

## Phantom Patterns and Online Misinformation with Megan Fritts

---

**[music: Blue Dot Sessions, Golden Grass]**

**Christiane Wisehart, host and producer:** I'm Christiane Wisehart. And this is Examining Ethics, brought to you by The Janet Prindle Institute for Ethics at DePauw University.

We take in massive amounts of information on a daily basis. Our brains use something called pattern recognition to try and sort through and make sense of this information. My guest today, the philosopher Megan Fritts, argues that in many cases, the stories we tell ourselves about the patterns we see aren't actually all that meaningful. And worse, these so-called phantom patterns can amplify the problem of misinformation.

**Megan Fritts:** The occurrence of the phantom patterns not only makes the problem of misinformation worse, but if used intentionally, as we think they are by certain algorithms used by social media companies or sites like YouTube, it can actually be a case of what we call epistemic exploitation. Some epistemic good of ours, by which I mean some part of our good life related to things we believe or things we know are taken advantage of.

**Christiane:** Stay tuned for our discussion on today's episode of Examining Ethics.

**[music fades out]**

**Christiane:** In her book, *The Right to Know*, the philosopher Lani Watson writes, "Knowing matters. What we know, or what we think we know, and what we don't know determines much of what we do. We decide what to buy, what to risk, who to trust, and who to vote for on the basis of what we know, or think we know, about the world around us."

What we know is being constantly challenged in the era of "Big Data" and online misinformation. Each one of us takes in an incredible amount of information, and our brains tend to try and make sense of that information by recognizing patterns. Some of these patterns end up being false, or as my guest Megan Fritts puts it, they are "phantom patterns." She argues that our ability to know information and recognize patterns gets exploited when online media outlets weaponize that instinct. She uses the term "epistemic exploitation" to describe this phenomenon. (And in case you've missed our other episodes on epistemology, "epistemic" means relating to knowledge.) I'll let Megan explain further.

**[interview begins]**

**Megan Fritts:** So in this paper, my co-author and I argue that one reason that misinformation, especially misinformation that's accessed digitally over social media or other websites is such a gripping problem right now. One reason that it seems like such a prevalent problem and it's hard

Examining Ethics is hosted by The Janet Prindle Institute for Ethics at DePauw University, and is produced by Christiane Wisehart. © 2022

to fight, is that we're exposed to vastly more information than we ever have been really at any point in history. All of this data that we're exposed to gives rise to what we call... Well, a term that we took from two data scientists, what they call phantom patterns.

A phantom pattern is something that looks like a meaningful pattern that requires an explanation, but is actually not meaningful. It's a phantom. And these phantom patterns can arise either randomly just as a result of us taking in so much information all the time that spurious correlations start to look meaningful, or they can actually be intentional phantom patterns kind of fed to us as a way of getting things to look meaningful or important, so people are more inclined to click on articles, people are more inclined to read them and feel more anxiety about the topic.

So what we argue is that the use of phantom patterns or the occurrence of the phantom patterns not only makes the problem of misinformation worse, but if used intentionally, as we think they are by certain algorithms used by social media companies or sites like YouTube, it can actually be a case of what we call epistemic exploitation. This is another term we borrowed from philosopher Lani Watson. In cases of epistemic exploitation what happens is some epistemic are good of ours, by which I mean some part of our good life related to things we believe or things we know, are taken advantage of. They're exploited for the benefit of someone else. So what we argue is that in some cases, phantom patterns are used intentionally to exploit people into being misled by misinformation, and that this constitutes a violation of a right.

**Christiane:** And you write that what's being taken advantage of here is humans' ability to recognize patterns and it's called the pattern recognition instinct. So could you flesh that out a little more? What is that and why is it so important to humans?

**Megan Fritts:** So we use patterns all the time, whether or not we recognize what we're doing. We use patterns when we make art, when we write poetry or novels, stories. There's certainly a pattern to narratives that people find meaningful. There's universal recognition of particular human narratives. Just our ability to recognize faces is an aspect of the pattern recognition instinct. If you've ever recognized a face in something that's not actually a face at all, if it's the front of a semi truck or something, they always look like faces to me. That's the pattern recognition instinct at work. This is a really good instinct that we have. In the early days in humanity, it was a crucial survival skill. It's maybe not so much a survival skill anymore, or a survival trait any longer, but rather it's what we call a humane instinct. It allows us to communicate with, and through the use of patterns in meaningful ways to other humans. So when this capacity starts to be used against us, when we are thrown into an environment where there are all these phantom patterns triggering our recognition of meaningfulness that actually aren't meaningful at all, the result is that this capacity becomes useless or even becomes a detriment to us.

**Christiane:** When you say that the pattern recognition instinct is an epistemic good, you mean that it's a way that we can collect information and make sense of information or collect knowledge and make sense of knowledge, right?

**Megan Fritts:** Yeah. Absolutely. Right. By epistemic good I mean it's some good capacity that we have that relates to our ability to know things or to come to have true beliefs or good beliefs.

**Christiane:** we're now in the era of what you and your co-author call "Big Data," where we humans are interacting with massive amounts of data, massive amounts of information. Wouldn't it be a good thing for our pattern recognition instinct to take over in this era of information overload?

**Megan Fritts:** So, that's certainly how it seems. We tend to operate under the assumption that the more information we have, the better off we are. Anyone who has spells of anxiety like myself, knows that when you're in the middle of an anxious episode, often the instinct is to go straight to information gathering, to get as much information as you possibly can on a topic. It feels like getting more information means becoming more in control, having better access to the truth.

This is actually not the case though. So when you get enough information, when you get enough data, as the amount increases, what also increases alongside it is the risk or the chance of getting spurious correlations. One funny example from the early days of the era of Big Data was in, I believe it was 2004 to 2006. There was a spurious correlation in the data that people didn't know was spurious at the time. We were gathering data on a whole bunch of things and noticed a correlation between the increase in murder rates and the increase in the sales of iPods.

So we measured this over the course of a couple years, and the rates of increase stayed very similar, nearly identical between the two. A few very well-respected institutes, The Urban Institute in Washington, hypothesized that these two rates were connected. They called it an iCrime wave and hypothesized there must be a relationship between these. Maybe people are being murdered for their iPods, maybe listening to your iPod too much made you homicidal, something like that. And as it happened, these two numbers weren't related at all, we had so much data that we were able to see things that looked like they must be connected, but it just turns out that they weren't at all. So once you get too much data, it becomes almost a statistical certainty that there will be spurious correlations that look meaningful. And the difficult part is figuring out which ones are meaningful and which ones aren't without letting that instinct take over and make that decision for you.

**Christiane:** Yeah. So that's a great example of maybe a phantom pattern, as you mentioned, that people are just crowdsourced more or less or came up with on their own, but you're also writing about companies or entities that are maybe exploiting this pattern recognition instinct, and exploiting the fact that we can find phantom patterns. And so I was wondering, how is it possible for... I mean, especially for artificial intelligence to do anything like exploit when it's just a machine or it's just artificial intelligence?

**Megan Fritts:** The answer to that comes down to how do these really complex machine-learning algorithms work? And I'm certainly not an expert on this. So I'm relying on the

wisdom of people I have read talking at a level that I can understand about how these algorithms really do function. But let's take a site like YouTube. YouTube has changed quite a lot since its very early days, where it was just sort of a fun, innocent place where people uploaded their homemade music videos and everything seemed really quite nice. Now, maybe some of the listeners are aware of this, but people have done experiments where they'll start playing some YouTube video, maybe that's vaguely political in nature.

They'll leave YouTube on and they'll leave it on auto-recommend where it just plays the very next video that it thinks you might want to see. And once you let it go for about five or six or seven videos, almost invariably, it's now automatically playing some kind of extremist content. And the reason that this happens is that these machine-learning algorithms have picked up on through years and years of users, certain behavioral patterns of ours. We are more likely to click on things or read things that connect with something that we're currently anxious about. So the algorithm has learned that if it plays a video that maybe preys on some kind of anxiety of ours, say preys on our anxiety about mass shootings or something like that, that they're more likely to keep people interested in the content if they stay on that line.

And in fact, if they show us something even more alarming, maybe something related to what we just watched, that that will increase that anxiety even more. It's not intentional in the sense that no human programmer sat down and thought, "how can I psychologically exploit these viewers," but it happens because the algorithms are programmed to keep people watching things. And the best way to do that is to increase our anxiety, to increase our interests by giving us things that seem meaningful, showing us events that seem like they may be related to one another.

**Christiane:** I understand how an algorithm can be exploitative even if there's no intention there, but kind of draw out the epistemic piece of that. How are we being epistemically exploited by machine algorithms?

**Megan Fritts:** So in cases of exploitation, it seems like an important aspect of what's going on is that there are at least two parties involved. It's some kind of transaction, some kind of trade off. Usually one of the parties is in some way made worse off by that exchange. Maybe overall, they don't end up worse off than they were before, but in some way, they end up worse off. A classic example of exploitation or potential exploitation could be like, people's worries about a market for human organs. The worries that the market would become exploitative, not because people wouldn't voluntarily sell their organs, but because those who would be inclined to do so are already vulnerable, enter the market in a vulnerable state and leave the market, at least in some ways, worse off. Usually worse off in terms of health outcomes.

That's a classic case of exploitation. So what does this kind of epistemic exploitation look like? Well, similar to the organ case, we have a transaction in the paper we argue that digital misinformation functions at least partially as a market, people are making lots and lots of money on misinformation, digital misinformation, I mean millions of dollars, and people are paying them, maybe not personally, but they're allowing them to get this kind of money by visiting their

website. Advertisers are really supplying the bulk of the money, but people are more likely to advertise on a site or YouTube channel that gets a lot of hits, that gets a lot of views, or a lot of reads. So we, as consumers of information and misinformation, are one of the transacting parties. We are getting something that we want, information or misinformation, they're getting something they want, data, on us.

And we enter this market in a kind of vulnerable state. Actually quite a vulnerable state. We enter the market as non-experts in most of the fields about which we're reading. This isn't our fault. It doesn't mean we're unintelligent. Knowledge has become wildly specialized over the course of the last century, far more so than it ever was before. And we're forced to rely on the words of experts really to know things or to act well. So we enter this market of information as non-experts already in an epistemically-vulnerable state. And what happens is that when this pattern recognition instinct starts to be triggered over and over again by phantom patterns, by spurious correlations, by misinformation, we leave the market, even though we've gotten something that we wanted, we've read the article we wanted, we saw the video we wanted, we end up leaving this transaction worse off, more vulnerable, whether we recognize it or not.

**Christiane:** I mean, I'm unfortunately such a black-and-white thinker. In my mind, I'm like, well, if that's the case, then do I just read books? Do I just read magazines? I mean, am I not allowed to read anything online at all?

**Megan Fritts:** It's not like this data's going away. It's not like the world is going to become simpler or we're just going to decide, you know what, Big Data isn't really serving us well in all the ways we want it to, let's pack it up and go home. That's certainly not going happen. So what can we do about this? Do we stop reading things online? That seems like a pretty big ask. It also seems pretty impractical to ask people to do that, and it's unclear that would really help. Some possible solutions that have been proposed are amendments to Section 230 of what's called the Communications Decency Act.

This proposal has not gotten a lot of academic attention, but it has gotten a lot of attention from politicians over the last probably four-ish years. And in fact, it seems to be a kind of bipartisan issue. People in both major parties have advocated, either changing or revoking this section of this particular act. So what is Section 230 of the Communications Decency Act? What it is, is a piece of legislation that says site owners can't be held legally liable for content posted on their site by other parties. So if I own YouTube and someone uploads literally anything, doesn't matter how horrible the content is, I'm not legally responsible for that as the site owner. Ironically, this part of the Communications Decency Act, which is drafted in the late '90s, was introduced in order to actually encourage content moderation.

This was the purpose of introducing it, but what was the case before, Section 230 is that a pair of court cases had set a weird legal precedent that if you do any content moderation on your site at all, you can be held legally responsible. If you do no content moderation, you won't be held legally responsible at all. And so Section 230 was meant to say, "Hey, you can do whatever you want, you won't be held legally responsible," as a way of getting site owners to do content

moderation. The worry is that because there's actually no incentives to do content moderation, and because there's so much more content now than there was 25 years ago, and it's such a larger task to do this content moderation, what Section 230 does essentially is, keeps the precedent in place.

It keeps the incentives there to not do any content moderation whatsoever. And so the worry is, well, this leaves the door wide open to this misinformation market. It's running rampant in part because no one is doing anything to stop it at all because they don't have to. There are probably far too many legal and pragmatic reasons that we can't actually get rid of the section altogether. Mostly it would make it almost impossible to own a website. But what we could do, is turn Section 230 into a quid pro quo. Turn these legal protections for site owners into a kind of trade deal where they need to do some amount of content moderation for extremist content, for democracy-undermining misinformation, for instance, in order to receive these legal benefits in return. So this is something that we propose as a possible, I don't want to say solution to the problem, but something that could help mitigate the amount of phantom patterns, the amount of epistemic exploitation that your average internet browser runs into on a day to day basis.

**Christiane:** There's a sort of backlash against so-called cancel culture. And in my mind, a lot of times I think people think of content moderation as cancel culture. Maybe I'm, maybe that's not right. I wondered if you had any thoughts about the backlash against content moderation, the backlash against consequences for extremist content.

**Megan Fritts:** Of course I would expect the largest question with regards to this proposal to turn 230 into a quid pro quo to be, how are we going to set standards for what is misinformation and what is not? Isn't the safest way to avoid authoritarianism, to avoid censorship, to just not have any legal definition or laws about misinformation at all? And I don't think this is a completely unfounded worry. Several years ago now, we saw something like this arising in the elections in Malaysia where these elections were one big issue that was in the middle of this election was a disagreement over fake news. So what had happened is the incumbent president had declared that digital fake news was illegal. That there were legal repercussions for printing it or writing it or accessing it.

But the problem was that there weren't very clear standards for what was fake news and the sitting president and those in the cabinet were calling things fake news that their political opponents thought, well, no, this isn't fake news, this is just criticism of those who are in power right now. I take it that this is the situation that most people want to avoid, reasonably so. So how do we come up with a good, maybe neutral understanding of, under what conditions is information misinformation, such that people can and should moderate it. And I don't know that there's a great answer to that question.

The best we can do is maybe appeal to some kind of what John Rawls would call public reasons. Some justification that the majority of people under ideal conditions when they're at their best self would agree this counts as misleading or democracy-undermining journalism and probably shouldn't be out there en masse. But I don't think that there's any possible way of

doing this in a way that's going to please everyone, and I'm not entirely sure that it's worth it to give up this kind of neutrality to moderate the content on websites. That's a trade off that it's hard for me to think about whether that trade off would be worth it.

So a further difficulty with solutions to the problem that rely on content moderation is the technological challenges of conducting this kind of vast content moderation. Doing content moderation through human moderators is probably an absolutely impossible task. The solution then seems to be: well, just use algorithms then to do this content moderation, maybe the algorithms can help us out this time. And a difficulty with that is that algorithms are really only as intelligent as the people who program them and are often easily outsmarted by human users.

So Facebook has in the past had algorithms to do some kind of content moderation and people who wanted to display their content that would have been moderated, or would've been kicked out by these moderators, simply found ways around this. They just find different keywords, different terms to use, to talk about what they want to talk about. So there's a lot of technological challenges to content moderating. Even if legally this kind of moderation was feasible, even if we got all the bills passed, could we actually do it, especially with huge sites like Facebook or Twitter? Could we actually conduct this content moderation? And I've been in conversation with some technologists working for various AI companies trying to ask them about this issue, and they seem as skeptical as I do, that this would really be something that we could do efficaciously.

**Christiane:** Has undertaking this work or writing about this work changed your own behavior in any significant way?

**Megan Fritts:** I think it definitely has. My co-author calls me a technological pessimist. I don't think that that's quite right. Technological advancements have been at least in many areas crucial and life preserving, but I think it has made me much more aware of the amount of time I spent online. I think it has made me more aware of my own vulnerabilities with respect to phantom patterns, with respect to being exploited epistemically. And it's also made me more okay with limiting the amount of information that I take and not feeling like I'm consigning myself to ignorance by not constantly taking in news media, rather, even though it's very contrary to our intuitions, I think that that can actually be a better epistemic practice than just taking it in constantly.

**Christiane:** So why do you care about this? What brought you to this work?

**Megan Fritts:** I think what brought both of us to an interest in this topic was having close relationships with people who saw the world very, very differently than we do. And by that, I don't just mean different worldviews, but rather having close relationships with people who thought that all of our sources for information were deceitful, were feeding us lies, and of course we thought the same thing about their sources of information. And the difficult part was in at least many of these cases, it's actually quite difficult to articulate a rational mistake that people who believe misinformation are making, because there's so much misinformation out there because they present evidence in the form of patterns to their readers, because there are so

many people in positions of authority, giving testimony that goes against what I believe is true on various topics.

They actually have sources of justification, epistemic justification, that are similar to the sources I have, but that are telling them the opposite thing. This is going to happen when you get so much information, so many people with a platform and the ability to spread their views. So we found that a lot of the discussion, especially in the philosophy literature on misinformation, was an attempt to articulate a kind of either a rational mistake or maybe bad character trait of people who do believe misinformation or find it compelling. And what we saw was that's not always the case. It's possible to be a rational, epistemically-virtuous person who just happens to find themselves in an environment where these rational capacities lead them into untruth. And that seems like a problem primarily with our digital environment, and what can we do about that?

**[Interview ends]**

**[music: Blue Dot Sessions, Pintle (1 Min)]**

**Christiane:** If you want to know more about our guest's work, or some of the things we mentioned in today's episode, check out our show notes page at [examiningethics.org](http://examiningethics.org).

Examining Ethics is hosted by The Janet Prindle Institute for Ethics at DePauw University. Christiane Wisehart wrote and produced the show. Our logo was created by Evie Brosius. Incessant clicking noise by the light that keeps off and on in this room. Our music is by Blue Dot Sessions and can be found online at [sessions.blue](http://sessions.blue). Examining Ethics is made possible by the generous support of DePauw Alumni, friends of the Prindle Institute, and you the listeners. Thank you for your support. The views expressed here are the opinions of the individual speakers alone. They do not represent the position of DePauw University or the Prindle Institute for Ethics.