



## Opinion Science Podcast

*Hosted by Andy Luttrell*

### Episode #13: Fake News with Gordon Pennycook

July 6<sup>th</sup>, 2020

Web: <http://opinionsciencepodcast.com/>

Twitter: [@OpinionSciPod](https://twitter.com/OpinionSciPod)

Facebook: [OpinionSciPod](https://www.facebook.com/OpinionSciPod)

#### **Andy Luttrell:**

There's a nasty virus making its way around the world, quietly jumping from person to person, sometimes leaving a wake of real destruction behind it. It's the novel coronavirus. You might have heard of it, and in our digital age, where information can move even faster than a virus, we're stuck in a confusing web of truths, half truths, and flat-out lies about this virus and its effects. Some articles discuss the science of face masks as a way to prevent spreading the virus. Others claim that Bill Gates will be injecting microchips into a vaccine in order to track our every move. And I'm really hoping you know which of those was the example of fake news. And it's obviously not just about health information. The 2016 election was a headache-inducing example of how quickly totally false information can spread online.

So, how do people navigate these misinformation minefields? Do they even care about the truth anymore? It might seem like there's no hope, that people are just too biased and irrational to sort fact from fiction, but luckily we might have reason to be a little more optimistic. You're listening to Opinion Science, the show about our opinions, where they come from, and how they change. I'm Andy Luttrell, and if you want to have a sound opinion, it needs to be based on good information. But do people know where to find that? I talked to Gordon Pennycook. He's an assistant professor of psychology at the University of Regina, and for the last several years he's been studying these questions. We'll talk about what fake news and misinformation are, why people share false information, and whether there's anything we can do to halt the spread of this dangerous content.

#### **Andy Luttrell:**

So, I was thinking that we could start by, if you wouldn't mind, giving like a background definition of what you mean when you say fake news or misinformation. To sort of contextualize all the stuff that you've done. What is it that we're really talking about?

#### **Gordon Pennycook:**

Right. So, fake news is a very specific form of misinformation I think, and part of the reason I got interested in it was because it's kind of a particularly egregious form of misinformation. So, a fake news headline is a headline that is an ostensible news headline that is just entirely made up. It's not just misleading, or sort of false, it's completely made up, like the Pope endorsed Donald Trump or something. That never happened. And so that's, whenever we use fake news in papers, I'm

talking specifically about fabricated news headlines, which is really, like I said, an egregious form of misinformation.

Of course, misinformation as a category is just things that are false, and some things that are misinformation aren't even intentionally false, like I guess fake news is a kind of form of disinformation, for that reason. It's intentionally deceptive. But misinformation doesn't have to be intentionally deceptive, and of course if someone's sharing a fake news headline that they don't realize is false, it's a form of kind of misinformation coming from that person. But ultimately, still fake news.

**Andy Luttrell:**

So, is it intention? Is that really kind of the seed that distinguishes fake news? That it is clearly and intentionally false.

**Gordon Pennycook:**

Yeah.

**Andy Luttrell:**

Whereas misinformation might be intentional but could just be a misunderstanding.

**Gordon Pennycook:**

Yeah. What people distinguish between misinformation and disinformation, usually intentionality is the key thing. Of course, it's almost impossible to determine that. Certainly, if you're just looking at the content. I mean, I guess you just assume that someone didn't unintentionally make something up, like you just know that it's not... There's no reason at all to believe to believe the thing, and so probably it was intentionally made up to be false. So, it's a bit of a sticky issue. This is why we usually just use misinformation as a kind of broad category, and then fake news is just kind of a test case, like a thing that we can use to kind of take the whole world of potential falsehoods and narrow it down to some category that we can investigate and try to find solutions to, and all that kind of stuff.

**Andy Luttrell:**

So, fake news is also definitely a term that has at least become really familiar in recent years. Is it your sense that that term is a product of recent events, or have people been using this idea of fake news for a while?

**Gordon Pennycook:**

Well, the actual, like the internet search history thing, it was not commonly used before the 2016 election. And there's a few kind of reasons that it became big. Part of it was because it was definitely a new thing, that the specific case where people were creating false headlines and putting them on... What they do is they just create a website that looks sort of like a pretty weak and not very good news website. It doesn't take that much effort to do that. You just kind of put things in the right spot on the website. And then you just start making things up. Make up your own content. Some of it might get spread on Facebook. If you get clicks, then you get ad revenue, and because there was so much going on with the election, there was a lot of stuff that was being created to take advantage of that. And so, that was a specific kind of new thing that came up in the realm of

misinformation. That's why the term was being used, because it was just... They were faking news headlines, and so it made sense.

And then, of course, after that, the term started being used as the president, Donald Trump, uses it. To say, "This is something I disagree with." And then, so if you ask any journalist about it, when they write an article with any modicum of controversy, most of the comments will say, "Oh, this is fake news," or something. Right? Which, of course it isn't, because it's not made up. So, it doesn't mean it's false or inaccurate. It means it's made up. So, it's used incorrectly probably more than it is used correctly these days, unfortunately.

**Andy Luttrell:**

So, when it comes to trying to understand how this kind of content spreads, and why people would believe it and pass it along to others, was there... What did we know at the time, right? Because you could say that, "Oh, this term fake news came up, but psychologists have already been basically studying this sort of thing before. We already had it nailed." Obviously, we didn't, because you've done a lot of work since then, but what did we know at the time this notion was bubbling up, but what was then left open for you to swoop in and help understand?

**Gordon Pennycook:**

Yeah, that's a good question, because the funny thing about it is that... So, there's lots of work on things that are related, like conspiracy theories, for example. Of course, it's very different than that. Especially conspiratorial kind of like ideation, like thinking about things as conspiracy theories, and like caring about them and doing research and all that kind of stuff, the kind of like typical conspiracy theory, what you expect someone in the basement with all the pictures on the wall or whatever.

It's a very different thing, because it's a passive form of being exposed to falsehoods. And there really wasn't actually a lot of work on that. Even forgetting about the fake news specific thing, but just people passively being exposed to misinformation. And I think the kind of... I don't know, positive thing, I guess, about the fact that fake news becomes such a big story is that it kind of got psychologists to realize that there's a lot of deliberate falsehoods out there that aren't as easily, let's say characterized, by being conspiratorial, or being a sort of type of belief. Of course, there's related research in ideology and all that kind of stuff, but the people believing falsehood, it wasn't something that was done a lot. Of course, there's lots of work in the realm of communication or journalism and other fields that have tackled these sorts of issues, but the underlying psychology of why people believe falsehoods just... I don't think psychologists like to say that, like make claims about what's true and false, you know what I mean? I think that was part of the big thing, is that we had something now that was so obviously false and made up that we can... We don't have to worry about telling participants that this is wrong, or this is right. Now it's already there for us, and so that I think might be part of the thing that opened the door.

But I'm just... That's a retrospective account. I have no...

**Andy Luttrell:**

Well, one of the things about that that I'm wondering is how much it matters that it is false, right? So, like as a person who studies persuasion, forever we've looked at how people approach information and evaluate that information, and whether they're going to use that information to

shape their viewpoint, and so in some sense, you could say like, “Well, by the time it gets to the person, it’s just information in front of them.” Right? Its actual truth is irrelevant, so I’m curious why you might say that it is relevant.

**Gordon Pennycook:**

That’s a great point. So, the way that I’ve been thinking about the fake news research is that we want to characterize something that’s occurring in the world, some sort of class of information, okay? And of course, we have... You can’t talk about people, whether they believe true or false stories, without talking about basic information processing, or talking about politics, or all the things that we’ve been talking about in psychology for decades, obviously. But if you want to try to understand a specific problem, something that’s happening in the world, you have to find some way to build a representative, or fairly representative kind of set of what that thing is, so that you can use it in studies to show people.

And so, it’s not really about why it matters that it’s false, per se. It’s to what extent can we take what we know already in psychology and apply it to this thing that has now arisen as a kind of new thing out there in the world? And so that, what we got in terms of some of the pushback for the first studies was that... As an example, I think it was the first paper that was published. I can’t exactly remember, but one of the first ones was that we showed that a single prior exposure to a fake news headline increases later belief in the headline because of the effect of familiarity. Of course, this is not at all a new phenomenon, right? The first paper on repetition for trivia statements was 1977, and so when some of the feedback we got was like, “Wow, this isn’t new.” It’s like it doesn’t have to be new. What matters is that it’s important and correct, you know what I mean?

The fact that there’s other evidence in other domains that supports this is a good thing, but it does tell us that we really need to worry about cutting off exposure to false headlines, because even things that are inconsistent with people’s ideology, like an anti-Hillary Clinton thing for Clinton supporters, is being believed more based on repetition alone. And so, to us that was an important finding, even though it’s not... The whole point of all of this is not to create a new psychology of fake news, but to create a psychology of fake news based on what we already know, and then learn some new things along the way, of course.

**Andy Luttrell:**

Yeah. I was curious what the fire under you is, if it’s sort of like just a cognitive interest of like, “Oh, this is interesting that people approach information this way,” or like, “Well, this is a real problem for the whole world, and so we need to sort of capture it right now.”

**Gordon Pennycook:**

That’s a great question. It’s kind of both, I guess. Maybe I’ll tell the backstory about how it happened in the first place. I mean, I’ve always been interested in people’s beliefs, and did work on mostly beliefs that... You know, you do psychology on what you don’t know and what you do know, kind of like I want to understand why people think things that are different than I do, you know what I mean? So, I did work on things like religious belief, and that was my first real interest in psychology, was like why do people believe when I don’t, basically. And then I just got more interested the kind of like foundations of belief, and religion was just my kind of step into that.

And so, I started my postdoc with Dave Rand, and at the time, he was doing almost all cooperation research, which I had no... I don't particularly care about that line of work and I'll be honest, and Dave knows this, I didn't read hardly any of those papers. I don't know what's going on there. I was like an outsider in the lab, basically, but I just knew that we would find something that we'd be both interested in, and then the election happened and all this fake news stuff was happening and I had published fairly recently before that a paper on pseudo-profound bullshit and why people think that random, abstract sounding sentences, like hidden meaning transforms unparalleled abstract duty, are profound. And so, there was a kind of obvious kind of intersect there, and also there was... Part of it was just the curiosity of what's going on.

But as it kind of got more and more developed within the first few months, the problems with society and people believing false things started becoming more real, and then as the research has developed over the years, we focused a lot more on interventions, like really more kind of applied work, really. Which, you always think when you write a grant, you're like, "Wow, this is what's going to happen, is that you're gonna take all these fundamental insights that we learn, and we're gonna come up with things to whatever." I never actually did this. I never even considered doing it, really. And then it just kind of... I don't know. I just felt like there was change, and that that was the thing that's been really driving a lot of the recent work, to try to actually do something to improve things.

**Andy Luttrell:**

Yeah, so let's scooch back a little bit and have you talk about what do we know about why people are open to things that are false, and who is most open to things that are false, and is it really... Is it just that it's convenient? If I go, "Oh, I ran across this information, fits what I want to think, I'll believe it." Right? But where do you come out on that?

**Gordon Pennycook:**

So, it's like just transporting myself back to 2016, 2017, when all of this stuff was happening, of course there were dozens and dozens of op-eds, and the primary claim was basically that this was a story about politics. The fact that people are believing and sharing fake news headlines is just another indication that we're a fractured, partisan society, and people are looking at things through their partisan brain and so on. And an academic who really crystallized that was Dan Kahan, who's a law prof actually at Yale, and he had written a paper saying that basically people are motivated consumers of misinformation. That's actually a direct quote from the paper. I've used it in so many talks I can now directly quote the paper.

That people kind of believe things because they want to believe them, and the account that comes from that is that it's what Dan has called it, identity protective cognition. Most people would just kind of think about it as motivated reasoning, where it's actual kind of explicit reasoning that people are using. They're kind of like using their capacity to engage in thinking to kind of delude themselves into believing what they want to believe. And then people share things that are politically consistent with their ideology, because it's good for them and that's what they want to do. They don't really care about what's true and false, they just kind of want to make some sort of partisan point or whatever.

And basically, the first wave and most of the work that we've done is just completely against all of that. Basically, what we show is that people who are more reflective, or if you get people to stop

and think, or if you do the opposite thing where, like get people to be more emotional, what you basically see is that people believe false stuff because they don't think about it that much. It coincides with their intuitive beliefs, which of course are going to be influenced by whatever partisan beliefs and ideology they have. But people who stop and reflect a bit more, they don't believe it as much, and they don't share it as much, and they are better at discerning between what's true and false, and it doesn't matter if it's consistent or inconsistent with their ideology, and it's true for both Democrats, Republicans. People who think more are just less likely to believe false things. Which again, is not some big, surprising finding, but it's just completely counter to this idea that people are engaging in motivated reasoning. They're just more accurate. It's just like classic information processing. They can figure out what's true and false.

And then what we've been showing in some more recent work is that part of the reason why people seem to share false content, fake news, and other things of that nature, is that they just forget, basically, or don't... They neglect to think about accuracy. They don't even think about it. In a paper that we did just recently on the COVID-19 fake news content, if you ask people in one group to rate whether they think it's accurate, in this case we just used yes or no, do you think this headline's true. That's not the exact wording of the question, but it doesn't really matter how you ask it. If you ask them to indicate whether they think it's true or false, they believe the true content 65, 70% of the time, which isn't great, but they are believing it the majority of the time. And the false content they only believe 25, 30% of the time. Something like that. Usually that's what we find for other content, as well, so often around 20%. It's a bit higher for the COVID stuff.

So, they're discerning between true and false, like there's a big gap there. They believe the true stuff way more than the false stuff. But if you ask them instead, like a different group of people, would you consider sharing this on social media? Much smaller difference. In fact, they would say they shared... This group, the same headlines, 35% of them. The other group, if you asked them which ones do you think are true, they only said 25% were true, right? Which means for sure, some people are sharing things that they could identify as being false if they bothered to stop and think about it, right? And so, that seems to be the bigger part of the problem to me, is that of course people share things that are aligned with their partisan ideology and all that kind of stuff, but the thing that we really try and emphasize and that we can actually help remediate is that people just don't stop and think about things. Whether or not they're partisan, they aren't reflecting enough. And part of that is probably because of the nature of social media, you know what I mean? It doesn't incentivize truth. It incentivizes other things besides that.

**Andy Luttrell:**

So, the cases where people could say that this is false, but are not directed to evaluate its truth, and then who share it, is the idea that there's just like a default assumption that it's probably true? Or is it really just like, "I do not care whether this is true or false, I would like to share this information."

**Gordon Pennycook:**

Well people, the thing I like so much about this is that it makes the least sense for experts and the most sense for people who aren't experts. Because people who are scientists, or even journalists when I talk to them about it, they're like, "How could you not care the most about what's true?" And in fact, if you ask people directly what's the most important thing when it comes to sharing, they say that it's accurate. They rate that more than whether it's consistent with ideology, whether

people will like it. Most everyone kind of really adopts that as a thing that is important, but what we don't appreciate is that people aren't in that sort of mindset when they're engaging with content on social media. I mean, maybe for scientists you don't get out of it, you don't stop ever thinking about whether things are true, because that's something that you build up and get used to. That's the kind of training that scientists get is to assess things and to make sure that you're not being led astray.

And given that almost everyone agrees that it's important to think about accuracy, I don't think that there's many people... I mean, of course there's some people who are sharing things because they're trolling, or they're whatever, but the majority of people care about accuracy, but they neglect to think about it. Think about what it's like on Facebook. You often, what people do is they go on to shut their brains off, right? They're taking a break from work. You're just scrolling through and it's like mostly pictures of dogs and babies and whatever, and then you see a news headline and you're like, "Oh, wow. Look at that." And then of course they don't read it, they don't click on the headline, they don't do anything else, and then what we're showing is they don't even think about whether it's true, and then they share it and then they move on with their lives, and that's it. It is what it is. And so, that seems to be at least a reasonable kind of caricature of what happens much of the time on people's Facebook feeds.

**Andy Luttrell:**

Do you have a sense of what they are using to decide what to share? Because in the scenario you described, they're not sharing all the dogs, they're not sharing all the babies, and they're not sharing all the news. And maybe this is just the other side of the coin that you haven't really zeroed in on yet, but it just raises the question, like if it's not about accuracy, what is it about that's making people do that?

**Gordon Pennycook:**

Yeah. Right. So, it is the other side of the coin. We mostly just want to get people to focus more on accuracy, and then all the other things... There's a myriad of other things. You know, like will people like this? How does it make me look? In many cases, does this reinforce some sort of political aim I have or whatever? But if you look at... What we're doing for the studies is we have now, just a bit of backstory for how we run these experiments, is we take actual true and false headlines from the world, and then use those in the studies. Of course, we could have a bias set that I just... If I'm the one selecting them, then I might select ones that are weird or whatever, and so what we do is we have a series of large pretests. We've done now I think five or six of them, where we take a set of headlines. The last one had like 74 false headlines and 74 true headlines. And some hyperpartisan headlines, which aren't really true or false, but they're misleading. And then you present them to huge number of people and they all just rate a subset of them, and so what you end up with is a big set of a bunch of different ratings, basically norms for news headlines. And then that's how we select which headlines to put in our studies. We match them on partisanship, and we do that, and then they come from that bigger set, and then we can vary them from study to study and all that kind of stuff.

But from all that data, we now have a bunch of different ratings on different elements of what the headlines are, and we know how much people would want to share them. And so, some of the things that are most impactful for sharing are things like whether they seem important, and do they cause an emotional reaction, would they be something that would be outrage provoking?

Partisanship is obviously a big one, although that sort of intersects with importance, obviously. And what else? We thought that whether they're funny and entertaining would have a bigger impact than it does, but that's mostly maybe because of the set that we have. They're all political things, and so importance is a salient aspect of that. But of course, people do share things that are funny.

**Andy Luttrell:**

And so, it sounds like most of this is political, right? You're saying?

**Gordon Pennycook:**

Yeah, the stuff those... The pretests I'm talking about is all political content. I mean, some cases we have some neutral things, but there wasn't a big component of the pretest. We don't have enough there to say anything definitive about it.

**Andy Luttrell:**

So, I'm curious. Just to go back to the idea that the accuracy perceptions, once you sort of draw people's attention to them, seem not to be biased by political attitudes. Is that really... How strong? Because we know that motivated reasoning happens, and we know that people will sort of use their predisposed opinion to color how they evaluate other stuff, and so how do you square that, what you're finding, with what we've seen before?

**Gordon Pennycook:**

Right. So, it's slightly more nuanced than that, because people do definitely believe things that are consistent with their ideology more than those that are inconsistent. But it doesn't mean... Depends what you mean by motivated reasoning. It doesn't necessarily support motivated reasoning with a capital R, and so this is work that we're doing with Ben Tappin, who is a postdoc in Dave's lab at MIT. And so, if you're going to interpret whether something's true or false, you have to by necessity bring to bear whatever knowledge you have to do that assessment, right? And so, imagine someone who only watched Fox News. Maybe by accident of history, like that's what their parents watched. They don't really care about, but that's what they're exposed to. Versus someone who only listened to MSNBC. The range of things that are going to be plausible for them in the political realm is going to be way different, right?

And it doesn't mean that they're being biased. It just means that they have a different basis of knowledge. And so, when people are making judgments about whether, and being reflective about it, and judging whether something is true or false, the fact that political partisans disagree doesn't mean that there's a bias. Just means that they have different basis of knowledges. And so, we don't know that that's motivated reasoning, per se, that's driving that.

**Andy Luttrell:**

Yeah, so I really like that, and I saw this when... Is it a preprint or is it out?

**Gordon Pennycook:**

I think the one that you're talking about is a preprint.

**Andy Luttrell:**

Got you. Yeah, and I forgot that you were involved in that, because I... The thing that always bothered me about... So, there's the classic study with how people evaluate scientific information



when it supports or opposes something they already think, and so one classic example is coffee drinkers. If you say, “There’s all this research now that shows that coffee’s actually bad for you.” People go, “Well, I don’t know how good that evidence is.” That looks like a bias, but again, like you’re saying, it’s like, “Well, if I’ve already come to this idea that it is safe, and I’ve reviewed the evidence, and I’ve really thought about it, and now you have one study to show me that it’s not, I have reason to doubt it, right?” Because it’s not that I’m being crazy and irrational. It’s just that, well, that doesn’t really make sense given what I already know about this. Which is basically what you’re arguing, right?

**Gordon Pennycook:**

Exactly. Yeah.

**Andy Luttrell:**

So, this suggests that motivated reasoning, as psychologists have been talking about it for a long time, may largely not be the thing that we’ve thought it was?

**Gordon Pennycook:**

Yeah.

**Andy Luttrell:**

Is that-

**Gordon Pennycook:**

Yeah.

**Andy Luttrell:**

Yeah. Okay.

**Gordon Pennycook:**

Exactly. And I think that when most people talk about motivated reasoning, they’re talking about a thing that either is extremely rare or just doesn’t really exist. Like where people are kind of self-deluding in a certain sort of way. It’s not to say that people are rational, of course, and that ideology doesn’t matter. It’s just that the locus that we were looking at in social psychology is wrong. Mostly what people cared about is the impact of identity, and the self, and all that kind of stuff. And I think what’s much more influential than that is the impact of prior experiences and prior knowledge, what you’ve been exposed to. The kind of things that people care about in the communication and journalism literatures. That, to me, seems like it’s much more impactful.

If you know that someone has been watching a lot of Fox News, that gives you more information than whether they identify as a Republican. Do you see what I mean?

**Andy Luttrell:**

Because that’s just the information that they have stored up.

**Gordon Pennycook:**

Yeah. And so, you have to use that from a straight information processing perspective to interpret what’s going on in the world, and identity tells you something about what kind of information is

stored in their brain, because of exposure and so on. But it doesn't tell you as much as knowing about that information itself, and what they've been watching, and all that kind of stuff.

**Andy Luttrell:**

So, it seems that... So, this sort of bias account of belief in false information doesn't seem to stack up well against sort of just a, "I'm not paying attention to the accuracy of it," right? Given that, is there any hope? What do we do to account for this belief in misinformation that seems rampant?

**Gordon Pennycook:**

To me, it's one of the most positive... I do so many studies that you just kind of feel sad at the end, you know? Like you do a study on prior exposure and like, "Oh, look. People do believe this stuff more if you repeat it, even if it's really unbelievable stuff." First, you're excited that you actually got the finding you thought you would get, and then you're kind of sad for the world or whatever.

In this case, there really is a positive thing. If it was the case that it was all partisanship, that people were just kind of protecting their partisan identities, and reasoning was being coopted to kind of reinforce that, then we would be pretty much hooped, right? We wouldn't be able... To solve the problem, we'd have to make people less partisan, and good luck with that, right? But what our research suggests is that people care about accuracy, but they just kind of don't remember or think about using it, or judging it, whether things or true, and so what that suggests is we can just get them to do that. You know, like remind them about accuracy. And that's what we're kind of pushing right now, is intervention. We've done a bunch of experiments where we, as an example, just at the start of the experiment, people in the control, they're just given a bunch of headlines and asked whether they would share them on social media or not. And in the treatment, the intervention is just at the start of the experiment, we give them a single headline, we say, "This is for a pretest. We want to know if you think it's accurate or not." Just to get them to think about accuracy. We don't really care what they respond or anything about it.

We just give them, it's just like a politically neutral sort of headline, just something, random headline, doesn't matter. We've done a bunch of different ones. It doesn't seem to matter what you ask about. And then after that, they do the same thing, where they asked about whether they would share these headlines. Some which are true and some which are false. They don't know that they are true or false ones, obviously. And the fact that when we ask them at the start to rate the accuracy of the single headline, that improves the extent to which they discern between true and false content when they're making judgments about sharing. Because basically in the control, there's nothing prompting them to think about accuracy, so they don't do a very good job at discerning between what's true and false. When we prompt them to think about accuracy by asking about that one headline, then they do a better job of it.

And we've done this a bunch of times. It works. Often what it does is it decreases the sharing of false stuff. We did one with COVID headlines more recently. It seemed to have more of a positive increase for true stuff. So, it kind of depends on which headlines you use, I think. But the important thing is that it's increasing the difference between what's true and false. Like you don't want people to just not share stuff online anymore, right? You want them to share true stuff, especially as it relates to a pandemic. You want them to actually be sharing that content.

And so, I guess the biggest criticism of all this would be like these are kind of like lab studies, where we're asking people what they would share, and we don't know if that's actually what would happen or whatever. And so, we did in a recent paper, and this is... I gotta give credit to Ziv Epstein and Mohsen Mosleh, who are also postdocs in Dave's lab. What we did is we did this large scale Twitter experiment, and so I'll just describe the experiment, because I think it's interesting.

So, basically, here's what we did. They created cooking bots. So, they were just like explicitly labeled as being bots, and so there's no deception there or anything. And then what the bots did is they followed a huge number of people who shared recently headlines from a low quality source. Well, actually two. We started off following people who shared content from Infowars, and also Breitbart, and then Infowars got booted off Twitter, and so it was just Breitbart after a while. So, Breitbart, I don't know if people know anything about Breitbart, but it's a rightwing, alt-right basically website that creates what would be called like hyperpartisan content. They don't make things up. It's not fake news. But it's pretty misleading and it's biased coverage. And if you ask fact checkers and professionals about it, it's not a very trusted source. So, it's like what we would describe as low-quality content on social media. So, we're getting now... We're building away from fake news and just talking about the quality, overall quality, of the content.

So, what we did is we created these bots. They followed a bunch of people who shared low-quality content, in this case from Breitbart, and then some people follow you back on Twitter, right? Even if it's Cooking Bot, some people, when you follow them, they follow you back. And so, we followed enough accounts that we got like 5,000-ish people to follow us back. And when someone follows you back, you can send them a direct message, and so we sent, from a cooking bot, this message saying, "Do you think this headline is accurate?" Right? Which is a super weird and dumb message to get from a cooking bot, and it doesn't matter, most people don't respond. Some people don't even open it. It's hard to tell whether people see it based on various things.

But that's our treatment. We just presented people, different subgroups of people on different days this message. And then, because Twitter has an open API, you can then look and see what people shared subsequently, okay? And so, what we did is we compared the quality of the news content that people shared if they got the message versus everybody else, who had not got the message yet. And what we show is that basically the quality of the content improved for 24 hours after we sent the message.

**Andy Luttrell:**

Regardless of whether they replied to it?

**Gordon Pennycook:**

Yeah. Oh yeah. We don't know. In fact, we don't even know... So, it's like in the paper, if you read it, we're underestimating the effect size, because we don't even know if they opened it. Right? And so, by the way, just one final piece that's important for this is the way that we quantify the quality of news is we just look at the sources. Because it's like we can't go through and fact check every single, thousands of links, obviously. And so, from a previous paper, we had fact checkers rate the quality of like 60 news sources. And so, we just take all the links that they have and then figure out what the average quality is, and then that's what's changing, is that basically what you see is they're sharing more stuff from New York Times, CNN, and less stuff from Breitbart and Daily Mail, or things of that nature. And so, that's our evidence that this thing, which is a pretty

straightforward intervention, just like think about accuracy a bit, even in a subtle way, and it actually improves things.

It's only surprising if you're cynical, but it does seem to have an impact, and so that, to me, seems important.

**Andy Luttrell:**

So, there are other efforts to mitigate sharing fake news, so I've noticed recently Facebook will... You'll see a friend or relative share something, and then a couple days later, there's like a giant image over it that says, "This is deemed invalid." Do those sorts of efforts, do those strike you as... Because in some ways, you say, "Well, those are prompting a concern for accuracy." I'm just curious to get your take on how are folks in industry trying to mitigate this, and in what ways are they doing it well, and in what case are they missing the mark?

**Gordon Pennycook:**

Yeah, so they do things like that. Also, Twitter had the thing where if you try to retweet the article, like opening it, then they'll prompt you to ask you...did you see that? Yeah, anyway, it's pretty interesting. So, mostly what Facebook in particular is doing, and what they have been focusing on ever since 2016 is fact checking. And some fact-checking-based intervention. I think that's good, and fact checking's great. We need to have something to point to when things are false, obviously. But it's not sufficient by itself, and there's a pretty simple reason for that, and it's like it's a lot easier to make shit up than it is to debunk it. And even if they were working at maximum efficiency, when something is put up and it's false, there's a few steps that... It has to be spread far enough that they get alerted that there's some falsehood. They have to then communicate with the fact-checking agencies that they're working with to do the fact-checking work. They have to do the fact-checking work. They have to communicate with Facebook about that fact-checking work. Facebook then has to implement the thing that you saw, where they kind of black out the thing. And then presumably do other things, like downrank it in the algorithm or whatever. But that all takes time.

By the time, probably a week or something, I don't know, certainly not within a day, and by the time that happens, most of the impressions have already been had for that content. And so, it's just not-

**Andy Luttrell:**

And as you've seen, it's just exposure, right? Even just exposure to that information is enough to seed that belief.

**Gordon Pennycook:**

Exactly, so you have to cut it off at the kind of source, and it's just not really doing that. And so, I think they're good to have, but that does not offer even close to a solution in any sort of way. I mean, a solution would be lots of things at once. And this is why we're kind of pushing for our idea, because if you just... I mean, really what needs to happen is people need to interact differently with the medium. They need to change the way they're interacting with social media, and maybe it is the case that when they see those sorts of warnings, they will think more about accuracy. But they could do so many different things to get people to do that.

We're trying to test simple ad-based messaging things, or simple prompts, interstitials they call them. You know, when something pops up and you ask them a question or whatever. There's lots of different things that you can think about, but so far, the companies haven't adopted anything. Maybe they're waiting for the papers to be published. I don't know.

**Andy Luttrell:**

I'm sure they're on the edge of their seats.

**Gordon Pennycook:**

Yeah. Exactly. Yeah, they're just... Yeah. Once the paper's published, then they're going to change the entire platform. That's what's gonna happen. So, yeah, they're doing some things and they're being more responsive than they have been in the past, that's for sure, but there's still so much more that can be done. And part of the biggest issue, I think, is that the extent to which they work, like actually work with academics is wanting. There's a longer backstory for that, but a lot more could be done for that kind of thing, to show that they're really taking it seriously and not just doing things that look like they make sense from a public opinion perspective.

**Andy Luttrell:**

The other thing I worry about those, too, is that the trust people have in those institutions, right? So, I mean I've seen examples of someone shares something and then Facebook says it's false, and then someone rolls their eyes and go, "Oh, well, I guess if Facebook says it's false, then it's not." And so, you go, "Well, there's another hill, too." It's another place where the bias could creep in, and so in some ways, I'm curious whether fact checking, from a persuasion angle, right? The source matters. So, if the fact checker is an independent third party, presumably, that will help a little bit. But if it's even a stereotype that this is an institution that has an agenda, then the fact checking may actually only increase willingness to share, to the extent that any of those identity markers are actually anything.

**Gordon Pennycook:**

Yeah. It's a big problem, for sure. I mean, all the fact checking is done by third party fact checkers on the Facebook thing. I don't think they communicate that very well to people. It's all, whatever, kind of hidden, and you have to click the buttons and go through the things. But that's the other issue, too, so not only do they have to kind of keep up with all the falsehood, but someone has to determine what's true or false, or false enough to be hidden, and people might not believe that. What we're trying to implement is something that is based on an individual's capacity to determine what's true and false themselves, which people can do more than we think.

I mean, of course there's going to be some people who, if you get them to think about accuracy, it's not gonna matter, because they don't know enough or whatever. But for most really false things, it's if you think about it, it's pretty obvious what's true and false. That it's kind of made up, or whatever. And it least it would improve things to some extent. It might stop them from sharing it, even if they're not sure if it's true or false, right? You don't have to actually get them to make sure they know it's not false to stop them from sharing it, obviously.

And so, it's like a person-centered approach, right? Like you're not being the arbiter of truth, which is what Facebook does not want to be, despite those fact-checking approaches. And you kind of put the power back into the individual, but you just kind of nudge them towards doing something

that they already, like I said, people care about accuracy, so it's not like you're even getting them to do something that they don't want to do. It's just something that the platform isn't built towards getting them to do readily, and so that part of it can be changed.

**Andy Luttrell:**

Yeah, the platforms reward other stuff.

**Gordon Pennycook:**

Yeah. So, you could even think about things that would reward accuracy. Of course, that's going to be difficult. Giving people points in a kind of Black Mirror sort of way. Because you still have to figure out what's true and false and all that kind of stuff, and so maybe... I always get emails about that kind of thing and I'm just like, "I don't know of a way to set that up that makes sense." But whatever.

**Andy Luttrell:**

So, just to wrap up, I'm curious. What does the future of this work look like? What are the still unanswered questions, problems yet to be solved? What's on the horizon?

**Gordon Pennycook:**

I always suck at this question, you know? I feel like I don't know what I'm gonna do until 9:00 that day or whatever, and then I end up doing it. There's no plan. I don't have a plan. What we're kind of working on now, I guess, is related to that kind of prompting people to think about accuracy as kind of optimizing it. Doing it in fair different ways. Ways that might be actually implemented by social media companies. We're working with Jigsaw, which is the R&D arm of Google, to do some work on that. We actually did some work with a group called Luminate, and they do.. .So, basically we have various groups that are looking at different ways to implement something to make people think more about accuracy on social media. But that's still more of the same, like it's not some fundamental shift, or just incremental sort of work.

And apart from that, just briefly note that I've gotten sort of consumed with this misinformation work partly for just kind of obvious reasons, because it seems important and all that kind of stuff. But I still have all these various projects that have nothing to do with it that I'm still interested in or focused on, but it's hard to... What they don't tell you about being a PI is how complicated it is to try to keep up with all the stuff, and I blame myself for it, because I don't say no to projects very well. And so, I've just got a bunch of random things.

**Andy Luttrell:**

And looking at your, I was looking at your CV, and there are a lot of areas charted out in the work that you've done, and there's a thread through them, but I have found myself in similar situations, where it's like, "Oh, but this other thing is so cool and it has nothing to do with anything. But why don't we do the study?"

**Gordon Pennycook:**

I know. In a certain sense I'm just a profoundly selfish research person, you know? Because I just... I mean, partly it's that I charted territory that I kind of go into any area, because I'm interested in basically just reasoning decision making, and so that's where I'll go for basically anything if you think about it, you know what I mean? So, I just follow my whims and mostly

research things, because I think it's entertaining and important. But that's what keeps you motivated and whatever, so it is what it is.

**Andy Luttrell:**

Well, I will cross my fingers that we solve the fake news problem right away, and all the institutions get on board with the right science.

**Gordon Pennycook:**

Although I should note that even if they were to listen to everything that we said, we would still slightly improve things. Which from my perspective of having never done anything to have any impact on anything at all, slight improvement, that's great. I'm on board with that.

**Andy Luttrell:**

Very good. Well, thanks for coming on and talking about the work that you do. We'll look forward to seeing what's next.

**Gordon Pennycook:**

My pleasure. Thanks.

**Andy Luttrell:**

That'll do it for this episode of Opinion Science. Thank you so much to Dr. Pennycook for sharing the work that he's been doing. Check out the show notes for a link to his website and more information on some of the research that he talked about. You'll also find a full transcript of this episode. For more about this show, head over to [OpinionSciencePodcast.com](http://OpinionSciencePodcast.com), and follow us on Facebook or Twitter @OpinionSciPod. And oh, is that the rate and review fairy? What's that? They should go to Apple Podcasts and leave a nice review of the show to help people find out about it? Thanks, rate and review fairy. All right, that's it for this week. Thanks for listening and come on back next week for more Opinion Science. Bye-bye.