well, it should increase to nearly 1 your confidence in any answer on which the clusters agree. What about questions on which the clusters disagree? If you're in the core of a cluster but have no symmetry-breaking evidence—no evidence concerning whether you're in a follower cluster or the truth cluster generated by the independent-minded people—then your confidence in a given answer should roughly equal the percentage of the clusters which agree with that answer. (This conclusion is trivial, following from the assumed lack of any symmetry-breaking evidence.) So, if there are two clusters, and they disagree on some question, then your confidence in your answer should go to around 0.5. In the two cluster case, then, with no symmetry-breaking evidence, your confidence in each answer on which the two clusters agree should increase from your competence level up to nearly 1, while your confidence in each answer on which the two clusters disagree should decrease from your competence level down to around 0.5. Note that the same is true for anyone who isn't in a cluster but learns of the existence of clusters and what answers those clusters arrived at: without any symmetry-breaking evidence, their confidence in a given answer should roughly equal the percentage of clusters which agree on that answer.

Having said that, you might have some symmetry-breaking evidence. Most importantly, there is the first-order evidence bearing directly on the political issues. For instance, with the question about the minimum wage, there is evidence concerning the likely effects of a minimum wage on employment, equality, and economic growth, as well as evidence in the form of philosophical arguments about the importance of equality, freedom of contract, and so on. The first-order evidence isn't represented in my model, but it is nonetheless important. As Kelly (2005) has persuasively argued, the first-order evidence should

not drop out of the picture in cases of disagreement, being wholly superseded by the higher-order evidence having to do with who holds which beliefs. It could be, then, that you can rationally be much more than 0.5 confident in your answer to a given question, even when there are two clusters which disagree about it, since it could be that the first-order evidence nonetheless supports the answer that you arrived at.

Introspective evidence is another possible symmetry-breaker. If you introspect and find that you formed your beliefs by following an opinion leader, then in my model, this provides decisive evidence that you're in a follower cluster. And so if there are two clusters, you should then become nearly²⁰ certain of the answer that the *other* cluster arrived at. And if you introspect and find that you formed your beliefs by being independent-minded, this provides strong (albeit not decisive evidence²¹) that you're in the truth cluster, and so you should increase your confidence in the answer that you arrived at.

A final kind of possible symmetry-breaking evidence has to do with the shapes of the various clusters. Picture the various possible sequences of answers arrayed on plane, with sequences close to each other in space to the extent that they agree on more answers. Above the plane is a surface, where height above the plane at a given point represents the number of people who give the corresponding sequence of answers. In my model, the truth cluster will be conical, while any follower clusters will be cylindrical. To see the former, recall the analogy with coin tosses. If everyone tosses a coin which is biased towards heads, then the all-heads sequence has highest expected frequency. But it is closely followed by the various all-heads-but-one sequences, which are in turn followed by the many more all-heads-but-two sequences, and so on. In this sense, the cluster formed by people tossing biased coins will be conical, and similarly for the truth cluster generated by independent-minded and better than random people. By contrast, a given follower cluster will be cylindrical. This is because following is an all-or-nothing affair. If you're a follower, you defer to your opinion leader on everything. And so there will be many people who give exactly the same sequence of answers as that opinion leader. But nearby sequences—those differing on one or two or three answers—will likely be much less common, though

²⁰Why only *nearly* certain? That's because it's possible that both clusters include opinion leaders but that your opinion leader happened to get everything right, and hence that you're in a follower cluster which coincides with the truth cluster.

²¹Why not decisive? That's because it's possible, though unlikely, that you independently arrived at the same beliefs as did an opinion leader, and hence that you're in a follower cluster despite having been independent-minded.

their exact expected frequency will depend on the proportion of right answers that they contain. Now, this difference in shape between the clusters is an artefact of my model. We could modify it and make following an opinion leader a probabilistic matter, so that if you follow an opinion leader, your probability of deferring on any given question is greater than 0.5, and whether you defer on one question is probabilistically independent of whether you defer on any other. In that case, even the follower clusters will be conical. In any case, this is just a proof of concept that details about the shapes of the various clusters could provide some symmetry-breaking evidence to indicate whether you're in the truth cluster generated by the independent-minded people, or instead in one of the follower clusters.

Seventh, and finally, how well does my model predict the details of the clustering that actually observe in the US? How well does it predict the number, sizes, shapes, and compositions of the clusters that actually exist? I won't attempt to answer this question in any detail, but I suspect that the model is unrealistic in a great many respects. Is that a problem for my argument? Here is a somewhat concessive answer. My model still illustrates that multiple clusters *could* arise in an epistemically innocuous way, such a way that learning that you're in a cluster shouldn't increase your confidence that you've made a mistake or that you've been irrational in forming your beliefs. But I also want to give a less concessive answer. As noted in the previous section, all that I need for my main conclusion—that learning that you're in a cluster should actually increase your confidence that you're right about everything—is the assumption that there are a bunch of people who are independent-minded and better than random at getting the right answer to any given question. Provided that this assumption is retained, it doesn't matter what other bells and whistles are added to make the model more realistic and more predictive of the clustering we actually observe.

6 Conclusion

People cluster around certain sequences of answers to political questions. This is curious, since the questions seem to be orthogonal. Some think that this makes the clustering epistemically suspicious, and so if you learn that you fall into one of several clusters, this should increase your confidence that you've been irrational and decrease your confidence that you're right about everything. I

have presented a simple model on which this is false. In my model, learning that you're in a cluster needn't increase your confidence that you've been irrational. And it should increase, rather than decrease, your confidence that you're right about everything. This is because if there are a bunch of independent-minded people who are better than random at getting the right answer to any given question, we expect them to cluster, and to cluster around the truth.

References

- Boland, Philip J, Frank Proschan, and Yung Liang Tong. 1989. Modelling dependence in simple and indirect majority systems. *Journal of Applied Probability* 26:81–88.
- Bramson, Aaron, Patrick Grim, Daniel J. Singer, Steven Fisher, William Berger, Graham Sack, and Carissa Flocken. 2016. Disambiguation of Social Polarization Concepts and Measures. *Journal of Mathematical Sociology* 40:80–111.
- Christensen, David. 2007. Epistemology of Disagreement: The Good News. *Philosophical Review* 116:187–217.
- Condorcet, Marie Jean Antoine Nicolas de Caritat, Marquis de. 1785/1976. An Essay on the Application of Probability Theory to Plurality Decision-Making. In Keith Michael Baker, editor, *Condorcet: Selected Writings*, 33–70. Bobbs-Merrill.
- Converse, Philip. 1964. The Nature of Belief systems in mass publics. In D.E. Apter, editor, *Ideology and its Discontent*, 206–61. Free Press of Glencoe.
- Freeborn, David Peter Wallis. 2024. Rational Factionalization for Agents with Probabilistically Related Beliefs. *Synthese* 203:1–27.
- Goodin, Robert E., and Kai Spiekermann. 2018. *An Epistemic Theory of Democracy*. Oxford, United Kingdom: Oxford University Press.
- Haidt, Jonathan. 2012. *The righteous mind: Why good people are divided by politics and religion*. Vintage.
- Hedden, Brian, and Kevin Dorst. 2022. (Almost) All Evidence is Higher-Order Evidence. *Analysis* 82:417–425.
- Huemer, Michael. 2016. Why People are Irrational about Politics. In Jonathan Anomaly, Geoffrey Brennan, Michael Munger, and Geoffrey Sayre-McCord, editors, *Philosophy, Politics, and Economics: An Anthology*, 456–467. Oxford University Press, New York.
- Joshi, Hrishikesh. 2020. What Are the Chances You’re Right About Everything? An Epistemic Challenge for Modern Partisanship. *Politics, Philosophy and Economics* 19:36–61.
- Kelly, Thomas. 2005. The Epistemic Significance of Disagreement. In Tamar Szabo Gendler and John Hawthorne, editors, *Oxford Studies in Epistemology Volume 1*, 167–196. Oxford University Press.
- Lakoff, George. 1996. *Moral politics: How liberals and conservatives think*. University of Chicago Press.
- Lewis, Verlan, and Hyrum Lewis. 2022. *The Myth of Left and Right: How the Political Spectrum Misleads and Harms America*. Oxford University Press.

- Russell, Bertrand. 1912. *The Problems of Philosophy*. New York: Barnes & Noble.
- Sliwa, Paulina. 2012. In Defense of Moral Testimony. *Philosophical Studies* 158:175–195.
- Weatherall, James Owen, and Cailin O'Connor. 2020. Endogenous Epistemic Factionalization. *Synthese* 198:6179–6200.