



SERRC
Social Epistemology
Review & Reply Collective

<http://social-epistemology.com>
ISSN: 2471-9560

Open-Mindedness, Rational Confidence, and Belief Change

Katia Vavova, Mount Holyoke College, evavova@mtholyoke.edu

Vavova, Katia. 2023. "Open-Mindedness, Rational Confidence, and Belief Change." *Social Epistemology Review and Reply Collective* 12 (2): 33–44. <https://wp.me/p1Bfg0-7BH>.

Abstract

It's intuitive to think that (a) the more sure you are of something, the harder it'll be to change your mind about it, and (b) you can't be open-minded about something if you're very sure about it. If these thoughts are right, then, with minimal assumptions, it follows that you can't be in a good position to both escape echo chambers and be rationally resistant to fake news: the former requires open-mindedness, but the latter is inimical to it. I argue that neither thought is true and that believing them will get us all mixed up. I show that you can be open-minded *and* have confidently held beliefs, and that beliefs in which you are less sure are not, thereby, more fragile. I close with some reflections on the nature of rational belief change and open-mindedness and a brief sketch about what might actually help us in the fight against misinformation and belief polarization.

1. To be Open-Minded and to be Sure

Here's an intuitive thought: the more sure you are of something, the harder it'll be to change your mind about it.

Here's another: you can't be open-minded about something if you're very sure about it.

If the first one is true, and if you're worried about succumbing to fake news, then you should shore up your confidence in what you believe. But if the second one is true, then to shore up your confidence is to close your mind. So, we might conclude, to resist misinformation, we should cultivate close-mindedness.¹

But close-mindedness isn't just annoying, it's dangerous. Cults and echo chambers cultivate it, reinforcing conspiracies and discrediting outsiders. To avoid getting stuck in some such sinister social structure, we should cultivate open-mindedness.²

So, it seems, we should cultivate both open-mindedness (to escape echo chambers) *and* close-mindedness (to resist fake news and other misleading evidence).

You might be tempted to take a side by giving, say, an unlimited defense of open or close-mindedness; or, you might, reasonably, opt for something more measured.³ Unfortunately, each option offends commonsense; fortunately, we don't have to choose. For neither intuitive thought is true, and they're what got us all mixed up in the first place.

I'll first present, more carefully, the supposed puzzle. I'll then show that belief stability and strength do not correlate in the way that these thoughts suggest: beliefs in which we're highly

¹ Cf. Fantl 2021

² Cf. Fantl 2021. I'll later argue that open-mindedness is a double-edged sword.

³ E.g., a limited defense of the open mind (Fantl 2018).

confident—even highly *rationaly* confident—are not automatically more resilient against apparent counterevidence. In fact, they can be easier to dislodge than ones we are less sure about. Once we see this, we’ll also see that open-mindedness doesn’t preclude holding strong, sure beliefs. It is indeed difficult both to resist misinformation and to escape echo chambers. But that’s not because resisting the one requires succumbing to the other; it’s just because rational belief is hard. I end with a brief, tentative sketch of what might actually help us both resist fake news and escape echo chambers.

2. Fantl’s Puzzle

Jeremy Fantl endorses both intuitive thoughts. He argues that escaping an echo chamber requires remaining open-minded toward counterarguments but that resisting subtle fake news requires stable and therefore strongly held beliefs. But if you should always be open-minded, he argues, then you should never have strongly held beliefs. So, you can’t be in a good position to both escape echo chambers and rationally resistant to fake news (Fantl 2021, 1).

An *echo chamber* is a community whose members distrust outside voices. Being in an echo chamber doesn’t mean you aren’t exposed to contrary evidence; it’s not a mere *epistemic bubble*. Echo chamber members’ distrust of outsiders inoculates them against counterevidence (Nguyen 2020, 153). Paradigm examples of echo chambers are conservative communities, like the Tea Party or Rush Limbaugh’s followers. Limbaugh encouraged followers to distrust the “mainstream media”—those most likely to report negatively about him (Nguyen 2018). But echo chambers aren’t essentially conservative: my students consistently cite, as echo chambers, both the right-wing communities they left behind and their liberal campus social groups. Echo chambers are a problem for everyone.

Most of us hope to avoid echo chambers. We hope to be responsive to compelling counterevidence and change our minds. Open-mindedness seems crucial here. You’re *open-minded*, according to Fantl, if you’re willing to reduce confidence in response to counterarguments when (a) your belief is controversial, and (b) the counterargument is compelling and apparently flawless (Fantl 3). Fantl thinks such flexibility of belief entails a cap on our confidence: “If you should always be open-minded toward counterarguments, then you shouldn’t ever have strongly held beliefs” (8). Since embracing open-mindedness means taking counterevidence seriously, it should prevent counterevidence inoculation and allow good arguments to penetrate echo chambers.⁴ Fantl thus concludes that if you want to be able to escape echo chambers, you should be open-minded.

A *fake news* story is meant to resemble real news. Its creators aim for it to be widely disseminated and occasionally convincing while either knowing it is false or not caring (Fantl 1).⁵ Fake news often appears as clickbait—something you want to be true or find surprising yet, perhaps, plausible. Something you’d forward to your friends. Paradigm examples include

⁴ Though presumably also some trick arguments; more on this later, in the section on open-mindedness.

⁵ See Fallis and Mathiesen (2019) for an illuminating overview and critique of definitions of fake news.

Pope Francis’s endorsement of Trump (Ritchie 2016), Hillary Clinton’s selling of weapons to ISIS, and Pizzagate (I know, but people still forward it, and one guy believed it enough bring a gun to the pizza parlor (Samuelson 2016)).

Sometimes we respond to fake news with suspicion or outright disbelief; sometimes we fall for it. How we respond depends at least on how subtle and convincing the fake news is, what we independently find plausible, and what we want to believe. Fantl argues that the strength of your belief also matters. If you truly believe that 9/11 was a terrorist attack and not a government conspiracy, *and* your belief is strongly held, then you won’t be easily convinced otherwise.⁶ Fantl thus concludes that resisting fake news requires holding beliefs strongly. Hence, the puzzle:

(A) You should always be open-minded toward counterarguments (to exit echo chambers).

(B) You should sometimes have strongly held beliefs (to resist non-obviously fake news).

(C) If you should always be open-minded toward counterarguments, then you shouldn’t ever have strongly held beliefs.

Therefore, you can’t both have an echo chamber exit plan and be resistant to subtly fake news.

Fantl rejects A: we needn’t always be open-minded because echo chambers aren’t all bad.

I reject B and C because of false underlying assumptions about belief strength and resilience. To hold a belief strongly, in Fantl’s sense, is to hold it with high confidence. But, as I’ve argued elsewhere, highly confident belief isn’t necessarily stable belief, resilient in the face of counterevidence.⁷ This is why strengthening our beliefs won’t automatically help us resist fake news and also why open-mindedness doesn’t prohibit strongly held beliefs. Once we see this, we can embrace a better understanding of echo chambers, fake news, and the difficulties these phenomena pose. But first, why are B and C false?

3. Strength of Belief is Independent of Belief Resilience

The first intuitive thought is tempting: why wouldn’t strengthening beliefs make them more resilient? It’s not clear why, but counterevidence can require us to lose confidence even in beliefs in which we are rationally *highly* confident. Furthermore, the beliefs we hold with *lower* rational confidence can be *more* stable—more resilient to counterevidence. Thus, strength of

⁶ Fantl’s characterization here overlooks the fact that some fake news stories concern matters that we don’t yet have an opinion about. This is, already, a reason to doubt that Fantl’s account can be right. I set it aside for the sake of argument but return to a case of the relevant sort of agnosticism in the final section.

⁷ See Vavova 2014.

belief is sometimes *inversely* related to belief resilience: the stronger the belief, the more surprising and powerful potential counterevidence can be.⁸

STRONG. You're about to flip a coin a bunch of times. You look at it, turn it over in your hand, and become highly rationally confident that it has heads on both sides. You therefore become highly rationally confident that the *second* flip (as well as the first, third, fiftieth, and so on) will land heads. You're about as sure of this as you can be of anything, and it's rational for you to be so sure. You flip and the coin lands tails. You're surprised, obviously (shocked, plausibly). Something has clearly gone wrong. You're no longer at all confident that the coin will land heads on the second flip.

The coin landing tails on the first flip is evidence that the coin is not two-headed after all. Something—be it with your vision or with the coin—has gone wrong. Antecedently, your belief wasn't just strong, but *very, very* strong and *rationally* so. Yet, upon receiving just one coin-sized bit of counterevidence, your confidence should plummet: from almost 1 to about 0.5 (agnosticism).

WEAK. You're highly rationally confident that the coin you're about to flip is slightly biased toward heads—that it'll land on heads about 51% of the time. You are thus *weakly* rationally confident (~0.51) that, on the second flip, the coin will land heads. You flip the coin and it lands tails. This is not a surprising result, and it doesn't much affect your confidence in the result of the second coin flip. You remain moderately confident (maybe a smidge less than 0.51) that the coin will land heads on the next flip.

Tails on the first flip is evidence that the coin isn't as biased as you thought. Antecedently, you weren't particularly confident that it'd land heads on the second try: you were just above agnosticism. Your belief is rational but not strong. Yet it barely budes in the face of counterevidence.

Together, these cases show that high confidence—even rational high confidence—doesn't inoculate a belief from counterevidence and, therefore, revision. In fact, the opposite can be true. In other words: your belief can be strong—*rationally very strong*—and yet quite fragile. Weak belief, on the other hand, can be stable and resilient. So not only do belief strength and stability come apart, they are sometimes inversely related.

These cases are important for at least two reasons.

First, they show that Fantl's claim that strengthening beliefs makes them more resilient is false. It would be a mistake, then, to shore up our confidence in the hope that this will help us resist misinformation. Against the first intuitive thought, then, being very sure of something *doesn't* make it harder to change our minds about it.

⁸ These cases are adapted from my 2014. Fantl 2021 discusses them on pp 7-8. I discuss his response to them below, in section 4.

Second, they show that we needn't, in general, fear having low or moderate confidence in our beliefs. Doing so needn't make them vulnerable. Nor need we fear having strongly held beliefs: being highly confident of something doesn't mean that we'll get stuck believing it. Against the second intuitive thought, then, you can be open-minded *and* sure.

Being open-minded thus needn't mean being an epistemic pushover, with easy to dislodge beliefs. This is true even if being open-minded requires you to have low confidence in your beliefs since low confidence doesn't entail more vulnerable beliefs. But being open-minded doesn't require low confidence because high confidence doesn't entail dogmatism.

Thus, B is false: strength of belief won't protect us from fake news; and C is false: open-mindedness doesn't require weakly held beliefs.

What about A? Should we always be open-minded toward counterarguments (to escape echo chambers)? That depends on what exactly it means to be open-minded. I'll return to that question shortly. First, Fantl's response to these cases.

4. Fantl's Response

Fantl is aware of these results and doesn't dispute them. Instead, he says, "there are *other* cases in which it seems that higher confidence in *p* is *the reason why* you will be more resistant to counterevidence" (Fantl 7, my emphasis):

TRUMP. You read a genuine-looking news story that Trump wears a toupée: it's from a respected source and includes a video of bald Trump chasing the toupee on a windy day. You are antecedently rationally highly confident that Trump doesn't wear a toupée because you are Trump. You resist the story.

In this case, you are rationally highly confident in your belief about Trump's hair, and that belief doesn't budge in response to counterevidence. As Trump, running your hands through your naturally hirsute head while watching the video, you should remain confident that you don't wear a toupée. In other words, neither the convincingly executed video nor the controversial state of your follicles should move you. (Likewise, the controversy about O.J. Simpson shouldn't make him doubt whether he killed Nicole Brown; presumably *he* knows.)

Isn't Fantl obviously correct? Who could deny these descriptions?

These are certainly paradigms of stable, resilient beliefs held with rational high confidence. These beliefs are also resistant to counterevidence. It doesn't follow, however, that high confidence is what makes them resistant, nor that they'd be less stable if more weakly held. And this is what, as Fantl grants, he needs to show: why strongly held beliefs would be harder to sway *in virtue of* their being strongly held (7).

Fantl recognizes this and attempts to say more. Comparing **TRUMP** and **STRONG**, he notes that in **TRUMP**, "you have a high confidence that Trump doesn't wear a toupée *while*

being disposed to be unsurprised if [counter]evidence is produced...” (my emphasis). In **STRONG**, however, “your high confidence [...] is conjoined with a disposition to be very surprised...” Thus, your belief in **TRUMP** isn’t just *strong* but also *strongly held*. A belief is *strongly held*, he stipulates, “when your high confidence in the belief is conjoined with a disposition to find evidence against the belief [un]surprising” (7).⁹ To be strongly held, then, is for a belief to be difficult to sway (stable, resilient).¹⁰

But this doesn’t explain anything (certainly not the relationship between belief strength and stability). It merely redescribes the case. Yes, this belief is strong *and* resilient, but *why*? What is the connection between *strong* and *strongly held* belief? Without one, Fantl’s case collapses.

The second, bigger, problem is that, as we already know from **WEAK** that stable beliefs needn’t be strong.¹¹ But if my belief is stable, why care whether it is strong?

Of course, once you are disposed to find evidence against a belief unsurprising, why does it matter whether the belief is strongly held – at least from the perspective of resisting fake news? If you would find a news story indicating that Trump wears a toupée unsurprising, then that news story won’t impact your confidence very much, whether [...] high or low. So it would seem that strongly held belief is not necessary for resisting fake news (Fantl, 8).

To defend his position, Fantl needs a case of a belief that is stable *because* it is strong and not for some independent, unrelated reason. (**TRUMP** doesn’t work; it raises the very objection Fantl himself articulates.)

Fantl grants this, and in a last attempt to defend his position, he argues that strength of belief is independently desirable—that “we want not only to be unsurprised by fake news stories, we want to do so in conjunction with strong belief in the denials of those stories” (8). Since fake news is, by definition, false, to deny it is to believe the truth.¹² And, he argues, we want to be highly confident in the truth: we want our true beliefs to be strong, not merely true.

But why? If strength doesn’t make our beliefs more stable, what does it do? Asked out of context whether I prefer *strong* or *weak* true beliefs, I might choose strong—*ceteris paribus*. But all else is often not equal. In particular, we sometimes lack sufficient evidence for rationally strong beliefs in the truth. What we *want* is, ultimately, immaterial. I should believe

⁹ The ‘un’ is crucial here; the original article omits it. This is a typo (Fantl, personal correspondence). There is a similar typo in the following paragraph. Here it is corrected: “All else being equal, the more kinds of evidence you’d find surprising and the more surprising you’d find a given kind of evidence, the [less] strongly held the belief is” (7).

¹⁰ See his p. 7 and fn. 9.

¹¹ In my 2014, I explain how belief stability can arise at least partly from the weakness—that the two have a common cause. See especially section four.

¹² See Fallis and Matheson for a critique of the idea that fake news is always false (p. 4).

confidently when the evidence justifies it and weakly when it doesn't. But that's okay. It doesn't make me more susceptible to fake news.

4.1 Empirical Coda

Fantl briefly presents one other defense of the idea that belief strength increases belief stability: “The consensus in empirical psychology,” he writes, “seems to be that increasing what philosophers generally call ‘strength of belief’ makes those beliefs more resistant to counterevidence” (6). This seems to conflict with what we learned from the coin cases, but that appearance is misleading. Some terminology, then I'll explain.

In philosophy, belief strength usually refers to confidence or credence level: how sure you are that *p*. Psychologists refer to this as *attitude certainty*: “the sense of conviction someone has about an attitude” (Tormala et al 2006, 423). (Confusingly, they use ‘attitude strength’ to refer to durability or resistance to attack.)

There were indeed studies showing that increasing attitude certainty can make an attitude more durable (e.g., Tormala et al, 2006, Bassili, 2996, Babad et al, 1991), which prompted the *crystallization hypothesis*, that “increasing attitude certainty *inherently* makes an attitude more durable” (Clarkson et al., 810, my emphasis). Translation: the stronger your belief, the more stable and resistant it should be to counterevidence—because of that increase in strength. “This has been the dominant, if not only, view of attitude certainty in *past* research” (810, my emphasis). But these authors, whom Fantl cites, go on to reject this hypothesis.

Being highly confident in what you believe doesn't *always* make your belief more stable, they argue. Rather, becoming more confident in what you believe just magnifies characteristics the belief had before. In particular, if the belief was already a stubborn one, increasing confidence in it will make it more stubborn; if it is, for whatever reason, a more flexible belief, then becoming more confident in it will make it more flexible and less durable. More carefully and in their own words, the authors argue that attitude certainty

... amplifies the dominant effect of the attitude [...]. If the dominant effect of an attitude is to be resistant to change, for instance, increasing attitude certainty should increase that attitudes resistance, as in past research. If the dominant effect of an attitude is to be susceptible to change, however, the amplification hypothesis proposes that increasing attitude certainty might increase that attitude's susceptibility. Thus, under some conditions, amplification might produce effects that look like attitude crystallization (e.g., increased attitude certainty leading to increased resistance), but under other conditions, amplification would produce effects that directly counter the idea of attitude crystallization (e.g., increased attitude certainty leading to decreased resistance) (Clarkson et. al., 812).

In short, this new *amplification hypothesis* suggests that the effect of increasing attitude certainty “depends on other salient aspects of those attitudes” (812).¹³ This is why sometimes, but only sometimes, increasing belief strength, or attitude certainty, increases belief stability, or attitude durability. Other times, increasing belief strength makes beliefs *less* stable.

Shoring up our confidence in our beliefs thus won’t necessarily make them more resistant to fake news. This is what both the coin cases *and* the empirical literature show. Strengthening of belief can *sometimes* increase belief stability. This makes sense: strong belief non-accidentally correlates with a wealth of evidence. But we’ve also seen that a wealth of evidence doesn’t always protect your attitude, as in the coin case. And belief can also be strong for all sorts of bad reasons: because it is personal, stubborn, or insulated from good evidence (perhaps it’s at the center of a conspiracy theorist’s web of belief). So, we need to know much more about when, how, and why strengthening a belief makes it more stable. At this point, we can’t conclude anything about what manipulations of confidence could help us resist misinformation. And crucially, again, we shouldn’t be thinking this way about what we ought to believe. It’s worth repeating: what we *want* is, ultimately, immaterial. We ought to believe as our evidence supports, not as it would be beneficial to believe (even if the benefit is retaining true beliefs).

5. Upshot

Zoom out and consider the lesson about strength and stability.

The coinflip cases show that Fantl’s claim that stronger beliefs are more resilient is false. It would be a mistake, then, to shore up our confidence in the hope that doing so will help us resist misinformation.

They also show that we needn’t fear having low or moderate confidence in our beliefs; it won’t make them vulnerable, easily dislodged by the evidential equivalent of a light summer breeze. Nor need we fear having strongly held beliefs: being highly confident in something doesn’t mean getting stuck believing it.

This shows that being open-minded needn’t mean being an epistemic pushover with easy to dislodge beliefs. This is true even if being open-minded required you to have low confidence in your beliefs. But it doesn’t, because high confidence doesn’t entail dogmatism.

Any remaining discomfort arises, I think, from an unfortunate fact about rationality. Insofar as we aim to be rational, we open ourselves to being misled because to believe rationally is to

¹³ Clarkson et al. explore one of those aspects: whether an attitude is ambivalent or univalent. They argue that increasing attitude certainty increases durability of univalent attitudes but decreases durability of ambivalent attitudes. It is worth thinking more carefully about this in order to see how their results might apply to the relevant fake news and echo chambers cases, but doing so is not straightforward. This is the clearest example the authors give to help us understand ambivalence and its relation to certainty: “an individual might be highly certain of both the positive (e.g., tastes good) and the negative (e.g., high in calories) features of chocolate, thus feeling certain of his or her ambivalent attitude toward the treat” (811). It’s difficult to know how to translate between an ambivalent attitude toward a comestible and the belief that a conspiracy theory is true or false.

believe what the evidence supports. But evidence can be subtly and convincingly misleading (just as fake news can be subtly and convincingly news-like). And so, in trying to be rational, we may be misled. This is just the cost of rational belief for fallible creatures like us.

We can see this more clearly if we translate Fantl’s original puzzle into one about open-mindedness. Plausibly, being more open-minded makes it more likely that you’ll change your mind. This is why open-mindedness makes escaping echo-chambers easier. But doesn’t open-mindedness *also and at the same time* increase the likelihood of *succumbing* to echo chambers?

Suppose I’m a member of Architects & Engineers for 9/11 truth, who think that 9/11 was a government conspiracy. Escaping this echo chamber necessitates taking seriously evidence that 9/11 was *not* a government conspiracy. If I’m open-minded, I won’t dismiss such evidence out of hand, and, if I can’t find flaws in it, I should change my mind. So open-mindedness really could help me escape this echo chamber.

But suppose I’m open-minded and *not* a member of Architects & Engineers for 9/11 truth. I read their pamphlet, can’t see anything obviously wrong with the argument, and so become slightly less confident that 9/11 wasn’t a government conspiracy. I open-mindedly attend a meeting to learn more and become completely convinced that 9/11 was a conspiracy.¹⁴

Open-mindedness is thus a double-edged sword. Perhaps, ideally, we’d want to be open-minded toward the truth but not toward falsehood.¹⁵ But open-mindedness isn’t asymmetrical. The very thing that makes it a virtue is what also prompts us to trust convincingly misleading evidence. It’s a virtue that, as Baldwin said in another context of all virtues, is “ambiguity itself” (30).

This is the root of Fantl’s puzzle. It’s not really about whether we should hold beliefs strongly or with an open mind. It’s about the fact that being rational requires responding to evidence, and misleading evidence is evidence. So, being rational can lead us away from the truth. In the end, that’s not puzzling; it’s just the unfortunate, inevitable cost of rational belief for fallible creatures like us.

6. What Hope There Is

If these reflections are correct, then what hope is there for believing truly? Is it just a matter of luck—of chancing on good, not misleading, evidence? No, we can do more by supplementing our rational habits with critical thinking skills.

In training my students to think critically, I arm them with a set of tools for seeing through confusion and nonsense—whether their own or someone else’s. This is what we aim to do

¹⁴ Cf. Senior 2021; about Bobby’s dad we want to say that perhaps he should have been a little less open-minded.

¹⁵ This, I think, is part of the temptation behind the thought that maybe echo chambers aren’t always bad; see Fantl and Lackey.

when teaching students how to spot fallacies, distinguish between correlation and causation, and see past popular science headlines. It's also what we hope to do when teaching close reading of difficult texts and why we urge them to charitably interpret even the most inane or offensive arguments: so they can see through nonsense and slay dragons, not take out the trash.

Critical thinking skills make students less gullible and more likely to reject an idea (hence 'critical'). Of course, being dogmatic also makes us likely to reject ideas, but it's better to be *curious*. Critical thinking starts with a "Hmm... that's interesting..." and continues with, "Could that be right?" Teaching my students to think about their own thinking and to approach ideas, whether coming from inside or out, with curiosity allows them to see that certain views aren't inevitable. That's the first step to dismantling the ones that are wrong.

We could say, then, that in training our students to think critically, we arm them with *bullshit detectors*, which better position them to defend against misleading evidence, whether it comes from the media, the news, their teachers (me included), or themselves.¹⁶ A well-honed bullshit detector strengthens their belief defense and offense: it is thus a crucial tool in the battle against misinformation and polarization. Let me explain.

First, a bullshit detector better positions you to resist echo chambers, whether you stand inside one or not. It can help reveal the subtle grooming that bosses like Limbaugh engage in when epistemically and emotionally isolating followers. Sometimes, what they're trying to convince you of will mirror something you're already inclined to believe, like that most major news outlets are biased; other times, it'll conflict with your beliefs, like that you can trust your friends and family. Either way, the bullshit detector can be a force for good, tempering both your own suspicions and the ones you might adopt.

Similarly for fake news. *Subtle* fake news is, by definition, hard to identify. A bullshit detector's job is to detect bad arguments and evidence. If the fake news conflicts with your belief that Obama was born in the United States, then the bullshit detector plays *defense*. If, instead, the fake news is something you already believe, or want to believe, the bullshit detector plays *offense*: opening your eyes to your own nonsense.

Of course, the best defense and offense sometimes fail; strengthening them is no less worthwhile.

Notice that having a good bullshit detector can make your beliefs more resilient and stable: less movable. But you can have a good bullshit detector and the accompanying epistemic stability without being highly rationally confident in your beliefs. (Unless it's your own bullshit you're detecting; then the detector will, hopefully, destabilize your beliefs.)

But you needn't even have an antecedent opinion for the bullshit detector to be of use.¹⁷ Suppose that I have never considered whether Obama called Trump an idiot, but I see a

¹⁶ I mean bullshit here in a looser sense than Frankfurt's.

¹⁷ Cf. fn. 6.

(unbeknownst to me) deepfake video of him saying it. The video is *so* convincing (Sample 2020). Do I fall for it? Depends. Do I know that there's such a thing as deepfakes? Do I realize that believing Obama said that would play into his opponent's hands? If I answer yes to both questions, then my bullshit detector will probably go off. (And notice: whether I know these things itself depends on how curious I am.)

If I don't know about deepfakes, it'll be harder to resist this evidence, but my bullshit detector could still lessen its effect. If I know saying this would be politically costly for Obama, and I consider the actions of someone with his self-control and attention to public image, then I might start to wonder. That moment of wonder is the first moment of critical thinking. It's the first beep of the bullshit detector.

Acknowledgements

I'm grateful to Hilary Kornblith, Sam Mitchell, Alejandro Pérez Carballo, Gillian Steinberg and three anonymous referees for helpful comments on earlier drafts. I wrote this during my 2021-2 sabbatical. Thank you to Mount Holyoke College and the Center for Advanced Study in the Behavioral Sciences at Stanford University for support during that year.

References

- Babad, Elisha Y. Ayala Ariav, Ilana Rosen, Gavriel Salomon. 1987. "Perseverance of Bias as a Function of Debriefing Conditions and Subjects' Confidence." *Social Behavior* 2 (3): 185–193.
- Baldwin, James. 1956. *Giovanni's Room*, Laurel: Dell Publishing.
- Bassili, John N. 1996. "Meta-Judgmental Versus Operative Indexes of Psychological Attributes: The Case of Measures of Attitude Strength." *Journal of Personality and Social Psychology* 71 (4): 637–653.
- Clarkson, Joshua J., Zakary L. Tormala, and Derek D. Rucker. 2008. "A New Look at the Consequences of Attitude Certainty: The Amplification Hypothesis." *Journal of Personality and Social Psychology* 95 (4): 810–825.
- Fallis, Don and Kay Mathiesen. 2019. "Fake News is Counterfeit News." *Inquiry* 62 (9-10): 1033-1065.
- Fantl, Jeremy. 2021. "Fake News vs. Echo Chambers" *Social Epistemology* 35 (6): 645-659.
- Fantl, Jeremy. 2018. *The Limitations of the Open Mind*. Oxford University Press.
- Frankfurt, Harry. 1988. "On Bullshit." In *The Importance of What We Care About* edited by Harry Frankfurt, 117-133. Cambridge University Press.
- Lackey, Jennifer. 2018. "True Story: Echo Chambers are Not the Problem." *Morning Consult* 19 November. <https://morningconsult.com/opinions/true-story-echo-chambers-not-problem/>. Accessed 8 April 2022.
- Nguyen, C. Thi. 2020. "Echo Chambers and Epistemic Bubbles." *Episteme* 17 (2): 141–161.
- Nguyen, C. Thi. 2018. "Escape the Echo Chamber." *Aeon* <https://aeon.co/essays/why-its-as-hard-to-escape-an-echo-chamber-as-it-is-to-flee-a-cult>. Accessed 8 April 2022.

- Ritchie, Hannah. 2016. "Read All about It: The Biggest Fake News Stories of 2016" *CNBC* 30 December. <https://www.cnn.com/2016/12/30/read-all-about-it-the-biggest-fake-news-stories-of-2016.html>.
- Sample, Ian. 2020. "What are Deepfakes – And How Can You Spot Them?" *The Guardian* 13 January. <https://www.theguardian.com/technology/2020/jan/13/what-are-deepfakes-and-how-can-you-spot-them>.
- Samuelson, Kate. 2016 "What to Know about Pizzagate, the Fake News Story With Real Consequences" *Time.com* 5 December. <https://time.com/4590255/pizzagate-fake-news-what-to-know/>.
- Senior, Jennifer. 2021. "What Bobby McIlvaine Left Behind." *The Atlantic* September 2021. Online 9 August 2021.
- Tormala, Zakary L., Joshua J. Clarkson, and Richard E. Petty. 2006. "Resisting Persuasion by the Skin of One's Teeth: The Hidden Success of Resisted Persuasive Messages." *Journal of Personality and Social Psychology* 91 (3): 423–435.
- Vavova, Katia. 2014. "Confidence, Evidence, and Disagreement" *Erkenntnis* 79 (S1): 173–183.