



Opinion Science Podcast

Hosted by Andy Luttrell

Geography of Bias with Eric Hehman

November 23rd, 2020

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

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

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

Andy Luttrell:

It's Thanksgiving! At least, it is the week this episode comes out...in the United States. If you're not familiar, this is an American holiday that is food-centric. As a vegetarian, the turkey at the center of the table wigs me out, but I—like many of my fellow Americans—am all about the side dishes.

But what's a traditional Thanksgiving favorite? Well, it depends on where you're from. The website Zippia consolidated data from Google searches across the country in November 2019 to see what Thanksgiving side dishes were being searched for most in each state. Mashed potatoes are a big favorite—from California to Minnesota, Illinois to Connecticut, mashed potatoes are the top pick. But then there's Louisiana's love for cornbread dressing, Florida's penchant for sweet potato casserole, the creamed corn obsession in Kansas, and Maine's feverish search history for...side salad. Even within quote-unquote "American culture," there's plenty of variation across the states.

This website has made other fun maps as well, including the states that love pickles the most—people in Maine love pickles and people in Hawaii don't care for them. There's also a map for sandwich preferences—from the understandable Lobster Rolls in Maine and Po Boys in Louisiana to the PB&Js of Nebraska and Bologna sandwiches of Ohio.

And sure, an unscientific analysis of Google trends is one thing, but the geographical spaces we inhabit are tightly tied to our psychology. Election maps show the political leanings in one district versus another. Classic research on the American South unveils its unique psychology, and analyses of many countries around the world reveal important tendencies among people in one place versus another.

New research is pushing this perspective even further, looking at how people's prejudices clump together geographically. And it raises an important question...where does prejudice live? In the minds of individuals or in the cultures they inhabit?

You're listening to Opinion Science, a show about the science of our opinions, where they come from, and how they change. I'm Andy Luttrell. And this week I talk to Eric Hehman. He's an assistant professor of psychology at McGill University. I met Eric a few years ago and quickly

became a fan of his work. And I'm not alone! One of the studies he talks about in our conversation just won an award from the Society for the Psychological Study of Social Issues.

Eric's been studying how prejudice varies from one place to the next. What's the average prejudice of Illinois versus Ohio? What about in one county in Ohio versus another? He and his colleagues look at these patterns and see what they're connected to...and when they change.

A quick bit of context before we get into things. First, we talk a bunch about implicit and explicit bias. You can think of that like prejudice. How much does someone have a preference for their race, gender, religion, etc., over another? Explicit bias is when people openly say they have this preference. Implicit bias is when the preference is automatic.

And we talk about this thing called "Project Implicit" a bunch, but I'm not sure we're ever clear about what that is. So I'll just tell you—it's a website where anyone can go and take various tests that measure implicit and explicit bias. The website's been around for years, and a ton of people have visited and given data for research. I'll put a link in the show notes if you want to check it out. For more about implicit bias and project implicit, check out Episode 16 with Mahzarin Banaji.

But this isn't Mahzarin's episode—this is Eric's episode! So let's jump right into our conversation...

Andy Luttrell:

So, I guess kind of the first question, just maybe to give an overview of the parts of the work that you've done that I'm interested in, because you've done all sorts of different things, but there's a thread through your work that looks at sort of regional variation in people's mostly racial biases and how those are connected to things about those regions, so I'm curious, even just to begin, like where did the idea of looking at that come from in terms of aggregating to this bigger level? Because psychologists were trained to look at you are a person and you have a bias or you don't, and this is quite a different perspective, so I'm wondering where that came from.

Eric Hehman:

I want to say it was due to The Guardian, which is a U.K. newspaper that you may be familiar with. And they had put together this database of police killings, essentially, and this was important because the police generally don't like to share this information in a way that's accessible to other folks, and that's why The Guardian was putting this together. And just kind of looking at it, they had a bunch of basic information about who was being killed and where they were. We were just kind of inspired that we could essentially address a very basic question that was like a regular, recurring narrative is to what extent racial bias might be involved in these police officer killings.

The other large piece, I think James Rae, who was a student working with Kristina Olson at the time, had put together this paper examining how bias was related to segregation, and he had used data from Project Implicit, as you know about, this massive database, a really unprecedented database in terms of the amount of psychological data that it has, and that has location information, as well. So, those are essentially the two ingredients. We could find and geolocate those police killings and we could geolocate these estimates of bias to essentially examine a correlation between the two of them.

And once I started thinking spatially, like that project took forever, and there was a lot of infrastructure to build up, and just starting to think about what this even means on a theoretical level, developing all that, I want to say we started that project in like 2015 and it was my first work doing anything spatially, and it wasn't published until like 2018. So, it was a long road, and a lot of other projects that I began well after the police killings paper were published in that interim. But that was what got us all started and kind of set me down this path, and at this point it's about half of my research program.

Andy Luttrell:

It strikes me that kind of that idea originally came from just these data that already existed. It's this rare instance of we could actually use psychological methods to understand things that are actually happening and actually unfolding in the world, which is something we don't normally get to do, right? To use the kinds of precise measurements that we like and also understand messy things that happen just naturally in the world and sort of this rise of resources, both the Project Implicit website and this Guardian project just sort of served you up this opportunity to connect the dots.

Eric Hehman:

Absolutely. And since then, a big portion of this work is just finding convenient data or data that's available. Sometimes it takes a lot of work to pull together, but it might be worth it. I do think this is a changing landscape in terms of more and more information is published online in forms that researchers and people can use such that this is gonna be a continuing thing, and more and more questions will be able to be asked, and that's why I'm excited about it. But a lot of times, you really want to ask a particular question and the data is just missing one tiny piece, and it doesn't let you really connect those different pieces of information to ask that question.

Andy Luttrell:

So, to maybe get a little more specific into what you found with the police killings, so I'm not sure if you said exactly what the relationship was, but if you look at, like what are you looking at on the person level, like what kind of information are you getting from that website, Project Implicit? And then what is it connected to in the world?

Eric Hehman:

Right. Exactly. So, a regular process is that all these individuals for this study, all across the United States, have visited the Project Implicit website, they've completed the implicit association test, which is this reaction time based measure of bias, as well as often an explicit rating of bias, as well, which is really just answering questions of like how much do you like, for example, Black people relative to white people. And we can take those measures and we, at the beginning, we had their IP addresses and limited geographic information, and we could geolocate them through a series of tools. Nowadays, they have that kind of regional information that's built into Project Implicit, making this process easier.

So, we can essentially place these individuals in space and time, and then if there are say like 10,000 people in a given area, we can essentially average them together to get a sense of what is the regional estimate of bias in this area? And the broader idea here is that any one of those

estimates may have a bunch of error involved, or individual variation, but if we average them all together, that random error or individual differences gets canceled out and what's leftover may be like a pretty good estimate of the culture or the bias in the region.

So, the unit of analysis is no longer people once we've done this averaging process for this approach. The unit of analysis is now like a given region. And once you get it to the region level, you can link up this measure of bias with presumably an infinite number of variables that are also at that region level, or that you've got to that level through various other processes. So, in this particular situation we also did the same process with The Guardian, the counted data, looked at the total amount of killings that happened in that area. We did control for population differences. So, for instance, if there's two Black men killed in situation A, versus situation B, if there's really just 20 Black men in situation or region A, versus like 2,000 in B, you would expect differences just based on the amount of population in the first place, so that's baked into our estimate. And then we just found a relationship between the implicit and explicit biases of white people in the region with the amount of Black people that were killed by police.

Andy Luttrell:

So, you have this, the average bias of an area connected to these things that are happening, that we would say are biased outcomes in an area. And so, when I think about that, you go, "Wow. That's so amazing." And then I stop to think like what does that actually mean? What does that signal? What does it mean that an area has bias and why would that be related? Because the other thing about this that I think is so interesting is that it's not necessarily police officers who are in the data set that are giving you that idea of bias, right? So, normally people might talk about like, "Oh, the person is acting with their own bias." But we might even getting an indication of that place's bias without ever considering the person at the center of that event that we care about.

So, what does it mean that an area can be biased?

Eric Hehman:

Absolutely. I think that's a really important point and all the work prior to that, understandably, was focusing on police officer biases. What I liked about this is that it was more of the... It showed that it was part of the local culture that might be contributing to this, rather than like the onus being on the individual police officers in the first place. So, what does it mean for a culture to be biased? I would say that's like a big, broad question that we're still working on, but there are a few different reasons why an area might be more biased. And this was nicely put forth in a paper by Peter Rentfrow.

Essentially, one possibility is that more biased people move to a certain area, which is why there are differences across different regions, meaning there's less biased areas and more biased areas. Explanation A would be like a bunch of biased people move here, a bunch of unbiased people move there, and that's why we are observing these differences. Another explanation is just that it has to do something with like the local environment when they get to it, and when I say environment here, I'm gonna be talking about like weather, mountains, literally the physical environment. That is a possibility.

And then the third possibility is related to this, but it's more about like the social environment. So, you move to an area, maybe there's a pre-existing culture for whatever reason, and that culture exerts an influence on you whether you were raised in that culture or whether you're a newly arrived person to that culture. And maybe these slowly shift over time, but I don't think that that would be a really rapid change.

So, I imagine to some extent both the first one, selective migration, is happening. Matt Motyl had a really cool paper on ideological migration and people moving to more conservative or liberal areas. But I don't know if it'd be like the meat and potatoes of this. I think a better explanation is more of like the social environment, whether that's... I don't know. You move to an area and there's local threats, you perceive threats to your economic outcomes, like you feel a threat to your job, you feel a threat to your culture because you're living in proximity to maybe another group, or you move to an area where there's just like lots of biases already present and that influences you in some small way, giving rise to these different clusters of prejudice or egalitarianism.

Andy Luttrell:

As you describe it, it reminds me, I mean obviously of Keith Payne's idea of bias of crowds, which I know there's been lots of kind of back and forth between the work that you've done and his work, and I don't know if I just came up with this metaphor or I read it in his paper, but it reminds me of something about a weather versus climate thing, where an individual person's bias is kind of like the weather on a given day that can be super variable and all over the place. But it kind of sounds like what you're saying is that when you average across all these different factors that might go into a place's climate, you're getting some sort of stable reflection, right?

So, I can measure the climate of a place and that'll tell me sort of in pretty solid terms what that place is likely to be like in a way that one or two measurements of weather aren't gonna do that. Do you know? Did I make that up or is that his?

Eric Hehman:

I've read that paper pretty closely and I don't think that analogy is in there, so I'm gonna credit you with that one. Yeah, I like it, and I think it perfectly captures what Keith is going for and what I believe to some extent.

Andy Luttrell:

So, was there always a connection between the work you were doing and his work? Because I sort of came to them at the same time and it just felt like obviously these were both born of the same stuff. But it sounds like maybe they were happening independently and just happened to come to similar kinds of ideas.

Eric Hehman:

I don't know. I haven't talked to Keith about this, and I'm really hesitant to say this, but I think maybe he was inspired by some of our work. We had maybe like four or five papers out there before the Bias of Crowds came out. I'd spoken with Heidi Vuletich, who is on that paper, as they were developing it, and some of the techniques that they used in that paper came from stuff that we were working on at the time. But in turn, we have been very inspired by that theoretical paper,

and some of the stuff that we're currently doing are totally based off that, so I like to think it's a back and forth.

Andy Luttrell:

Could you describe what that perspective is? It sort of maps onto what you were saying, but sort of what does it mean? As best you can tell. I'm not asking you to sort of summarize for him, but in your interpretation, their idea of a bias of crowds. What is it about?

Eric Hehman:

No, totally fair. So, Keith in his paper presents a bunch of phenomena that are associated with implicit bias. These tend to be that it has a lot of measurement error, meaning that it's very variable, even in a single individual over time. You take a measure of implicit bias at time A and a measure of implicit bias at time B and you will get like pretty decently different scores. Also, it doesn't... It seems to develop pretty early in the developmental process, like children are developing an implicit bias. It seems to be relatively consistent.

And then he also brings up that even though on an individual basis there will be a lot of variability, if we just gave a measure of implicit bias to 100 people over and over again, on average we would generally find say a pro-white bias. And that's like an extremely robust and replicable finding. So, he kind of presents a number of these phenomena and then puts forward a solution that might kind of neatly solve or address some of them, which is that implicit bias isn't really something that an individual harbors. I mean, it is, but it's more influenced by the context or the environment. And social psychology has always paid a lot of attention to how the context and the environment will shape our attitudes, but I would say this paper is kind of refocusing a major contributor to our implicit biases to the environment.

And this means, for instance, that if we're in situation A, and our implicit bias is like... I don't know, a 0.2. I'm just throwing out a number. And then we move to another environment, it might shift as we are maybe subtly influenced by both the people around us, maybe like the structural biases that are around us, that cause us to think in a certain way. And they present some evidence of this as a reanalysis of other data that's existing in that paper.

So, it was a really beautifully crafted theory kind of addressing a bunch of the puzzles that implicit bias researchers hadn't really figured out over the years and it pulls them all together. I will point out there's other people that are kind of pushing back against this and it was like a really interesting psychological inquiry paper, or article, where the model is like a single theoretical paper and then a bunch of responses, and then a response to the responses. And the responses continue.

Andy Luttrell:

Well, it's interesting to me, so as someone who is more focused myself on just attitudes and opinions in general, my impression of that work and other work that looks at this geographical perspective is one of kind of curiosity that so much of it is looking at these kinds of social biases. Race bias probably most predominant. And so, I wonder how much of that is because it's just people who were already doing racial bias stuff that then took on this perspective, versus there's something about a social bias that is most amenable to this geographical perspective. As opposed

to like I prefer pie over cake and that's probably my own preference and I'm not picking it up from the Ohio air that I breathe.

Eric Hehman:

Maybe, but maybe you've been... Maybe you've lived a pie-saturated life.

Andy Luttrell:

I could have. Yeah.

Eric Hehman:

They've been around influencing you subtly over that time. Yeah, maybe. Certainly, that is the case for me, like I was already interested. I'm interested in intergroup stuff in general, but often that that is race based, and so that was why a lot of our initial questions were going that way, but there has been... We've done some anti-gay biases work, still social. Peter Rentfrow again and Sam Gosling do have a paper looking at spatial distributions of personality and... But I feel like for the most part, so far it's been adopted by social science researchers interested in prejudices, but maybe that will change over time.

I do think spatial approach in general is like relatively newer to psychology, and I think questions like this will continue, but maybe... It'll be interesting to see where it branches out. I don't see things like self-esteem or I don't know, say things like cognitive dissonance varying in meaningful ways spatially, but I've done some work with Michael Slepian, who studies secrecy, and just like he's a good friend of mine, and we were just kind of poking around at some data out of curiosity, looking at whether... like what secrets people hold seems to vary geographically, and it does. So, maybe when we're talking about personality factors that might be influenced by our local environment, we'll be more likely to see spatial variation.

Andy Luttrell:

I wonder too whether the kind of region you're looking at matters. So, in this paper that is relatively recent of yours, you sort of go through the gamut of looking at counties, versus looking at this sort of... I don't pretend to understand what the core based whatever that is called is versus a state level, and so does it... Those are obviously very convenient boundaries, but is there any sense... I mean, you've actually seen differences between those boundaries in terms of how reliable the measures are, so kind of what might those differences be and what do they mean for us in terms of thinking about the geography of bias?

Eric Hehman:

Right. Yeah, so there's an issue in this area of research and others, but it's very salient here, called the ecological fallacy, which is the idea that you might observe a particular relationship, say at the level of the individual, but then you might observe a very different relationship if you had those same variables at the level of say the state. So, in all of our work we have thought it's important to kind of like test at multiple levels, to see is this effect fairly consistent across all of them? And I think what region you want to use in this work depends on your question, so we have a paper on how same sex marriage caused changes in local anti-gay biases, and we did a state-level analysis for that work, because these laws were being passed at the state level.

And so, to use a county level, when there's like a broader thing that's applying to all counties, wouldn't really make sense there. But they do have different pros and cons, so if we stuck with a state-level analysis, for instance on say prejudice, I think Texas would be a good example, where Austin is like this liberal bastion surrounded by a sea of red. You're kind of like glossing over really big differences on smaller units that are within that broader region. The core base statistical area, one that you brought up, which you can essentially think of as like cities and their suburbs, tries to get around this, and that's why it's been constructed specifically. So, these have been defined in a way, it's like the people living in those suburbs are commuting into the city.

So, if you were to draw the boundary at where the city actually ends, you'd be kind of like missing the people that are just in that city all the time and surely influenced by that. And in spatial analysis, this is known as like the aerial unit problem, as like if we're gonna call this thing a unit, are we capturing all of the people that are in it? Are we inappropriately drawing a line between two groups that are culturally similar? We don't really have this issue with people, because we are just self-contained little units walking around.

Andy Luttrell:

Boundaries are clear.

Eric Hehman:

Right, right. So, it's unique to this sort of analysis.

Andy Luttrell:

Are there things about a particular place too that might make it more likely to show these regional biases? So, I don't know if you've looked at like in general, one example would be this really nice graph that is coming to mind of sort of state by state, the relationship between implicit and explicit biases, right? And you just see this nice little scatter plot, where the states that tend to be lower on implicit bias also tend to be lower on explicit bias. And so, that is treating each state as though they contribute equally to the equation, right?

And so, what I wonder in terms of research on culture, we know that there are these differences between tight and loose cultures, places where people are more likely to conform with whatever sort of the cultural leanings are, versus cultures where people are more likely to adopt just their own idiosyncratic perspectives. And so, part of me wonders whether you might see variation in how good or how useful the aggregated data are depending on where those data are coming from. Does that make sense?

Eric Hehman:

Yeah, it does. I think it's an awesome question. So, we're pretty loose, right? In the United States and North America in general. But other cultures that are maybe more tight or more homogeneous, I do think that questions of like whether regional estimates of anything are gonna predict outcomes. Really depends on having that variation in the first place, so I can't imagine, say we go to a place that's like all the different regions are fairly homogeneous, everybody has the same attitude, I can see some of the relationships that are demonstrated, say in the U.S. or Canada, not holding in those areas, perhaps.

I think it's a great question whether variation in attitudes regionally even kind of maps onto this looseness versus tightness idea. But we don't know yet. Most of the work is based on just like what data is available. The U.S. is pretty good about data availability. We're doing work here in Canada, as well, and in the U.K., and Europe more broadly. But it's harder to get access to this sort of data to ask these questions. I mean, definitely not out of China, who doesn't really share any data. But like in other regions more broadly, I think that's where this should go eventually.

Andy Luttrell:

I was gonna ask, you mentioned the study with same sex marriage legalization, which is another just very cool example of what this kind of perspective brings. So, would you mind kind of just summarizing what that project did and what it found?

Eric Hehman:

Absolutely. So, we became interested in the role of government legislation causing changes in bias, and even zooming out a little further, this is a broader question of norms. And can government laws be perceived as norms in a way that are gonna be influencing the attitudes of people who exist locally in that area, governed by that norm? So, the norm that we're looking at here is whether it's appropriate or legal for same sex people to get married.

The U.S. was a really great place to study it in a causal way, so many folks have learned like correlation is not causation. In questions like this, we can't manipulate the variable, which is like what is the norm in your area. We can't manipulate the law saying it's okay for same sex couples to be married. But there's like 50 different states and they all passed same sex marriage legalization in some functional form at different periods of time. And that's great for a researcher, because it really lets you rule out alternative explanations and you have like 50 different units in which you can essentially compare like what was going on with bias ahead of time and what was going on with bias afterwards?

So, because of this design, in terms of observational data, lets you make like a really strong case for a causal argument that the law is causing changes in bias specifically. And that's essentially what we did, so we looked at there's both implicit and explicit measure through Project Implicit. We also incorporated data from the American National Electorate survey. This is a different source of data asking about attitudes towards gay men and lesbians. And we just plotted these trends over time, like before the law was passed and after the law was passed, essentially across all 50 states, and we found that while anti-gay bias was decreasing even prior to this law being passed, following same sex legalization, it was decreasing at a sharper rate.

And so, because of this design, we can attribute that change in slope, essentially, or the rate of change over time to this legislation.

Andy Luttrell:

And the time scale is pretty quick, too, actually. Right? Like we're not talking about centuries of change that need to happen. And it's not even decades, right? I mean, the data you have are within, what, a couple years? As those laws rolled out? And so, that's a pretty quick downturn, right? Like probably, your data probably can't speak to it, but even within the next, what? The next week or two weeks? When were you able to capture that downturn?

Eric Hehman:

Yeah, I mean the trend after the law or legislation has been passed locally is kind of estimated on everything afterwards. So, essentially we're saying that that is happening immediately, though of course, that's not happening immediately for everybody. And it does seem to be happening fairly quickly. I do want to point out that it seemed to be responsible for about like 3% of the variance overall, which at first I was quite disappointed by, and I think getting used to running experiments in the lab, you get used to your models essentially explaining... I don't know, 60% of the variance. 3% of the variance seems like a really small number in terms of how much it's influencing overall people's biases.

But for that paper in particular, kind of in the review process, we started looking up field interventions and what is a reasonable number to expect in the first place, and that 3% number is like pretty great, from what I understand. So, there's considered a gold standard intervention for young Black boys living in cities, and it's about reading, and it's like... receives millions of dollars in funding from the federal government, generally supported by lots of research and people in that area, and it's responsible for about 3% of the change in terms of those individual children's reading ability.

And just like another example, one that I always like, so if you were a baseball manager and you had to put in your best hitter versus your worst hitter for like the final bit of a game, that decision is like .0003% of the variance in terms of whether there's going to be a hit or not, but this is something that like every baseball manager would endorse everywhere. So, even though these things can have tiny little effects when we think about them as a percentage of variance, I guess you can think about this effect influencing people's biases in the presence of like the whole slew of other things that are happening in their life, in their culture, in their personal variables and what they have going on, and so it doesn't seem so bad when I think about it like that. But I want to keep it in perspective, I guess. We're not changing these people totally. They're just like little, tiny tweaks.

Andy Luttrell:

But even still, I mean the comparison to that probably high-intensity intervention on reading skills, right? That's like working directly with kids to improve their reading. This is change a law that may not even affect your life directly and all of a sudden the sort of trajectory of people's opinions takes a turn. Maybe not a huge turn, but that it can happen just at sort of like a decision off into the distance that doesn't... No one reached out to me. No one did anything, tried to persuade me necessarily. It's just now there's a change in the culture and my views might take a slight change, as well. That makes it all the more impressive, I feel like.

Eric Hehman:

Huge. Yeah, and I think I've really, through this study, came to appreciate the value of norm research more. This has been a thing around in psychology forever. But you just... I just started like seeing and appreciating more in terms of like evaluating our entire field, and like what are the big, giant effects that are consistently moving around humanity and their attitudes? And I think norms are a huge one, and I think we can observe the opposite direction, as well, though I don't have firm data on this, like norms conveyed by political leaders that are maybe endorsing

xenophobia can push around attitudes in the opposite direction. And I would bet there's a lot of evidence out there for that to be the case right now, though I haven't seen a paper.

Chris Crandall has done some work looking at how pre and post the election of Donald Trump, people were much more comfortable expressing prejudice towards a number of groups that have been targeted by Trump, and a lot of his campaign rhetoric, and not towards the group that he hadn't targeted during his campaign rhetoric. But no official paper exists that I know about yet. It goes both ways.

Andy Luttrell:

It raises the question, though, that idea that people felt more free to acknowledge a particular opinion is another question I had for you, which is when I had thought about this kind of geographic work, it often seemed about implicit biases, probably because the data come from Project Implicit. That is just sort of the place where implicit bias data come from. But a lot of what you find, you also get for these explicit biases. So, just to kind of back up a little bit, implicit bias is this idea that people can have these subtle, automatic preferences for one group or another that we can capture with a simple, like you said, reaction time task, which we never have to ask people, "Do you prefer one group over another?" But we could just ask people, right? Do you have a preference for this one social group over another?

And so, I wonder, my impression is that often you're seeing similar things for both of those ways of getting at bias. So, I guess the question is, is that a fair assessment? And if not, is there a reason why we might expect to see this kind of aggregated value more for one kind of bias than another?

Eric Hehman:

Right. Yeah, so I have a stance that implicit and explicit estimates of bias are not useful to think about separately at the regional level. So, typically at the individual level, implicit and explicit bias correlate at like 0.2, 0.3, and this low correlation has often been interpreted as evidence that these are like separate constructs, or values, or ideas that exist in our mind. For instance, at the state level it's dramatically different. So, regional aggregates are correlating at like 0.8 or 0.9, and then that goes down the smaller the regional unit is. So, at the core base statistical area that we mentioned before, it's like a 0.6 or 0.7, but you're totally right. Almost everything that I've studied is equally predicted by implicit and explicit, and I just think of them as like two different ways of measuring bias, and that's really what I care about.

We haven't found consistently that implicit bias will be predicting one type of regional outcome, whereas explicit bias will be predicting another type of regional outcome. So, for the purposes of the regional work, I just think they're like both useful information that should be incorporated simultaneously.

And kind of going back to the Bias of Crowds model, one critique I do have is that a lot of the patterns that Keith Payne and colleagues pointed out about implicit bias at the regional level are also the case for explicit bias at the regional level, and yet the phenomena that he describes at the implicit level aren't there for explicit in the first place. We have high test-retest reliability with explicit bias. It doesn't seem to vary as much over time. So, I don't think we yet know the whole

story about how the environment is shaping our biases, and maybe it's shaping our explicit biases more than we might think, as well.

Andy Luttrell:

Yeah. I would have maybe thought, without ever looking at what the results are, that there'd be an argument that it would be implicit bias that's really this whole thing, because people have talked about what implicit bias really is, or these unconscious associations that get built up... I mean, people explicitly say from the culture that you're steeped in, right? And so, it would have seemed like, "Oh, well, this is the kind of thing." Whereas explicit, I go, "Well, if I'm answering the survey question, I have a lot more say in what my answer is. And I can sort of deviate and sort of add new thoughts to that question." And so, it is interesting that kind of both of those ways of getting at bias seem to be doing kind of the same thing, like you said, at the regional level.

Eric Hehman:

Yeah, me too. And I was definitely initially surprised by that. I think that's one of the bigger findings that was very surprising to us has provided a puzzle that we've definitely thought and talked about a ton since we first realized that. There's work coming out right now where we show that's basically the case for any type of bias that we look at. It doesn't seem to be constrained to like Black-white bias, for instance.

And one possible explanation is related to this idea of culture, so at the individual level there's all this measurement error associated with both implicit and explicit bias. This is getting very slightly statistical, but a bunch of noise in your measure provides a ceiling to how high two variables can correlate. And when you aggregate a million people, all that random noise gets canceled out, so in a way you're raising that ceiling to potentially reveal a stronger correlation in the first place. So, that would be like another explanation that maybe the relationship at the individual level truly is higher, but we are so bad at measuring bias that we are not able to find that. But I'm not sure I believe that either. There's a lot of great work showing that implicit bias seems to be qualitatively different than explicit bias.

So, I would say that the jury is still out.

Andy Luttrell:

So, that's one big open question. Are there any other? Just to kind of wrap up, sort of big questions that are next in terms of thinking about this. What don't we yet know? And you also mentioned mountains and rivers earlier. I'm curious what you have to say about... I know you have some data looking at sort of just like how far can we push this idea that aspects of a region would be related to the stuff that floats around in people's heads?

Eric Hehman:

Right, so we did one very exploratory paper and what this paper did basically is build a giant pot of variables, like we had about... I forget. Maybe it's between 800 and 1,000 variables, and these were all characteristics of regions that we clumped into four categories. One which is just like stuff that happens in the environment, and that might be drunk driver related fatalities. Some of them are like aspects of the environment. This would be like the mountains, how many trees are around, how much it's raining, what's the elevation. Then we had mostly from the census a whole bunch

of population stuff, and this really drilled down, so this might be like the percentage of Hmong people who speak English okay.

And then we had... What was the final? Oh, and then just like things in population that were unrelated to population exactly, so how many dentists are in the region? What's the percentage of healthcare workers? And we... Essentially, the goal... What this statistical technique does is you throw everything into the pot, and it tells you like which variables you should keep, and that means that these are the things that are pushing around biases. And so, we did this on a whole bunch of different types of biases, like Black and white biases, gay-straight biases, but moving into a domain that I've done less work in, such as anti-obese biases or anti-atheist biases, anti-Jewish biases, and our goal was to find out do we find any consistent predictors across all of these biases?

And if we do, it's unlikely that it's just chance that that variable was uncovered. And you should definitely with this approach expect some variables to just be randomly associated. But to the extent that it's showing up consistently across all these areas, that becomes increasingly unlikely.

The downside afterwards is like figuring out what the heck is going on with your results, so some things we found were very consistent. So, one major cluster of variables that predicted more bias was just like is my life bad and hard? And the things that were in this cluster were workplace-related injuries were higher, income was lower, there was more drunk driving fatalities, there's like more violence in the area, so it seemed that people in those areas that are maybe under various types of threat had more biases in general across multiple types of biases. That's consistent with some literature that we know already that threat causes biases against different others.

The one that was like unexpected and continues to be unexpected was a major cluster in the opposite direction. So, is my life bad predicted more bias. And it was more of like the percentage or the concentration of healthcare providers in general had a consistent and negative effect on biases. So, the more that they are healthcare workers in a variety of domains, this wasn't one variable. This was like six or seven variables. People's biases were lower. And this is not explained by things like population density, or distance from the coast. This is very specifically healthcare providers. Except I will point out, because it resonates with my personal hatred, dentists were consistently and positively associated with more bias, and I don't know why. I'm just joking about dentists. I'm sure dentists are wonderful people. But it was consistently associated with more bias everywhere there were more dentists.

Andy Luttrell:

Listeners can't see, but you put up quotation fingers when you said, "good people." No, no.

Eric Hehman:

Caught.

Andy Luttrell:

So, we're kind of pushing the boundaries of where we can detect regional level variation in bias. I'm curious what's on the horizon. What sorts of things are you looking into now?

Eric Hehman:

I'm very interested in some of the questions we talked about before, so we have great evidence at this point that different areas vary in their biases. I think a big question is why. That definitely has not been resolved yet. So, that's like a big focus in the lab, is like why is this area different than that area? What's causing the rise in bias in the first place?

Another focus is on similar or maybe inspired by this same sex marriage paper, is I really want to get at causal effects, so we're starting to move into more longitudinal things that with this sort of data allow you to make causal conclusions that are difficult otherwise. Because a big advantage of this approach is that you can study stuff that you can't in the lab, you have to turn to these databases. So, you can't study people getting killed in the lab. You can't study people dying from cardiac arrest in the lab. And I think that's like a strength of this work, but you also need like the right data to examine this over time. So, those are kind of like the big, ongoing projects in the lab right now.

Andy Luttrell:

Well, great. Well, thanks for taking the time to talk about the work you're doing, and we'll keep an eye out for all that new stuff.

Eric Hehman:

All right. Thanks for having me.

Andy Luttrell:

Alright, that'll do it for another episode of Opinion Science. Thanks so much to Eric Hehman for talking about the geography of bias. As always, check out the show notes for a link to his website and the particular things we talked about. You'll also find a link to a transcript of this episode.

You can learn more about the podcast at OpinionSciencePodcast.com or following @OpinionSciPod on Facebook or Twitter. If you're liking what we're doing, leaving a review on Apple Podcasts is the best way to show your support. Thanks!

And just a reminder—my new audio course on The Science of Persuasion is available on the new app, *Knowable*. It's like Spotify for learning. You can find a link in the show notes or go straight to Knowable.fyi.

Ok, that's all. Happy Thanksgiving. Stay safe. And see you soon for more Opinion Science. Buh bye!