



Opinion Science

Hosted by Andy Luttrell

Fighting Misinformation with Sander van der Linden

January 30th, 2023

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

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

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

Andy Luttrell:

In 1690 the first multi-page newspaper in the English colonies was printed in Boston. It had the alluring name: Publick Occurrences Both Forreign and Domestick. The paper was published by British immigrant Benjamin Harris. He had come to the colonies just a few years prior, which of course left him with a dilemma—how could he assure his new community to trust a foreigner to share the news of the day? Well, he would just say so. He'd make it the central mission of his new journalistic enterprise.

He wrote at the top of his first issue that he wouldn't publish anything other than “what we have reason to believe is true, repairing to the best fountains for our information.” Any mistakes, he said, would be corrected in the next issue. He even deliberately acknowledged the elephant in the room, writing: “there are many False Reports, maliciously made, and spread among us.” These Public Occurrences would be the capital-T truth.

So there in the first issue of his newspaper, the fine people of Boston would read about confrontations between Native Americans and local settlers, the colony's military activities against Canada, outbreaks of smallpox and “epidemical fevers,” a fire that destroyed a bunch of houses, a local widower's suicide. But he also reported a surprising and salacious fact about the king of France, King Louis the 14th. The king was in trouble with his son because he'd been sleeping with his son's wife.

Well, the local authorities weren't too keen on bad-mouthing the French king...wouldn't be great for British-French relations. Plus, Harris was just sorta printing a newspaper on his own, so they shut down the newspaper just four days after its first issue.

But here's the thing. The king wasn't sleeping with his daughter-in-law. For one, he didn't have a daughter-in-law at the time. His only daughter-in-law died months ago—something Harris would have surely known. Even still, his late daughter-in-law was like super devout, and she was bedridden for years before her death. There's also still no historical evidence that anything like that happened. Seems like Harris just made it up. And he would have had reason to. Harris was a committed Protestant, and he knew that Louis XIV was ramping up the persecution of Protestants to make the country Catholic. Back in London, Harris was facing a bunch of legal trouble for printing anti-Catholic pamphlets. It's why he fled to the colonies.

So it all seems pretty calculated. The king already had a reputation for being a dude with a libido, so sleeping with his son's wife wouldn't seem all that unbelievable. And Harris would clearly be happy if his Bostonian neighbors saw his religious enemy as a monster. Benjamin Harris, the apparent champion of truth and crusader against false reports prints politically motivated misinformation.

So yeah, that's how far back fake news goes in the United States...to the first issue of its first newspaper.

You're listening to Opinion Science, the show about our opinions, where they come from, and how they change. I'm Andy Luttrell. This week I'm excited to share my conversation with Sander van der Linden. He's a Professor of Social Psychology at the University of Cambridge. Among the things he studies is the psychology of misinformation. Why do we believe false things and what can we do about it? This has become a major area of research in psychology as social media has allowed falsehoods to spread faster than ever before. Sander has a book coming out soon called "FOOLPROOF: Why Misinformation Infects Our Brains, and How to Build Immunity." It comes out February 16th in the UK and March 21st in the U.S.

It may be worth tying some connections in the grand Opinion Science web. First, I'll note that Episode 13 of this show from July 2020 features Gordon Pennycook and his work on fake news and controlling misinformation. I would listen to both if I were you. And also, a key idea in Sander's approach to preventing misinformation is this classic idea from psychology called "inoculation theory." You'll hear plenty about it here, but savvy Opinion Science fans will also remember that Episode 49 from October 2021 features Josh Compton and his take on the many varieties of psychological inoculation. So put those next in your listening queue, but for now let's jump into my chat with Sander van der Linden and what we can do about our pressing infodemic.

Andy Luttrell:

Nice. I was thinking about it, though, like when you must have been writing this... I mean, there's a lot of very up-to-date stuff in here. When did you pitch the book? When did this whole thing actually start?

Sander van der Linden:

Years ago.

Andy Luttrell:

Okay.

Sander van der Linden:

Yeah. Yeah. It was years ago. And I had the idea for years, and I just couldn't get around to writing a book. I don't know how you feel about this as an academic, but you're just so used to writing papers that the idea of writing a popular book... I mean, I got approached by a number of fairly academic publishers on the topic all the time, and I was thinking, "Oh, I could do an edited book." But you know, is that really what I would want to do if I'm gonna write a book? I'm already appealing to other scientists with my articles. Is it really gonna make that much of a difference to

really write sort of an academic book on it? And then I thought, “Hm, maybe I should do a more popular book. I’ve always wanted to write a book even before I got into academia.”

And so, that kind of resurfaced for me, the idea of writing a book, and then I just couldn’t get around to actually doing it. So, I had it in my head for two, three years, and I just couldn’t get actually to sit and write it down. Also, because it’s such an uncertain endeavor in some ways. It’s sort of like, “Well, is this gonna turn into a feasible project? Is anyone gonna read it? Is this really turning into something?” And it takes a lot of your time. You really have to sit down and write every day for however long it takes. And so, I was putting it off, and then it really started off for me when I got an agent who helped me shape this up into a proposal and then actually turn it into a book.

Because yeah, I did do a lot of popular science writing. I had a blog for a while in Scientific American Mind, and other outlets, and I enjoyed that process, but it’s... You know, one of the things I’ve learned, it’s so different from writing a 1,000-word essay for the public versus a book, where you have to think in advance about all the stories, and all the chapters. So, to answer your question in a shorter way, I started writing this, it must have been 2017, 2018 that I started on the book, just right after the 2016 craziness, and then... Well, I guess unfortunate thing about writing a book on this topic is that every time something new happened, then I felt I just had to rewrite the whole thing, and so that at some point I had to draw a line under it and say, “Well, I’ve covered so much misinformation, I think that I’ve covered all the major examples that one can think of.” Wars, pandemic, elections, and so yeah, that’s kind of where I just said, “Well, I have to just finish it now.” Even though there’s always gonna be a new strain of misinformation coming up.

Andy Luttrell:

At the time you started, though, had you really been doing much of the work on controlling it? Like the inoculation stuff? It seems like by 2017, that hadn’t really been off the ground yet. Or maybe it had been. Was that the early days of doing that work?

Sander van der Linden:

Yeah. That was the early days of doing that. In fact, it was much earlier, and again, this is something that I started... It was an idea I had during my graduate degree, so a very long time ago, and I talked to some people in my department who were interested, but I just... You know, it was kind of a vague idea. It was more about... My research question at the time was how do you inoculate somebody who’s already been inoculated with misinformation? And at the time I sort of called that reverse inoculation, which I know you have Josh Compton on the show, right? With whom I collaborated very much later, turning out therapeutic inoculation, which I think was a much better term than my initial reverse inoculation.

But it just... Yeah, I just parked that idea and did something else for my PhD, and then once I finished the PhD, I came back to the idea of inoculation, and this was before the 2016 election. It took us a long time to design that study. So, initially it was in the context of climate change coming back to the Yale program where I was at the time, and it took us at least a year to design it, because we really wanted to do a two-phase study. We wanted to understand the common misperceptions that people had, or the common forms of misinformation about climate that were out there, and then how we could inoculate people against it. And it’s kind of funny, because at the time it wasn’t

so obvious that exposing people to misinformation would have this huge negative effect, because there wasn't a whole lot of research on it.

I mean, there was this sort of old school research on propaganda and stuff, but it wasn't so clear to us. Now, there's so much research on it, but at the time that was a big thing just to see what's the effect of exposing people to misinformation on climate change. And then, you know, if it has a negative effect, then could we inoculate people? It's kind of a multi-phased study. So, we did that whole thing, then we wrote it up, and as you know, I transitioned to postdoc somewhere else, and so I had new responsibilities, new papers. We submitted this to... Yeah, kind of a big journal. It took a long time. It was a Nature journal. But they came back basically asking to change the paper in such a fundamental way that we felt it wasn't really our paper anymore. And so, yeah, it was tricky for me. Took a long time for me to get that paper out.

And so, we spent more than a year and a half, almost two years in revision with this paper, trying to decide what we want to do with this work, and it was just a coincidence that when we finally published it the 2016 election was ongoing and it just got picked up, this idea of you can inoculate people against misinformation. And so, it became kind of the perfect storm of the whole problem of misinformation, and the term fake news was getting a lot of traction, and then we published this study about this idea of, "Well, you can inoculate people against it," using a modern example like climate change. And that's really when it got picked up and we decided to invest more in this line of research. Yeah.

Andy Luttrell:

What is it that that journal wanted you to... What was the paper they wanted you to have submitted?

Sander van der Linden:

Yeah. Well, it's interesting. This is kind of a debate in my area. I'm not sure how... Perhaps you're familiar with it. I'm not sure if you know Dan Kahan, who works on cultural cognition theory? And so, at the time, this was a very influential idea, and we were working on measuring and figuring out how people perceive consensus in science, on things like climate change and how you could inoculate people from attempts to politicize science using this information. But Dan's idea was really that people aren't so responsive to information. People are motivated, right?

Well, I think the idea of motivated cognition is not so controversial. I mean, everyone processes information that's congenial to their own political sort of beliefs. But his hypothesis was a bit more extreme in the sense that the people who are the most educated, who are the most literate and numerate, will polarize the most when exposed to factual information. And we're kind of debating back and forth on how reasonable that hypothesis is, but at the time I think it was a very influential idea, sort of the whole stuff when cultural cognition was fairly new, and I think all of the reviewers were convinced that that should be the conclusion of the paper, that people don't respond to evidence.

And the fact that we were able to inoculate people and that even republicans and conservatives in our sample responded positively to the inoculation about the scientific consensus, and they just said, "Well, in light of what we know about cultural cognition, this just wouldn't happen." And

so, we were like, “Isn’t that a motivated sort of process?” In a way, they were proving Dan right. But not in terms of our data. I mean, our data was just so clear that everyone was sort of updating their beliefs in the sample.

Now, I don’t think we changed the minds of people who were deeply skeptical of climate change. But I do think we were able to neutralize attempts to further sort of entrench them in these ideas using the example we had in our study. So, that was the main problem. They wanted to reframe the whole paper around cultural cognition. And for us, it was sort of like, “Well, it’s an interesting idea, but it’s not really core to the idea of inoculation or the purpose of this study.” And so, we struggled with being at a prominent journal and having to reframe your whole paper, so we decided not to and just go somewhere else.

Andy Luttrell:

To that paper, so I recently in my persuasion class for undergrads taught a day on misinformation, which was new for this semester, and I showed the results of that paper, which it’s not that often that a five or six condition study makes that much sense, right? That graph is beautiful because you go, “Condition A, B, C, D, and E all look exactly like what you’d expect,” right? And it just... It makes sense. Show the consensus information. People believe it more. Show the criticism, people believe it less. Show both, nothing happens. But inoculate at two different levels and you protect against misinformation. So, it’s just a really nice... It’s just so clean, the way that it looks to tell that story.

Sander van der Linden:

Yeah. I love that you say that because I was sitting in my... I was a postdoc at Princeton. I was sitting in my office by myself looking at the data and I was like, “This has never happened to me.” And I was like, “You know, I run experiments, and they fail all the time. It looks like a mess.” And it’s a hard time publishing stuff that’s messy, and so I was looking at this data and I was like, “This looks beautiful.” And I was like, “It’s all stacking up.” It wasn’t perfect, but mostly it’s a very sensible progression of how things should happen according to the conditions. And I was like, “Wow. This has really never happened to me. I think this is really cool.”

And then I thought, “Well, maybe it’s a fluke.” But you know, it was a large representative sample, and then we ran some other studies that John Cook, who was a colleague who ran the same study, more or less, with a slightly different treatment independently, and we were kind of back and forth, and he was finding the same results. And so, it gave me more confidence that this was actually what was happening. And I was very optimistic at the time that because it was such a sensible pattern that seemed very straightforward, that it was gonna be an easy time publishing this in Nature Climate Change it was at the time.

But yeah, it’s funny how you say that, because that’s what it was to me, and I was so happy with the result, and my co-authors were, as well. But we did spend almost a year and a half thinking about these conditions pretesting, making sure that it made sense, that we got it right. Which is unusual for a study, I would say. I mean, I’d love to take a year and a half every study that I run, but you know, most of the time it’s a few weeks, or sometimes even half a day, or sometimes a few months, but it took us a long time to really think this study out. And then I thought, “Well, it

does pay to really deeply think about experimental studies, have multiple people design it, critique it, and then in the end you do get something that's perhaps worth publishing.”

Yeah, so that was great, but we were fortunate because it seemed so useful to the debate that was going on at the time around misinformation. And so, luckily, I think it still... People still found it useful. I think especially in the more practical sphere where people were looking for solutions. And so, I'm glad that it still got some-

Andy Luttrell:

The delay might have made it even more impactful, right? Because it ended up making it released at a time when people cared about it. So, let's maybe take a half step backward and start with what misinformation is, so to orient people to what we're talking about. If you were to define misinformation and explain why this is something we ought to be worried about, what would be your answer to that?

Sander van der Linden:

Yeah. Well, the definition that I maintain is misinformation is just any information that's false or incorrect. I think in practice we often operationalize it as something that runs counter to a well-established scientific consensus. For example, in the domain of science, but it could also be something that's just proven to be incorrect. Now, then the tricky thing is it's become so politicized, and people questioning scientific consensus, and mainstream scientific explanations, and so it gets tricky, but I think the most straightforward definition is just something that's false.

And disinformation, I think, is misinformation coupled with some psychological intention to actually actively harm or deceive people, so it's really not just false, but somebody trying to dupe somebody else intentionally. And now propaganda is disinformation with a political agenda specifically. That's kind of how I think about it in trying to separate all of these terms out in my own head.

Obviously, the taxonomy is not going to be perfect, but that's kind of how I distinguish these three forms. And I think disinformation and propaganda are much more concerning than misinformation in a way because misinformation could just be a simple honest error or mistake. But I think when people use the term, they often imply or think about disinformation rather than misinformation, and so I think it's useful to maybe split these terms out. And the same with fake news. It's become such a political term that a lot of us try to avoid it, but yeah, I think when people say fake news they often mean disinformation rather than misinformation.

Andy Luttrell:

It's tricky because you're implying that you have the correct answers. When you talk about misinformation, you go, “I know the truth about everything, and so I can like...” It's this omniscience that we don't actually have on the ground, right? We're just sorting through information. But it seems that just what's concerning is that we know that at least some proportion of the information that we share and use to form our judgments is actually just demonstrably false, right? And so, we're concerned about how easily people can incorporate that false information into their worldview.

So, as an example, the climate change stuff, what was the misinformation that you were considering in that work?

Sander van der Linden:

Yeah. So, the misinformation we were considering in that work was a petition called the Global Warming Oregon Petition, and maybe it's useful here to sort of introduce our philosophy on this, and I think you're right when you said people assume that you know the truth or there's some ground truth, and I think it's so important in science to acknowledge uncertainties and that we don't know everything, but at this point in time, at least given the weight of the available evidence, we can say something is more likely to be correct than something else. And I think that's really what we mean here.

And for us, the way we've operationalized it is not to tell people what's true or false, because there's a lot of research in this domain that uses articles that are fact checked and just we kind of, in a way, we delegate responsibility to the fact checkers to determine what's true or not and we just say, "Well, okay. This has been fact checked as true. This is false. You have a correct or an incorrect belief."

The approach that we've taken is more we just want to help people calibrate their judgments as to how reliable a piece of information is. You know, assuming that we're not the ultimate arbiters of truth. We can say that when something is presented to people, it can be presented in more or less misleading ways, and it would be useful for people to be aware of the ways in which they can be misled so they can tune their own judgments in a way that's perhaps more consistent with what is accurate versus what is manipulated. And that's kind of the paradigm that we use.

So, the Global Warming Oregon Petition is a real petition, so it's a real website that gathers signatures claiming that there's no scientific consensus on global warming. And in fact, a lot of scientists have signed on saying that climate change isn't real and that sometimes it can even be a good thing. And so, inherently that's contradictory because you're first you're saying it's not real and then you're saying it can sometimes be a good thing, but beyond that, there are manipulation techniques that are present in that petition. So, we're not saying that this petition doesn't exist or that it's not real. It's real. It exists. But there are techniques that are used to try to dupe people.

So, what are some of the issues with this petition? One is that a lot of the signatories are bogus, so the signatories are unverified, and so these have now been removed, but you know, people who signed up were members of the Spice Girls, so it was Dr. Geri Halliwell, Star Wars, Charles Darwin, people were just... A lot of jokesters would just sign up. And then there's the issue that what they call scientists are basically anyone with an undergraduate degree. So, not people who have a PhD in a relevant area or publishing in a relevant area, but just really anyone with a science degree. And so, they claim that I think 35,000 scientists have signed on, but even if you took that number at face value, that's less than 0.1% of all science graduates in the U.S. in a given year, so it's really a tiny number.

But they use this framing strategy, though, because it sounds like a big number, 35,000, and so they use the absolute kind of fraction to persuade people. But I think their biggest issue really is what we call the fake expert technique. So, in the past they used a template from the National

Academy of Science to present this petition, and the National Academy of Science actually had to put out a press release in the '90s saying, "This is not us. This is some impersonator. Climate change is real. We're not endorsing this petition." And then there are a lot of tricks that they use like that that make it seem scientific.

So, the Global Warming Oregon Petition is published by I think the Oregon Institute of Science and Medicine or something like that. It's basically like a prepper farm somewhere in Oregon, so it's not like a real sort of science institution, so there's a lot of fakery going on to try to mislead people, and that's what we try to make clear so that people can then make up their own mind as to whether they want to trust it or not.

Andy Luttrell:

Yeah, so this is a great transition to these strategies for manipulation that you've all come upon. And I didn't quite draw the connection before, but you draw the connection in the book to Cialdini's 6 Principles of Influence, right? So, this is... Listeners of this show will probably be familiar with those. And so, I'm curious. You pivot from principles of influence to principles of manipulation, and what to you is the difference between those?

Sander van der Linden:

Yeah. It's so interesting because I talked to Bob about this and he told me this wonderful story about Bill McGuire, who if people don't know, was the originator of the inoculation analogy, who visited Cialdini once before he did this research, and Bob sort of asked, "Yeah, what do you think? How should people study social influence?" He said, "Well, if you really want to know about how people are influenced, you should go and ask the people who do it for a living." And so, that kind of motivated Bob to go undercover and sort of explore these six principles of influence.

And it kind of motivated me to think about what are the principles that are used to systematically deceive and manipulate people with misinformation. Can we uncover those and then inoculate people against them? So, the difference between the two, and I just kind of footnoted this in the book because I didn't want to make it a whole thing, but I think a lot of the principles that Bob highlights, like expert authority can influence people, or conformity, or something like that, is fine in a sense. There's something that I refer to as ethical persuasion, that if you try to persuade people, you make it clear that you're trying to persuade them, and you're being ethical about it in a sense that you're using real experts who are just stating what's within their expertise, yeah, you're trying to influence people, but I wouldn't consider it bad or manipulative influence.

So, the flip side of that would be using fake experts, so you use people that have fake credentials, or you pretend that you have more expertise than you really have, or you have somebody who has no expertise in climate science talk about climate change, but because they have a doctorate in some other discipline, it sounds scientific. Well, you know, Alex Jones had this guy on his TV show, or his radio show, pushing these supplementary vitamins all the time, and he would call him this doctor from MIT, but in reality this guy did like a summer program at MIT. He didn't study medicine or anything remotely related.

And so, you know, this is what I would call the fake expert technique, which is clearly manipulative and influences people in probably undesirable ways, whereas if you're clear about people's

credentials, that they're an expert, and you're being transparent, and you're trying to influence people, I think there's nothing wrong with that as a social psychologist. I have colleagues who still think there's something wrong with that, but you know, who are all about this sort of inform, not persuade, but personally I think persuading is fine as long as we're all in on it. I think the idea is that if you're trying to influence someone without their knowledge using manipulative tactics, I think that gets into the realm of the dark arts. That's why I think it's useful to inoculate people.

Andy Luttrell:

Yeah. The influence... All of those are sort of premised on if you read Cialdini's book, it's very much like, "Okay, but don't... You can't lie to people." Like, "Yes, people are more influenced by authority figures, by consensus information, but you can't just tell people oh yeah, everyone in your neighborhood has already donated," if they haven't, right? That's persuasive but it's manipulative when I'm intentionally deceiving you to achieve some end of my own. Is that the difference?

Sander van der Linden:

Exactly. I think that's the difference. Yeah. Yeah. You can make up the social norm to influence people, but that would clearly be deception or misinformation, in fact, if it's just not true. And I think that's where I would draw the line. I think that's also where Bob would draw the line in terms of these principles of influence.

And I think what he liked about the book was that to sort of give people a guide of how they can spot the reverse. Okay, here there's some principles of knowing when to resist influence when it's not ethical.

Andy Luttrell:

So, to generate this set of six strategies for manipulation, did you do what he did? I mean, did you talk to people who engage in this kind of work? Where did this list of six strategies come from?

Sander van der Linden:

Yeah. Yeah. It's not some magical list. It's not exhaustive. It's just... John Roozenbeek, who's a postdoc or research associate now in my lab, and he's been with us for actually quite a few years now, since the beginning of a lot of this research, we spend a lot of time trying to uncover the techniques that people use, both through literature, so we looked at report from NATO and other organizations who have done a lot of systematic reviews on information warfare and tactics that are being used in the field. We listened to... So, we both had a mutual colleague who set up these events, and he would invite people, including some of those people who worked as a troll in a Russian disinformation farm, essentially, and we would hear about what their strategies were, how they handled it, what was going on, and how people were duped.

So, we went around and talked to the people who do this for a living, as well as some systematic research in the scholarly literature, and then we sort of combined what we thought were some of the most prevalent techniques at the time, and those were the six degrees of manipulation. And in fact, a lot of people ask me about this. Aren't they kind of broad? And it's like yeah, they were intentionally broad because the more work we did later on, the more specific we got in terms of

the techniques, but we often find that they're just subcomponents of these sort of larger strategies that we initially identified.

One of the techniques, might be useful for people, is called impersonation. But that can happen in lots of different ways. Fake expert just being one example of impersonation. And so, that's kind of how we stumbled across these larger, overarching categories that are being used, and that's kind of where we started, yeah.

And then, you know, there are more that I discuss later on in the book, but they can all... most of them can be traced back to some of these larger, overarching strategies. And sure, people ask emotion, right? But it isn't just emotion. It's using emotions to manipulate people, things like outrage, or fearmongering, and then... Yeah. Some people are like, "Well, can't you use emotions to manipulate people for good causes like climate change or poverty?" And in a way, the answer is yes, but it would still be the same in the sense that you're using a technique to influence people that might be misleading and is not transparent, and so we should be careful about that.

Andy Luttrell:

Can you give a flavor of some of the other ones, as well?

Sander van der Linden:

Yeah, so the six are impersonating, that's one. The use of negative emotions to scare people or to cause outrage. Polarization, so polarizing groups in society. So, you know, what people often don't realize is that misinformation... There's this interesting side debate whether the real problem in society is polarization or misinformation, but I think often that the connection here is that actually that a lot of misinformation plays into polarization. So, a lot of misinformation is framed in a way to specifically tap into polarized debates that are happening in society, so polarization is one. Not just political, but just trying to tear divisions between groups in society.

And then we have trolling, so trolling is very common during elections and influence campaigns. Basically, setting up troll farms to sow divisions and use some of these other tactics like polarization. But it's also a technique in itself, so you try to manipulate public perception of a debate by creating fake accounts and falsely amplifying debates by creating bots. And then sophistication of the bot is kind of where trolling comes in, so you can have automated accounts that are very low quality, or you can have humans pretending to be somebody expressing an opinion online.

So, for example, a lot of Russian trolls will pretend to be Americans and post about things like provocative content about relations between African Americans and white people, and issues that are on the nexus between, like gun violence, and things like that, so these were real people pretending to be, you know, I'm just an American sitting at home and I have this idea about, "Hey, we should all not get vaccinated." Or then they would also say, "We should all get vaccinated," to try to stir up a debate. And so, that's really the essence of trolling and falsely amplifying societal divisions.

And then there's conspiracy theories, so those are really part and parcel of misinformation now. But the thing I think that was interesting for us is that conspiracy theories often leverage a grain

of truth. So, what we try to do in our interventions where we try to inoculate people is tell people if you're being too ridiculous with some outlandish conspiracy, nobody's gonna pay attention, right? It has to be grounded in something real in order for a conspiracy to take off, and that's why they're often so influential. So, conspiracies are obviously a big one.

And I think the last one is called discrediting, which is really about saying you're fake news and trying to discredit other people using a range of techniques. Discrediting the mainstream media. I think Trump has these famous tweets where he basically says, "NBC, CNN, New York Times, they're all fake news." And that's part of the discrediting of official organizations and mainstream media. And sometimes that just turns into whatever opinion you have that isn't aligned with mine, that's fake news. I'm gonna try to discredit you.

And so, those are some of the strategies that we inoculate people against.

Andy Luttrell:

And it seems like they're not mutually exclusive, like they can bundle up, right? So, trolling leverages negative emotion to get people to divide themselves politically, and so polarization, trolling, and emotion all sort of can get wrapped up in one another. The one that I was the most uncertain about was the conspiracy one. So, I wasn't quite sure how to take that, so is this a strategy for misinformation or a type of misinformation? How do we construe conspiracy as kind of on par with impersonation?

Sander van der Linden:

Yeah. That's a good question. And I think you're right, a lot of them can be combined. And so, in some of our interventions, especially a trolling batch, we kind of put it all together for people because trolling uses some of these strategies. I think conspiracy theories can both be a form of misinformation and also a technique to construct misinformation. So, when you're stepping into the shoes of somebody who's trying to produce misinformation, one of the things that you want to do to make it more effective could be to write it in such a way that it leverages the idea of a conspiracy. And so, you might want to say... You know, you could just have misinformation saying, "Oh, taking vitamin C will help reduce the effects of the coronavirus," which is false, right? And that doesn't use any particular technique. But you could also say, "Oh, big pharma is trying to hide that vitamin C is really curing the coronavirus because if people just take vitamin C then they'll lose all their profits from producing the vaccine," or something like that.

So, that uses conspiracy as a technique in constructing misinformation in the hopes of it going further and reaching more people than it would have otherwise because you didn't use a tactic. So, that's I think where you can use it as an active ingredient in creating misinformation, and that's kind of what we expose in our interventions, but certainly you can also just descriptively analyze misinformation and call some of it conspiracy theories as a type of misinformation. And so, I think you can do both.

Andy Luttrell:

Is it that we're just naturally inclined to like conspiracies? It wouldn't be a strategy if it didn't work, right? Is that sort of the implication, that people are just drawn to this form of information?

Sander van der Linden:

Yeah. I think similar to appealing to people's emotions or outrage, there's a lot of studies showing that there's entertainment value to conspiracy theories, and so people find them inherently interesting, and I think because they lure you in with this simplicity that is some complex event, but you have a simple causal explanation, it's really appealing to people. Like, "Oh, you know, you have this explanation for why the things are happening, and it's sort of all connected, and it also seems well researched, and there's a lot to it." So, it's often dressed up as something intriguing for people and I think the other factor is that it plays into a lot of fears, and doubts, and uncertainties that people have about the world, so it gives them a sense of agency and control over what's happening, and we can pinpoint to the victim and the evildoers, and it sort of works it all out.

I think that's why it's so intriguing for people that there's something in there for most people depending on what your psychological needs are. And we often do this. When you fictionally sort of create a conspiracy from scratch, and the sort of birds aren't real one, I love that one because it's so good at explaining what's happening. If people haven't heard of this, there's this satire conspiracy called birds aren't real, and the idea here, and this is why I think conspiracies work, because it sounds ludicrous. The idea that some birds are really government drones here to spy on us. Or in fact, they're just all drones in this conspiracy. The birds aren't real because they're all drones. But the way they explain the questions that people might have and how they explain that away with very simple causal explanations is what I think makes it intriguing.

So, for example, why do birds sit on power lines? Because they need to recharge. Because they're drones. It all just makes a lot of sense now. And so, even if this is satire, it's just so gratifying I think for people to think about, "Okay, what else could be part of..." And so, okay, I'm already suspicious of the government, and privacy, and people spying on me, so this totally fits into my worldview. Even though it's a joke, I think it was constructed in a way that really illustrates that the person who came up with this understands the appeal of conspiracy theories.

I think, my personal two cents, the way I would contribute to this conspiracy is that you shouldn't say that all birds are drones, because that can easily be falsified. To make the conspiracy more intriguing, it should be that only some birds are drones, and then you can have a whole thing about what those birds look like, and the different spots that they have, and the indicators of when it's a drone. I think that would make it even more realistic for people.

Andy Luttrell:

Okay. Well, hopefully that's the one that takes off. So, if we have these tools of manipulation that folks use to spread this kind of information, the next step is to try to counter that, right? So, I want to talk about the work that you've done in that domain, and in particular the bad news game. And so, just to give us a starting point, can you just sort of describe the basic idea of this game and why it should have anything to do with stopping misinformation in its tracks?

Sander van der Linden:

Yeah. Yeah. So, this game we created Bad News, which was sort of a pun, really originated from the idea that or out of the idea that we could inoculate people. But the challenge was with the climate stuff, people come into the lab or online, and they read a 600-word essay preemptively before they see the misinformation, and I had a student at the time, so Jon Roozenbeek again, who

said, “It’s a really cool idea but how do you scale this in the real world? How can we actually... People aren’t gonna come into the lab, read a 600-word essay, be inoculated. How could we really do this in a way that’s also fun and entertaining for people?”

And so, he came up with this idea that maybe... So, he and a friend of his were thinking about this idea of a disinformation simulator, and the more we were talking about this, the more we realized we could combine the two ideas. So, you could have inoculation within the idea of a disinformation simulator, and so you could combine this. I thought this idea of simulating what’s fake news or what misinformation strategies are often being used could be really powerful when you take the inoculation idea, and you create a weakened dose out of that. So, the simulation would be the weakened dose of the kind of tactics that are used against you, and then we needed to have an environment in which to make that work.

And so, obviously there were a lot of parties involved. Game designers, software engineers, media organization we worked with, but ultimately we came up with Bad News. Bad News simulates a social media feed. We started out with Twitter, so it simulates Twitter. You have a follower meter and a credibility meter, and the idea is to gain as many followers as possible while not losing your credibility by making use of weakened doses, controlled doses of some of these techniques of the six degrees of manipulation in scenarios that are seemingly real on social media.

And so, that’s really what the game is about, and it tries to inoculate you against each of these six strategies by exposing you to weakened doses and showing you how to spot them through this principle that we kind of termed active inoculation. So, active inoculation is really premised on the idea that people generate their own antibodies, so you step into the shoes of somebody who’s trying to deceive you and you’re actively generating your own content. And we thought that was so much more important than passive inoculation, which is kind of like coming into the lab, somebody gives you arguments that you can refute when you later come across a falsehood, but you know, you have to rely on somebody else’s arguments, you have to remember them, you have to be motivated. And so, we thought, “Well, what if we let people do it on their own?”

It’s a bit more chaotic. Bit more uncontrolled. But what if we let people just do it on their own and create their own weakened dose? Give them a few scenarios they can choose from. Whatever people like, using a heavy dose of sarcasm and humor and see how we get on. And that was really the original idea, and then we got a prototype, and then we started testing, and then it kind of went from there.

Andy Luttrell:

I’m curious about the actual building of the game. So, I shared this with my students, I mentioned, and I was like, “This is the rare game built by academics that’s actually worth playing.” It actually looks like a game. It actually functions like a game. It’s not just like a glorified slide show.

Sander van der Linden:

Right, right.

Andy Luttrell:

Which is a lot of how academic games turn out. So, what was the process of actually taking the seed of a clear, science-based idea, but making it a game that the world would actually be interested in playing?

Sander van der Linden:

Yeah. Yeah. It was such a fascinating idea because at the time, I had just moved to Cambridge. I was a new assistant professor. It's almost seven, eight years ago now. It was quite a while, actually. And Jon and I were talking, and he had this idea of the disinformation simulator, but it was more... It was not the weakened dose version. So, I was slightly concerned about the initial idea that he and his friend from college had. But the more we started talking about it, the more it turned into something that we could do.

So, I decided to use some of my startup funding that I had to start your own lab to throw into this idea. And I thought, "Either this is all crazy and it is gonna turn out into something terrible or this could be really interesting." And so, it was sort of a high-risk project, but initially we started off with a bare bones version of the game, and we have these techniques, so the first person was researching these techniques, right? And so, we were kind of doing things in parallel, so we had a software graphics design company mockup some ideas, and some scenarios, and what it would look like, and then we were doing this research on the side, and then at the end when it sort of came together, Jon and I started writing the scenarios, and inputting things, and seeing how it worked out, and doing some testing.

And Jon's colleague started his own sort of media company at the same time, and they were interested in producing innovative solutions, real, practical solutions to help people, and so that's really where we felt the pressure to make this into something that's going to be fun for people. Because the way I had originally maybe envisioned it was something like a boring slideshow, something like an academic would come up, right? As you were saying.

And I think Jon was quite instrumental having all these sort of slightly controversial ideas in making sure that it didn't turn into sort of a boring science project. So, I had a risk management strategy, which primarily was to try to reign in some of his Dutch humor, which wasn't necessarily suited for a large international audience. Quite direct sort of humor. So, I was a content moderator, but also things... Making sure that we stick to the science while also dealing with the creative company and their ideas. And they were really good at just putting out an intervention that seemed fun and interesting to play. But it also made us a bit nervous, you know?

It was interesting, because when we had a working prototype, I loved it. The idea that we were able to simulate a social media feed I think was really powerful, that we could show these weakened examples. I mean, this was kind of the first interactive fake news game at the time, so I think that was really novel in the way that they programmed it and were able to come up with things. But then there was the challenge of embedding the science and making sure that it follows these strategies, and what's the right sort of way of doing it, and it's different from a standard lab experiment, right? Because we made it a choose your own adventure game, which means that we can't control the exact path for every individual user. So, if you're really stuck in sort of a causal inference paradigm, this is really challenging because even though conditions control treatment

might be randomized, we don't have full control over what every player sees all the time in the game.

And so, it's a 15 to 20-minute deep dive, and so it's a bit more messy, but we really wanted it to be a field sort of experiment, and a real thing that people would enjoy. Because Jon sort of said if it's not a choose your own adventure game, it's not cool, man. It's gonna be boring. If there's only one thing you can pick, and I think he really had a good sort of pulse on the gaming sort of industry and what people were doing. So, yeah, I think all of that was really instrumental just to make sure that this is not gonna turn into the yawn factor for people, so that it's a real game, it's fun, but still has the science to go with it.

And you have to be lucky, because the creative partners that we were working with, they were also interested in actually evaluating this and making sure it has impact, and they were also very interested in the science part and adjusting things based on the science rather than vice versa. I mean, sometimes we had to compromise, and they would come up with something, and I think there were... Some of my students came up with memes that weren't endorsed by the creative team, and so they felt really strongly about some cat meme or something in the game and they were like, "There are no cats in the game. This makes no sense from a creative point of view why the last screenshot of the game should be some kind of funny cat meme or something." And they were like, "This makes no sense." And so, we had to do things, but also embedding... So, the whole thing about we had the prototype, we all agreed it was working, it took a long time to build this, but then we had it working, and then we said, "Okay. Well, now we have to evaluate it, so we need a pre-post evaluation at the least."

And they were sort of like, "Whoa. Okay. That's a whole different ballgame." Because they're creative people. They don't really do that. And so, yeah, that required a lot of back and forth, and what's possible, and yeah, some of the fair critiques people have about some of our initial studies were just the result of... You know, the programmers, we just had to work with what they could deliver in terms of the engine. And so, things would slow down the engine a lot if we had too many items, so you know, you can't have your 30 psychometric sort of instrument built into the game. We could ask maybe five, six questions at the beginning, and the anchors on the Likert scales would sometimes float around a bit.

And you know, we had to get that fixed and make sure that that was all working, and so it was an interesting experience also for them. And so, once we had a reasonable pre-post test, that's when we had it set up. But then, you know, they said, "Look, we can really build this out into a scientific evaluation component, but it's gonna take a lot of funding." And so, we had to find funding elsewhere to try to build it out. But now 10 years later, or eight years later, our games now have a kind of built-in component where we have a pre and a post test, and we actually have a dashboard so that the data comes in continuously, it does your ANOVA t-test automatically. It just gives us real-time performance data, so it's really... You know, we were able to make this highly advanced. Still with the same team, so we've learned together over the last few years in how to make this both fun and scientific.

I don't know. It's a long story, but that's kind of part of it. Yeah.

Andy Luttrell:

So, you're saying you were paying this graphic design team out of your startup fund?

Sander van der Linden:

Yeah.

Andy Luttrell:

You were having to request administrative support for paying these game designers. And I have to imagine they're like, "What is this guy doing? Why did we hire this guy? I thought he was gonna do science."

Sander van der Linden:

Yeah. Exactly. And it wasn't just a little bit. It was most of my startup funding. Because you know, creating a game like that was expensive. And so, it took a lot of... I mean, to be fair to the game developers, they really did this on the dime. It was exploratory. They gave us a huge discount for what... You know, if you go to Silicon Valley and ask them to produce this thing it would be millions. And so, we really did it... They gave us the cheapest version possible, like a prototype to work with, but it still required most of my startup funding. So, I didn't have admin support for a while myself, and other things, but Jon and I were confident.

And I don't know, when you meet a PhD student and they're so excited, and they have all these ideas, and I was like, "It's great." And we were just spending nights and days in the pub here and drawing things out, and it was just so much fun that I thought, "Let's just do it." I mean, if it doesn't pan out, what else is startup funding for, right? You gotta take some risks. Yeah. We were worried that it wasn't gonna pan out or that it wasn't working at all, and so we were relieved when initially it seemed like we were on the right track.

But I should tell you that before we did the game, we created a school version, like a board game. That's where we initially started. So, the story isn't very linear, but it wasn't like we all of a sudden decided, "Oh, we were gonna do this social media game." We were much more basic. So, the first idea we had was to create a board game, and we went into high schools with that board game and did some initial testing, like a pilot test, and that gave us some confidence that this idea could work, and that's kind of when we went online and decided that there was a major component missing to our whole idea here.

Because it always seems when people say on their stories like, "Oh, they did this thing and then they were successful," but not really. We created lots of different versions. We were fooling around with these cards, and giving people cards of the techniques, and then they have to create their own headline with different sort of... Okay, if you're the conspiracy theorist, what image would you use? The sort of burning car and chaos in society or this factual headline? And so, we saw students were kind of enjoying it, and kind of learning about the tactics, but then we thought, "You know, we just spent a year going into schools and all this fieldwork and we got 100 students in our sample, and this is all very difficult and expensive," and that's why we said, "Well, let's go online." That was kind of the backstory to that.

Andy Luttrell:

Okay. Good. Because I was wondering if you had any indication that you were onto something. Because to get to the psychology of it, it's sort of a pivot from classic inoculation theory, right? So, classically inoculation theory, as you've been talking about, is you believe something in particular, and then you're sort of coached to question when people try to challenge that belief of yours. And so, it's all very specific. It's like training you to counter argue particular kinds of attacks on a belief that you hold. Whereas this is quite a different idea, right? That it's more that there are techniques that people will use for lots of different forms of persuasion that you want to be attentive to, right?

And so, could you walk through just to flesh that out a little bit, like what is the premise psychologically of why a game like this would actually be helpful?

Sander van der Linden:

Yeah, so there's I think two departures from inoculation, the original inoculation theory that are worth noting. One that we already kind of touched on with the climate study was that the original idea was that you would have a belief, kind of a "healthy" belief, like getting vaccinated helps prevent against disease, and then you would be attacked on this idea, and the whole premise of idea is that you could bolster people's attitudes preventatively by telling them about these upcoming attacks and inoculating them so that they maintain their positive attitude about vaccinations. But my interest already in graduate school was about this idea of, "Well, it's an interesting idea, but I think the problem with McGuire's theory, he kind of wrote in the 60s. He said look, I'm quite confident that this could work when people have the right attitude. But when they don't, it gets a bit more tricky, and I'd love for future research to explore how this could work out." And then he never did anything on it, which is kind of surprising.

So, I interviewed John Jost for the book, who was McGuire's last PhD student, to get some sense of what happened there, and he kind of said like, "Oh, he loved moving on to new research projects and he kind of just saw that this was a thing that he came up with and it was for other people to continue." But then we saw with climate that it didn't really matter so much what your prior attitudes were. You could still be inoculated even if it... But then what does it mean theoretically for somebody who already has the disease to then be inoculated against the virus?

I think the way that we think about it is that in attitude terms, you could think about it as that you both change the original attitude people have, and then protect it from further attacks, which is kind of called therapeutic inoculation. I initially called it reverse inoculation and then I think we all agree Josh had a better term. Therapeutic inoculation, which is that you... Just as with some viruses, when people are already infected, cancer, HPV, you can still boost immune response by giving people the vaccine, and so it could still be useful even when people already have been exposed to a falsehood.

Clearly, when people have already been fully radicalized, it's gonna lose some of its or most of its value. And so, here's where I started theorizing about this, the inoculation metaphor, that I actually think it's better if we think of it as the where people fall in terms of their infection status. So, for some people it will be fully prophylactic, they've never been exposed, and they have the right attitude. Some people have already been exposed, maybe they haven't fully changed their attitudes

yet, but they've been exposed and are playing with the misinformation in their heads. Some people have been fully duped by the misinformation.

And so, I think therapeutic inoculation can have a really positive role in sort of the middle of the spectrum, whereas sometimes I think you just end up deradicalizing or debunking if you're dealing with people who are already fully entrenched in the particular belief system, right? But inoculation doesn't purely have to be prophylactic, and I think that was the big insight of that study and also later studies that kind of showed that you don't need to have a particular attitude.

And the reason why I think this whole infection status spectrum thing is important to answer your question is that when you go into the real world, you don't know what people have been exposed to, how often, where. Even in some of our work with social media companies, we asked Facebook. We're doing some pre-bunking on Facebook and Facebook can't tell us exactly what... I mean, can't or won't. I don't know. But you know, I think they probably don't know what every user has been exposed to, how much misinformation on a topic they've seen. That's difficult. That takes huge amounts of data and computation. They may not be doing it because that's not something they're currently... I think theoretically maybe they could do it, but it's not something, for example, they might be actively doing.

And so, that takes a lot of time, so we figured with these games we don't really know who's... I mean, we know in terms of we can measure it and stuff when people come into the games, but what would we ask people? The game touches on all of these issues and techniques. It's not like a one thing. We can't ask people about their attitudes on GMOs, and climate, or these techniques. It wasn't really clear what single thing we would ask people about that would count as their sort of prior exposure variable, so we kind of just assumed that people come into this game having been exposed to lots of different kinds of information. But the most practical thing is that we sort of want to pre-bunk these techniques.

We found that it's effective regardless of people's starting point, so to speak, and Josh Compton refers to this idea that therapeutic inoculation tends to be effective also in other sorts of experiments, and so for the practitioner it doesn't really matter if the inoculation is prophylactic or therapeutic. It helps people. That's the bottom line. I think for the theory, for us social psychologists, it's interesting to tease out where people are on the infection sort of spectrum. Is it fully prophylactic, is it therapeutic? Because there might be slightly different ways in which it works in terms of the mechanisms.

And so, that's part of it. And then the other part of your question is about the technique level, so inoculation has traditionally been on the issue level, so that it's clear that you have an issue and people have attitudes towards this issue, and then you inoculate them. But we wanted to scale this up and that kind of corresponds with the idea of what McGuire called refutational different and same. So, refutational different would be an attack that doesn't reference the specific content that you've been inoculated against. But traditionally, that's been a fairly narrow spectrum, so I inoculate you against one misleading argument, and I don't show you a related argument that wasn't raised in the inoculation itself. With this intervention, it was pretty broad, right? We inoculate people against a core concept of a conspiracy theory, but then we test them on a whole range of specific conspiracy theories that make use of the technique and then see if they've been

inoculated. And that's why we expected that maybe the effect sizes would be a bit lower, because it's not one to one in terms of what people are being inoculated against, but that hopefully it would still work.

And so, that's really the idea behind the technique level inoculation, is a broad spectrum vaccine, so you expose people to weakened doses of the technique, and then you test them on a whole range of specific variants of that technique and see if they still show resistance. And that's kind of the innovation there in terms of technique level inoculation that could be prophylactic for some people, might be more therapeutic for other people. And I agree from a theory point of view, that's interesting to sort out, but for the people on the ground it often doesn't really matter that much.

But the interesting point maybe for academics is that when I initially presented this, there was a lot of resistance from people who only read McGuire's original paradigm, and so they kind of said, "That's not the original McGuire experiment." And I sort of said, "Well, if we only did the original McGuire experiment we would never get anywhere in terms of new theories and new ideas and extending it." And it's funny that Josh Compton wrote this article, how he talked about the analogy being instructive, not prescriptive, because there's so much more about this viral analogy that's left unexplored.

I think John Jost told me the same, that McGuire was actually quite flexible in the way that he thought about it. I mean, he had ideas, specific ideas about his own paradigms, but he was flexible in terms of how this metaphor could actually work. Because to me, he never... As a social psychologist he never explored what for me is the true question of this line of work, is herd immunity. The social nature of inoculation. And yeah, he never really touched on that. And so, this is why I think it's so interesting.

Now we're doing kind of computational models of what happens if we inoculate X% of a social network and we achieve herd immunity against misinfo, and sort of taking us to the next level. That's all of these exciting open questions. Yeah, so that's kind of the idea. But I think it's the same principle. You just do it on a technique level so that it offers broader resistance, and we have to be a bit more flexible about the assumptions of people's prior exposure levels coming into the experiment.

So, we've done more work trying to sort out, for example, people who perform, who seem more susceptible on the pretest actually benefit more, so that does seem like it benefits the groups that are most in need of the vaccination. But there are other things, like if somebody really scores high on an index of polarization, they might not be inoculated as much on the inoculation technique, so they might have some prior beliefs that are interfering with the inoculation. We've done things with social cues, so if we put lots of social cues on the post that we show people at the end, so even though they've been inoculated against the post let's say that shows a conspiracy, if we say that a million people have liked it, that kind of interferes with the efficacy of the inoculation a little bit. So, it doesn't eliminate it, but it reduces it a bit.

I think there's lots of ways in which the inoculation, where context is relevant for the efficacy of the inoculation, including your prior belief status, source, norm cues, and all these interesting things that we need to explore further.

Andy Luttrell:

Well, that's great. I'll keep an eye out for all that stuff that's coming down the pike and this has been great, so thanks for taking the time to talk about this work.

Sander van der Linden:

Yeah. My pleasure. Thanks so much for having me.

Andy Luttrell:

Alrighty, that'll do it for another episode of Opinion Science. Thanks so much to Sander van der Linden for taking the time to talk about his work. Check out the episode webpage for a link to his lab's website where you can find more about his research. His new book is FOOLPROOF: Why Misinformation Infects Our Brains, and How to Build Immunity. If you're listening to this today...I mean, I guess you'll always be listening to this in your today, but if you're listening in my today, you can preorder the book online. But if your today is after February 16th in the UK or March 21st in the U.S., then you can just buy it right now. Like, you could go to the store and they might have it just like sitting there.

Also, at the beginning of this episode, you heard the story of Benjamin Harris' inaugural issue of America's first newspaper. In writing that, I owe a huge debt to a recent book by Andie Tucher, a professor at the Columbia Journalism School. Her book is called: "Not Exactly Lying: Fake News and Fake Journalism in American History." You'll find a link to the book in the show notes plus a few other sources I turned to. One thing about that, by the way, is that I guess the whole idea that journalists shouldn't like make stuff up was not always a prevailing mentality. It wasn't until the late 19th century or so that some newspapers emerged with a real commitment to objective, sound reporting. Before then, reading the newspaper was just like a breezy pastime, and part of the fun was picking the papers that reported the news from the perspective you liked. Embellishing and fabricating were just par for the course.

But for all the very real and trustworthy news about why and how we hold opinions, well gosh, this podcast has what you're looking for. Opinion Science is the name. Listen to it wherever you like, just subscribe to it, okay, so you don't miss any. Go to OpinionSciencePodcast.com for transcripts, links to the cool stuff we talk about, and ways to support the show. Have fun with that, and I'll see you back here in a couple weeks for more Opinion Science. Buh bye...