



Opinion Science Podcast

Hosted by Andy Luttrell

Social Media Polarization with Chris Bail

June 7th, 2021

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

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

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

Andy Luttrell:

There are robots everywhere. And not like the housekeeping robot in the Jetsons. I'm talking internet bots. They're a little less exciting than the robots of sci-fi, but they might be more powerful. Bots are little bits of software that take care of automated tasks online. Like search engines use bots to take walks through the internet and pick up new websites. And bots can be bad—they can be used by cyber-attackers to take down websites and extort the owners. Spambots clog up your inbox. But other bots are pretty helpful, like ones that send you updates about the weather in your area.

But could bots help bring about societal change? The 2016 election was marked by lots of commentary about how politically divided the United States had become. The face off between Hillary Clinton and Donald Trump, and the aftermath of Trump's election to the presidency fueled what seemed like the most polarized public we'd ever seen. What made us so divided? And how could we ease those tensions?

A group of researchers thought that bots could give us a clue. One type of bot is a social bot—accounts on social media that aren't actually people but are instead little computer programs that strategically share certain kinds of information. Like, there's a Twitter account that just tweets a random frame from the TV show *The Simpsons* every 30 minutes. That one may or may not solve society's issues, but these researchers programmed two new Twitter bots to help them understand political opinion. One of them was programmed to mostly retweet messages that prominent liberals had shared and the other was programmed to mostly retweet messages from prominent conservatives. The idea was that if we could just break out of our echo chambers—see what people are saying outside our bubble—maybe it could chip away at polarization.

They paid a bunch of everyday citizens to follow whichever bot disagreed with their own political leanings. Of course, this wasn't super obvious. These people just thought they were getting paid to follow a particular Twitter account—no mention was made of politics. And for the first couple days, the account just shared pictures of nature landscapes, so it wasn't obviously a political account.

After a month of Democrats getting conservative tweets in their feeds and Republicans getting liberal tweets in their feeds, everyone completed a survey. And as far as they knew, the survey was

totally unrelated to the Twitter thing—it came from different people and had a totally different vibe.

So what happened? Did breaking outside of their political bubbles soften people’s views? No. In fact, it was just the opposite. Republicans who spent a month with liberal ideas in their Twitter feeds became significantly more conservative. And the more attention they paid to the bot’s retweets, the more conservative they became. It was similar for Democrats, but less pronounced. Democrats who followed a conservative bot became somewhat more liberal, but not significantly so. But they sure as heck didn’t become any less liberal after breaking out of their bubbles.

So there you have it. Bots 1, Society 0. What seemed like a viable strategy for easing polarization just created more of it. And the big question is: why?

You’re listening to Opinion Science, the show about our opinions, where they come from, and how they change. I’m Andy Luttrell. And this week I talk to Chris Bail. He’s a computational social scientist and the lead researcher on the Twitter bot study I just described. He’s a professor of sociology and public policy at Duke University, where he directs the Polarization Lab. Chris had a book out recently where he follows up on the bot study and explores the ways in which social media distorts our sense of other people’s opinions. It’s called *Breaking the Social Media Prism: How to Make Our Platforms Less Polarizing*.

I talk to Chris about his research, the themes in his book, and what hope there is for navigating political polarization online.

Andy Luttrell:

So, I thought one place to start is I think you might be the first true sociologist that I’ve talked to for this show, and I’m curious if you could do a little breakdown. In the book, you mention you define sociology, and as I read it, I thought, “Gosh, that sounds a heck of a lot like social psychology,” which is my bread and butter. So, I wondered if you could just even paint a picture of what sociology is, like what are the assumptions that make you a sociologist instead of some other kind of social scientist? And then also, what is the computational part of it? Because that seems pretty central to your academic identity.

Chris Bail:

Yeah. I mean, you know, sociology, I define sociology as the science of social relationships. And that’s a little bit distinct from social psychology. Certainly, social psychology is shaped by social relationships, but for sociologists, there’s a kind of geometry to social relations. We tend to do things like map social networks or think about how large groups of people interacting with each other create kind of systems-level phenomenon, so where we think of a lot of human behavior as emergent, meaning like what happens at one time and the characteristics of people at one time can’t necessarily be used to predict the outcomes at another time. And so, one thing that I think’s unique about sociology is we capture some of the complexity of social relationships, especially as they unfold over time.

And you know, ironically, I think social psychology if we go way back grew out of a lot of classical sociology. So, the first mention of social psychology that I’ve ever been able to find in the canon

is the work of Norbert Elias, a very long-lost founding figure in sociology who really thought that, yeah, what we really need is a kind of collective psychology in contradistinction to the then prominent Freudian stuff going on.

Andy Luttrell:

And so, where does the computational part come into it? So, sociologists, sociology 50 years ago looks different than the work that you're doing.

Chris Bail:

Absolutely. Yeah.

Andy Luttrell:

Because of the tools that were available.

Chris Bail:

Yeah. I mean, in some ways it's funny. In the 1950s, the question on the tip of every sociologist's tongue, or a lot of them, was mass persuasion. So, we were exiting World War II and there's other parallels, too, like radio was really coming into a golden era, TV was, broadcasting was becoming so prominent, and there were actually a lot of questions about technology and mass persuasion that a lot of people were looking at in interesting ways. Of course, though, the last half century kind of upended everything. In the social sciences, I think we often think of ourselves as kind of data poor. We don't have the particle accelerator and the colliders that physicists have, or the tightly controlled experiments that we might want to do a lot of the time, but on the other hand the kind of explosion of data not only from social media and the internet, but also from the mass digitization of human archives has really prompted some of us to say, especially the great Duncan Watts, a really prominent sociologist, to say social science has finally found its telescope.

And I think by that, he means that with all this data, we can finally start to unlock some of these population-level processes that were once kind of largely impossible to study. So, this new era that a lot of us are calling computational social science really is trying to work towards some of those large-scale questions using these new types of data sources. And also, of course, parallel advances in machine learning and computing power that enabled these kinds of new types of models.

Andy Luttrell:

Yeah. I like the idea of these new data stores being the opportunity for social scientists to finally go, "Aha! We've got a way to do some of this stuff that we've always dreamed of doing and thought we could try to do, and now we actually have a handle on it."

Chris Bail:

Exactly.

Andy Luttrell:

And not to maybe jump into the limitations too early, but it can feel like this approach is like, "Ah, finally, we can just solve all the problems that we've been thinking about." But what are some of the... Even just thinking about them at this stage, what are some of the challenges that we still have to be a little bit cautious before we go all in on the conclusions we can draw from these data?

Chris Bail:

Oh, there's so many. I mean, our telescope may be better described as like a crappy pair of binoculars right now. And you know, this happens in any kind of paradigmatic shift in science. We reached the peak of expectations when five, six years ago people were talking about for example getting rid of public opinion surveys, because maybe Twitter data could be used to take the pulse of the public, right? And one of the most obvious things in my research, and I've had a kind of circuitous path into computational social science, I've come out of a lot of qualitative work, actually in-depth interviewing work, and so one of the really fascinating things to me is to compare, use multiple methods to try to triangulate social science. And what we see very clearly there is this just stunningly huge gap between social media and real life in particular.

And you know, even though we might be tempted to say use Twitter, or TikTok, or whatever kind of data we can scrounge together to make sense of human relationships, really, we're just seeing like a small part of it, and it turned out in my research that understanding that gap became probably the single most important explanation of social media and political polarization, which is what I'm currently studying.

Andy Luttrell:

Yeah. I want to get to the new interview stuff that you've been doing for the understanding these dynamics, but to sort of get that rolling, we kind of have to talk about what instigated all of that, which is this big study that attempted to break echo chambers. And so, I'm curious where the impetus for this came from truly, to think like, "Oh, just all we have to do is just show people the other side of the conversation and people are gonna come around." Did it seem as obvious as I'm making it out to be at the time?

Chris Bail:

Well, I mean, we've known for 50 years that birds of a feather sing together, to use the kind of axiom of social science identified by the sociologist, Paul Lazarsfeld, in the 1940s and 1950s, right? People tend to surround themselves with likeminded people, and the concern, of course, translated into politics, is that this can create a kind of myopia where we can't see that there's two sides to every story, or we can't empathize with the other side, and so on and so forth. And you know, in 2017, when we launched the study that you were describing, it was a pretty tidy explanation of what had happened, right? Liberals were shocked that Trump got elected. People were shocked that Brexit happened in the U.K.

And the idea of an echo chamber, that it insulated us from opposing views, provided a really convenient explanation. You know, "Oh, it's just Facebook's fault. It was Twitter's fault." And you know, I was largely in support of that view at the time and had the hubris of many kind of data scientists of my generation and computational social scientists, to say, "Well, we can just dive into the Twitter data and figure this out." And so, we quickly realized that we could see patterns of echo chambers. Turns out there's a lot of questions about how big they are and how persistent they are, but we could see them. But what we couldn't know is whether people were kind of creating echo chambers around themselves or whether the echo chambers were influencing their views.

And so, we're kind of trapped in some circular reasoning. We can see that some people are talking the same way because they're surrounded by the same type of people, but what we'd really like to

do is break them out of their echo chamber. We'd really like to see what happens, as you said, when you take someone outside their echo chamber. And we did think that there was a good case that this would make people moderate. I mean, there's several decades of social psychology going all the way back to at least Gordon Allport in the 1950s, which suggest that when you encounter someone from another group, especially in person, you learn that the prejudice or stereotypes you might have had about them aren't true, and then you kind of... You change your views.

So, that was a very promising potential hypothesis. There was some indication at the time though too of what are known as backfire effects. You know, attempts to persuade people actually making them double down in their preexisting views. So, we did have some inkling that the experiment might not go the way that so much social science and social psychology might have suggested.

Andy Luttrell:

So, when you saw what the results were, to what extent did you go, "Yeah, there's that backfire effect," versus like whoa, that's not what we thought was gonna happen. Even if it was possible, I truly didn't think that was gonna happen.

Chris Bail:

Yeah. It was certainly a bummer. I mean, you know, because had we found that people moderated their views, then there's a simple solution. You know, just dial up the exposure to opposing groups. And at the time, this is something that Jack Dorsey told the Washington Post that he was thinking about doing, so we were really hoping for that. Whether... You know, the extent to which we were surprised, I guess the big question is does social media create the kind of conditions necessary for reasonable and rational deliberation and mutual understanding? And you know, when you begin to think about it, Ezra Klein when he wrote about our study, I think really nailed it. He said something like think about the last time that you encountered someone with an opposing view online. Were you calmly considering an alternative viewpoint or were you really pissed that someone called Alexandria Ocasio-Cortez a communist? Or if you're on the other side, that someone called Brett Kavanaugh a rapist, right? These are the types of things that are kind of typical on social media. Expressions of identity, not really expressions of ideas.

And so, we want social media to be a competition of ideas, but really it's a competition of identities. And so, that's really the crux of it, I think.

Andy Luttrell:

So, at what point did that then transition into the second wave of this work that is really what kind of centers and grounds the book project, which are these interviews? And I hadn't realized that that kind of qualitative work was something that you had a background in, and so this is a more natural direction than I thought maybe it was. And so, what was it about this that seemed like... You know, the only way we can really get a grasp on this is to do that quantitative computational stuff again, but complement it with this kind of in-depth qualitative interviewing method?

Chris Bail:

You know, I think again we social scientists are usually content to stick with what we're comfortable with. You know, if you run lab experiments, well then run another experiment. If you do surveys, run another survey. And we're not unlike medical doctors in that way. If you go to a

surgeon, you're gonna get surgery, right? Same kind of thing. And I think for me, the real inspiration to do this came out of this moment in social science where really, we're warring with each other about specifying the tiniest causal effect using the most valid design and recruiting principles of statistics that are almost impossible to observe at large scale in the real world, combined with the kind of deep interpretive tradition where the goal isn't really necessarily to identify general principles, but to explore kind of new possible mechanisms that can inspire future research. And in this way, I think the best kind of social science goes through cycles of deductive and inductive reasoning.

You know, we test theories, but you also need to generate theories. Otherwise, we wind up with this kind of well-known spotlight problem, right? It's like as if everything's dark and a spotlight illuminates one part of the world and that's what we study predictively, right? So, for me the qualitative work really was an attempt to say, "Okay, here we have a somewhat unexpected finding, and to be frank, we didn't know what was driving it." On the one hand, we had all this data. We had millions of tweets and could track people's networks as they're unfolding over time. We had surveyed them at multiple points in time. We had looked for patterns in the language that they were exposed to in our experiment when we paid these Republicans to follow a Democratic Twitter bot and Democrats to follow a Republican bot, and we looked at the data every which way using the fanciest techniques in machine learning and really there was no smoking gun.

And so, I thought, "Well, I don't know what it was like for these people to step outside their echo chamber and I don't even know who these people are. Why don't we go talk to them?" The thought of it, right? And you know, it's kind of a strange thing to do to do a qualitative field experiment, which is ultimately what we did, where really it's a mix of you're actually trying to randomize people in real life, so in this case we once again paid Republicans and Democrats to follow a bot that we created that exposed them to politicians, journalists, and media organizations from the other side, but we're also getting to know them before, during, and after, both online and off, through much more in-depth analysis of their social media behavior, but also getting to know where they grew up, what they think of Trump, what are their thoughts on climate change, what did they do the last time they logged into social media?

These are the kind of data that we just don't have when we look at patterns from 30,000 feet using the tools of computational social science.

Andy Luttrell:

So, how many of those interviews did you do? I forget sort of the scope of the qualitative part of things.

Chris Bail:

It's about 154 total. Yeah.

Andy Luttrell:

Okay.

Chris Bail:

Multiple interviews.

Andy Luttrell:

I hadn't thought of it as a qualitative field experiment. Is there any real precedent for that? Are there examples that come to mind of other people who have done similar things?

Chris Bail:

I mean, the one I love is Betsy Paluck, who's a social psychologist at Princeton, who did a very famous study of television messaging and ethnic reconciliation in sub-Saharan African, where she was kind of embedded in two communities that either were exposed to this TV messaging and were not. So, that was certainly an inspiration for me. You know, I think part of the problem is we don't have a lot of people who have training in both fields. Often, the people in the qualitative end of the spectrum really are very critical of experimental methods and think that, again, they're estimating tiny effects among tiny parts of the world. And you know, conversely, the people with the experimental skills don't really... maybe they're not comfortable talking to people. That could be part of it. But you know, also they think, well, that's a small N, right? That's a small sample size. I can't make anything of that.

So, you know, really, I think the most interesting questions are in between, and we need more people to try to find that middle range spot, which is another thing that sociology I think kind of somewhat uniquely contributes.

Andy Luttrell:

It's reminding me, I recently talked to Robert Cialdini for this podcast, and he has this notion in social psychology of what he calls the full cycle approach to social psychology, where you look to the world to get your questions, you bring them into the lab to take them apart, and then you go back out into the world to make sure you still understand what's going on. And there's part of this that feels like it's doing all of that in some ways simultaneously, where you're sort of just exploring qualitatively, while you're manipulating variables, while you're measuring stuff quantitatively, and at the end, presumably, you get a fuller, more accurate real picture of what's happening.

Chris Bail:

Yeah. It's really messy. You know, for example, like in our qualitative field experiment, one of the things that shocked me was one particular individual, his attitudes on race just completely flip flopped. He had fairly negative views about African Americans and then within a few months he had all these positive views of African Americans. I was like, "What was he seeing? What did the bot retweet?" I was desperately using the survey data to try to figure it out. It turns out he started a romantic relationship with an African American woman, right? This data point that you just don't see if you don't get to know people and get to fill in these blanks.

You know, there's many other examples that make the research messy, and then in the end you really realize, well, you sacrifice breadth for depth, and you really want to understand. You have to care about those idiosyncratic details in a person's life to understand exactly how they're experiencing your treatment. Maybe in that case, it made this particular guy more receptive to the racial appeals that the bot was making than the next person.

Andy Luttrell:

Are there... I was kind of wondering. At a few points, I think you mentioned like moments of surprise, or having learned yourself from these interviews. Are there things about U.S. public opinion that you ignored, took for granted, didn't appreciate until looking at what people actually said when doing these interviews?

Chris Bail:

Yeah. I think one thing that so many of us are guilty of in the ivory tower is to just assume that everybody shares our perspective on the world, especially in the realm of politics. If there's one thing we know from survey research, it's that most people avoid politics at all costs, and I certainly know that and was familiar with that research, but I think really understanding how for a lot of the people we met, sports, video games, TV shows have a lot more centrality to their view of the world and what they care about than what Mitt Romney said or what Hillary Clinton said. And in fact, they're often avoiding politics. So, one of the women that I profile in the book, this pretty interesting woman who I call Patty, and she's just... She's an unenthusiastic Democrat. She really... She doesn't really like either party, but if you ask her who do you vote for, she says, "Well, I usually vote for Democrats when I vote." So, she's not the type of person who we kind of idealize in social science, which is this rational creature who's capable of expressing a political opinion on any topic at any time, right?

She's like most people. She just... She doesn't care much for politics. She gets frustrated when politics comes up on TV. She thinks that both sides are unreasonable. And interestingly, even though she's a Democrat, she had some kind of conservative-leaning views, so she was mildly anti-immigrant. She thought she was worried immigrants were really destroying American culture and taking away American jobs. She was kind of concerned about government overregulation of the economy and taxation. In a lot of ways, she seemed like the type of person which journalistic accounts have profiled, the type of person who would be sympathetic to a Trump-style populism to someone who really could be turned.

And so, when we began to kind of turn up the dial of exposing her to more and more Republicans, I was really curious. Is this person gonna kind of buck the trend and maybe I finally found the person who will respond to this stuff? And what we saw was exactly what we saw in the quantitative experiment. So, instead of say calmly considering the alternate ideas of a David Brooks, a center-right figure, and maybe moderating her views on liberal issues, she tuned right into the most extreme parts of the continuum. So, you know, folks like Ted Cruz owning the libs and so on and so forth, all this kind of... In other words, she gets exposed to partisan warfare by stepping outside of her echo chamber, and somewhat paradoxically this then makes her feel like she has to choose sides.

So, again, not engaging in a competition of ideas, but a competition of identities. And really what we saw is this identity priming effect that came to inform a lot of the rest of the book and a lot of the rest of the research.

Andy Luttrell:

Yeah. There's a way that you frame it that realizing there's a war out there, and not only that, but like I'm in it, like I didn't realize that I'm part of this battle, so I better gear up and get ready to get in there.

Chris Bail:

Exactly.

Andy Luttrell:

One of the pieces of research recently that's really kind of reshaped the way I think about these things goes to what you're saying, and you said it in the book too, which is as someone who reads political psychology, who's at least somewhat interested in how politics unfolds, it feels like oh yeah, everyone has their side, everyone knows what they're doing. But the public opinion data that show like huge amounts of the public either don't choose liberal or conservative when offered the option, or they say, "I don't know," or, "I'm right in the middle." You go, "That's crazy." Based on everything that looks like... Oh, there's this huge war that's a brewing with all these extreme views.

Chris Bail:

Right.

Andy Luttrell:

And so, that transitions a little bit into the metaphor that you use throughout, which is what social media is, so you refer to it as a prism, which is sort of a contrast to this notion of just a pure, simple mirror that reflects out. And so, that contrast is grounded in this notion of looking glass self, which is a classic sociology idea, so I wondered if you could sort of give us a little glimpse into what looking glass self is, sort of just as a concept, and how social media is or isn't that.

Chris Bail:

Sure. You know, I think what makes us unique as humans is our tendency to cultivate identities, and to care so much about what other people think about us. I mean, if nothing else is distinctively human, I think that's probably it, right? And you know, this concept of the looking glass self comes around the turn of the 20th century, and basically the idea is as follows. Each day, we kind of wake up, we knowingly or unknowingly, we're presenting different versions of ourselves. We're kind of like, "Today I'm in professor mode. Tomorrow I'm in uncle mode." Whatever it is, we kind of enact these roles.

But then importantly, we're constantly scanning our social environment for evidence of which ones of these identities are kind of working, which ones make other people like us, and which ones give us a sense of social status. And then we tend to cultivate, again knowingly or unknowingly, those identities that make us fit in and make us feel good about ourselves. The great tragedy of human nature, especially from a sociological perspective, is that we're terrible at doing this, right? We constantly misread the reactions of others. We cultivate identities that aren't necessarily the ones that would give us the most status. And you know, but the key question is misreading the social environment. How does that happen?

And that's where I came up with the metaphor of the social media prism. We might like to use social media as a mirror, and again, what we were just talking about. You know, everybody's so extreme out there. How can everybody be not picking a side? It's because those of us who spend time on Twitter are seeing a profoundly distorted segment or portrait of the American population. We know that 73% of posts about politics are created by just 6% of Twitter users, and those 6% of Twitter users have pretty extreme views. So, the net effect is that we all think that things are much more polarized than they really are.

So, when we think about the looking glass self and our tendency to cultivate identities, and this is where the sociology comes in. This is where understanding how social relationships shape this identity cultivation process, we really need to think about how the structure of our platforms changes this all too human process. And I think it does so in two ways. The first is that it gives us unprecedented flexibility to try out different identities. You know, some platforms let us be completely anonymous. Others at least allow us to give selective accounts of what's going on in our lives.

And then second, we have these much more efficient means for monitoring our social environment. Likes, follows, comments, we have terms like getting ratioed, right? We have all these new metrics that kind of help us figure out whether our presentation of self is working. Now, of course, those metrics aren't the perfect measure of what most people would think of us. Rather, they're measures of what very active people on social media think of us. But we tend to normalize that, and we tend to misunderstand it as evidence that whatever presentation of self we're using, to borrow the phrase of the legendary sociologist Erving Goffman, we tend to cultivate that identity. And unfortunately, that has far-reaching consequences that encourage extremism and make moderates seem invisible, I think.

Andy Luttrell:

This sort of struck me at the beginning of the pandemic, which then I saw you come to that at the end of the book, which is I was so struck by the difference between the polling that showed that this wasn't a partisan issue. Everyone was like, "Oh my God, this is awful. We gotta buckle down and do something." And yet you go, "Oh, but it just looks like there's this partisan issue." The tricky thing that I am trying to grapple with is that that was a case where eventually the world caught up with the social media distortion. And so, is there a sense that that is what happens, that people do end up shaping their own views based on these extremists? Versus I think you kind of talked a little bit about how people might retreat and go, "I don't want to get caught in the middle of this. I'm just gonna get thrown under the bus for even expressing an opinion."

Chris Bail:

Right. So, how much we care about this social media feedback, and you know, for social psychologists, you might call this social learning. You can think of it as a giant Skinner machine, where we perform this thing and if we get the pellet of food, then we do it again, right? But we have to ask about the motivation. So, what types of people would care about gaining status on social media? And the answer, our research indicates, is people who lack social status in their offline lives.

So, the story that's most vivid in my mind from all the many, many people we met is a guy that I called Ray, and Ray is an interesting individual. You know, when you first meet him, he's very polite, even deferential. He goes out of his way to say criticize incivility online. He says he's a moderate conservative, but he says mostly he avoids politics. Then we go to look at him online, especially on Twitter, and I was shocked to discover that this is basically the most prolific political troll I've ever seen in a decade of doing this stuff, right? So, like who is this guy? And unspeakable, vile stuff that he's saying about Obama, and other leading Democrats, and you know, we go back to try to make sense of him, especially... and this is where the triangulation of multiple data sources becomes useful. You know, our survey firm can give us data on his demographics, his location. We're able to discover that he's actually a middle-aged man who lives with his mom in a city that's mostly liberal and a profession that's mostly liberal, so he's an outcast. Social media for him is a kind of refuge where he's gaining something that he can't get in his offline life.

And so, of course the kind of micro celebrity he's able to achieve online isn't very meaningful, like the people liking his posts are not having him over for dinner or asking for his autograph or anything like that, right? But it fulfills for him I think a kind of psychic need for status that we all need. So, what I'm not saying is that all of us have the potential to do this kind of Dr. Jekyll into Mr. Hyde style transformation that we see in this guy Ray. You know, the important question is where do you get your sense of status from, and for a lot of us, it's not social media. And to the contrary, talking about politics on social media is more of a liability. It can complicate the hard-fought status we've earned in our life and our career, or in our family, or friendships, or other kinds of personal relationships. And so, the other thing I do in the book is profile a lot of moderates and try to as you say explain why they really disengage. And of course, the consequence for all of us is that we've come to misunderstand these extremists like Ray as kind of representative of moderate people on the other side, when the moderate people have just largely gone dark.

Andy Luttrell:

So, your description of being able to know this guy so well did make me think, are there... Despite all the value of doing this kind of work, are there ethical tradeoffs? Have you gotten any pushback on like is this a little invasive that you can sort of stitch together this guy's online and offline life in a way that we can sort of make guesses about why he's doing what he's doing?

Chris Bail:

Sure. There are absolutely ethical issues in this new field, computational social science, and I think we've seen largely unethical experiments that were done without people's consent, especially on social media platforms like Facebook. But one of the unique parts about the researchers leading the effort rather than the platforms leading the effort is we are all beholden to ethical standards by institutional review boards. Now, these aren't perfect. They have to evolve to meet these changing threats to people's confidentiality. But, you know, we at the very least can ask for permission to do all of these things. So, people who participated in these studies were told that we were gonna have access to all of their data, so we could still, I think, rightly ask questions about privacy, but at the very least, I think informed consent is something that is really important whenever possible and almost always should be collected by researchers.

So, there are ethical issues here. A lot of them are complex, too, because especially once we start to do experiments that extend beyond the lab, that extend to real people in real life, you can't

control people like you might like to in a lab, and you can't control them, for example, talking to each other. And this creates a lot of challenges not only for ethics, but also for analysis of the data. So, we for example saw people ganging up on the bot and really egging each other on, so like actually observing peer effects, so to speak, in the wild. So, there are ethical issues. We're really... I always tell my students it's not just about satisfying the lawyers at your university and institutional review board. It's really thinking about unknown unknowns. And in a lot of cases, we've had to share less than we might like, which is why we go to great lengths in the book to change details of people's professions, and names, and geographical locations, and things of that nature to further protect their confidentiality.

Andy Luttrell:

I was wondering if you could talk a little bit about some of the prior work that you'd done, so you reference a few times work that you've done on I think it's anti-Muslim bias and how it sort of emerged, which is not something I'm super familiar with. So, maybe you could give a background, and also the ways in which it might have leaked into how you think about the polarization on social media.

Chris Bail:

Yeah. Yeah. So, my first book is called *Terrified: How Anti-Muslim Fringe Organizations Became Mainstream*, and really more broadly, I'm not an expert in Islam. I'm an expert in how ideas compete for legitimacy in the public sphere and that's really what I wanted to do with that book, is develop a theory of why certain ideas come to dominate public discourse. And this was again during the early ages of computational social science and I recognized that there was a potential to use algorithms to track this spread of text across the public sphere. So, basically what I did is amassed a large sample of press releases that were kind of competing to define Muslims. Are Muslims a peaceful religion or are they secretly harboring... are they an enemy within, right? Which was... There are many competing narratives about, or maybe some Muslims are good and other Muslims are bad, right?

So, I didn't have a horse in the race for which one of these ideas was right. I just wanted to see which one competes and how. And so, I recognized that by comparing these press releases to hundreds of thousands of newspaper articles, television transcripts, and other kinds of text, I could discover which ones of these narratives is picking up, getting kind of selected, and would up articulating kind of an evolutionary theory that explains how fringe ideas become mainstream through a process in which there's a kind of selection, an evolutionary selection of fringe, particularly emotional fringe ideas.

And that book also used this kind of mixed method approach, where I was seeing this patterns in the data of how these once fringe ideas were becoming mainstream, but also talking about the people who were actually enabling these transformations, the social networks that were enabling these transformations, the fundraising networks that were enabling these transformations, and so that was very much an attempt to explain how extremism becomes normalized, and certainly on social media we see that left and right these days.

Andy Luttrell:

So, the value of the press release is that these are organizations who are kind of spitting out possible stories and it's up to individual media outlets whether they're gonna pick it up, right? That's why it's sort of a which of these narratives is getting seized upon?

Chris Bail:

Exactly. And you know, we don't think a lot about ideas that die. We tend to focus on the ideas that survive. And those, of course, are in the minority. And so, when we try to do an epidemiology of ideas, we are often left with the lone surviving idea, right? And then we tend to generalize characteristics of that idea to a broader theory. So, for example, like why did Occupy Wall Street become so successful? Well, like 99% was a great slogan or something like that, right? But there were any number of other competing ideas about social inequality at the time, and that's in the end a pretty impoverished way of explaining why some ideas survive when most ideas die. And so, the cool thing about this type of plagiarism detection strategy that I used, I think, is that you're able to explain why most ideas die and why only a small minority survive, and then the consequences of that survival for the evolution of public discourse thereafter.

Andy Luttrell:

It reminds me, we relatively recently did the grad student selection process where you're often in that same boat, where you go, "All we have access to is information about students who already made the cut first, right?"

Chris Bail:

Exactly.

Andy Luttrell:

And then we can come up with all these theories about what it takes to be a successful grad student, but we were missing information about all the students who didn't make that cut for whatever legitimate or illegitimate reason.

Chris Bail:

Exactly.

Andy Luttrell:

So, it sounds a little bit like that, but sort of at this global scale.

Chris Bail:

Exactly. Yeah.

Andy Luttrell:

So, in what ways... Do you see parallels between what you discovered through those kinds of analyses and what you see just on general political rhetoric on social media?

Chris Bail:

Well, I think one thing that was happening as I was doing that analysis was social media was really coming into its own. So, in that book I tried in one of the penultimate chapters I think it was to see,

well, what does social media mean for all this? Because back then, this was the aftermath of September 11th, you still have media gatekeepers. You still have people whose job it is to decide what should count as news and these people have enormous influence over the evolution of public debate. And one thing that has so obviously shifted is sure, if the New York Times retweets something a lot more, if Fox News retweets something, a lot more people are gonna pay attention. But there's also a potential for virtually anyone to go viral and virtually anyone to start a far-ranging debate, and I think that is where we've seen this power of social media and the design of our platforms to create these kind of somewhat perverse incentive structures for status seeking.

You know, really, we're optimized for status seeking, which is not the same as being optimized for democracy, much less democratic deliberation. So, for me, that's the biggest change. And there's good parts of that change, too, right? Democratization of debate is probably a good thing, but it also means that there's much less control and that has obvious consequences for the evolution of our public conversations.

Andy Luttrell:

That's a better transition than I anticipated for the last thing I was gonna ask about, which is-

Chris Bail:

Okay, cool.

Andy Luttrell:

The Discuss It website.

Chris Bail:

Right.

Andy Luttrell:

Or app. And so, as though bots online, and big data, and qualitative interviews weren't enough, your lab thought, "Let's make our own social media platform."

Chris Bail:

Yeah. We are gluttons for punishment, I guess. Yeah.

Andy Luttrell:

So, that just even as a thing to do, I'm curious, did that... Where did that spring from? How did it actually get pulled off? And then also, what is it and why might it be valuable?

Chris Bail:

Right. So, I mean here's the big problem I think we face in this research space, is we need to ask fundamental questions about how the design of social media platforms are shaping social relationships. You know, things like the algorithms inside the guts of the platform, the publicness of the platform, how people are brought into conversation with each other, how they find each other, and all of these things are completely off limits if we go to the platforms and say, "Hey, Facebook. I'd like to make 300 users anonymous today. Could you help me out?" That's a quick and pretty easy no on their part for a lot of different reasons. Some of them reasonable. They can't

compromise their user experience. They don't want to deal with the PR issues. They don't want to deal with the legal issues. Possible ethical issues. It's just too much risk for them.

But on the other hand, we've been content to allow Facebook, Instagram, Twitter, other platforms to evolve in this largely ad hoc and chaotic manner, and nobody's asking the question if we could redesign social media from scratch, how would we do it? And so, we decided that if we wanted to ask these fundamental questions about how we should reform social media, we would have to create our own platform, pay people to use it, and then we would have full control over all the features of the platform and try to identify which ones might be polarizing and which ones might be less polarizing. And so, our hope in this effort is that we've created a space that not only our group, but other researchers can run all sorts of field experiments that really simulate a real social media platform.

So, we offered a large group of people money to complete a long survey about their views, and then a day or two later, they received a seemingly invitation, seeming to be coming from another group of researchers, that basically invited them to test, "test a new social media platform." And so, we didn't give them any incentive to use the platform in one way or the other. We simply gave them an invite code, which unbeknownst to them paired them with a member of the other party to engage in an anonymous conversation about either immigration or gun control. And we wanted to study anonymity because there's this interesting puzzle with anonymity. On the one hand, anonymity enables that guy Ray, the Dr. Jekyll into Mr. Hyde guy, who hides behind anonymity and says things on social media that he'll never say in real life, and so there's an obvious danger of anonymity. It can lead to some of the worst human behavior. We've all heard like don't read the comments, right? That's because the comments are usually anonymous.

On the other hand, if identity is really spearheading our experience on social media and it's such a key motivation of how we use social media, then anonymity also has this feature which is allowing us to take identity out of the equation. And if we do that, the thinking was perhaps people would be able to focus more on the content of each other's ideas, rather than the identity, like the us/them dynamic that we see so clearly on most platforms.

And so, what we did is we compared people who used the anonymous app to people who didn't, and we were pretty surprised to see that people who used the app depolarized a lot. And perhaps even more surprisingly, Republicans depolarized at about twice the rate of Democrats, suggesting I think that some of what we're seeing is the result of peer pressure. You know, Republicans may feel because of the orthodoxy and the strong ideology attached to their side it's really hard to explore unpopular ideas in the context of a very public social media space like Twitter or maybe even Facebook. And so, this suggests that maybe one of the solutions is to create some spaces... You know, I'm certainly not saying that we need to make Facebook anonymous or Twitter anonymous, but that there could be a new kind of social media site that enables anonymous conversation and perhaps incentivizes people to convince each other across party lines, creating a new kind of status, right?

We can't really... The funny thing about identity, I think, is you can't tell people, "Stop being extreme." That has the opposite effect. But what you can do is try to create new types of identities and new forms of status and then kind of nudge people to try to pursue them. And so, I think a

structure like this would be really interesting for that small part of the population that actually wants to talk about politics. Obviously, there's a huge part that would not have any interest in talking about politics anonymously online, but that's always been the case. So, why not allow political conversations to splinter onto a new platform much like we've allowed our hobbies to go to Pinterest, or our videos to go to TikTok, or whatever it is.

Andy Luttrell:

Did the app look good? I always think about when psychologists design video games to test. You go, "No one would voluntarily choose..."

Chris Bail:

We spent a lot of time working with a professional graphic designer and a lot of UX or user experience people to try to make it user friendly. I think in the end, it succeeded. It simulated like a reasonable upstart social media company. We even had to staff a user support line for a week, which was fun. But you know, yeah, these things... This final issue of translating science into public tools is a real challenge, and for this case, we were really just trying to run a research study. But the other mission of the polarization lab is to create public tools that anyone can use to kind of try to attack polarization from the bottom up.

So, for example, we have tools that allow people to identify and avoid extremists and boost moderation, so that we can all kind of become aware of this tendency for false polarization to occur, for us to overestimate the extremity of the other side and lose sight of moderation. We have bots that retweet people on a bipartisanship leaderboard. We have other tools that allow you to, for example, estimate what your politics look like to other people based upon the content of your tweets. Things like that, because you know, at the end of the day, even though these top-down solutions would be great, Facebook and Twitter aren't going anywhere, right?

And so, yes, there's things that Facebook and Twitter could do. For example, they could optimize or boost posts that appeal to people across social divides instead of kind of preaching to the choir, which is what we're seeing right now. But you know, at the end of the day it's gonna be up to us social media users. We are the supply side of this problem, you know? And at present, I think an underappreciated amount of polarization is coming from us. And so, trying to encourage a more thoughtful introspective social media user, but also give them tools that allow them to enact and make non-polarizing habits into core parts of their daily experience on social media is I think where we need to go.

Andy Luttrell:

Nice. Well, I will keep my fingers crossed that we save ourselves from full collapse.

Chris Bail:

Yeah. Yeah.

Andy Luttrell:

And we make social media a pleasant place to be. Just wanted to say thanks so much for the time to talk about this. This was really fun.

Chris Bail:

Sure. Thanks so much for having me. It was a pleasure. Thanks.

Andy Luttrell:

That'll do it for another episode of Opinion Science. Thanks to Dr. Chris Bail for talking about his research. Once again, his book is *Breaking the Social Media Prism: How to Make Our Platforms Less Polarizing*. You can find a link in the show notes. You'll also find links to Chris' website and the Polarization Lab at Duke University.

For more about this show, head on over to OpinionSciencePodcast.com. Follow the show @OpinionSciPod on Twitter and Facebook and anywhere you get podcasts. And if you like this stuff, it would mean a lot if you left a review on Apple Podcasts or any other app.

Learn more about the science of opinions and persuasion with one of my online courses. More info at OpinionSciencePodcast.com.

Okay, that's it. Thanks for tuning in, and I'll see you in a couple weeks for more Opinion Science. Buh bye!