Opinion Science Podcast

Hosted by Andy Luttrell

When a Society Changes Its Mind with Tessa Charlesworth

February 27th, 2023

Web: http://opinionsciencepodcast.com/
Twitter: @OpinionSciPod
Facebook: OpinionSciPod

Andy Luttrell:
I've been poking around my family history lately. It started—like all good things—when I Googled my name to see what came up. That's how I found this book—"Luttrells in America" 1671-1900. Obviously I ordered it, and it's this huge book self-published by someone who had been busy trying to trace her family genealogy back to their Irish roots. It was a couple Luttrells from Ireland who ended up in the Colony of Virginia that kicked off her research, but before these immigrants to the colonies, we Luttrells go back to Sir Geoffrey Luttrell in England. We had like a castle and everything.

You know, I say "we," but I actually can't find my own family in this giant book of Luttrells. I have names of great great grandparents that my dad's picked up from his family, and having caught the genealogy bug, I did some internet sleuthing and was able to trace us back one generation further than we already knew. But then it's a dead end. And it's pretty much Indiana farmers all the way back. Although if you leave the Luttrell line, eventually you meet a distant relative in Germany in the 1600s who apparently was beheaded after being accused of witchcraft. Sorry to hear that, grandma.

In any case, the process of digging into the past puts you in a funny frame of mind, imagining the lives of everyday people long gone. Looking at Indiana census records from 1850, seeing the name of your great great great great grandparents and their neighbors, you wonder, "what did they think of the world?" Like sure, I wonder what their day-to-day life was like and how they got by without a refrigerator. But I'm almost more interested in how they felt about things. What did they disagree with their neighbors about? What was the whispered gossip on Sunday afternoons?

Historians have plenty of tools for unearthing facts of the past. Who signed what document? What objects did a group of travelers leave in their wake? But the minds of our ancestors are more ephemeral. Is there any way to go back in time to see how beliefs and attitudes change in the grand scheme? Well, yeah. Social scientists have been developing new ways to peer into the past to get a better sense of how a society changes its mind over time.

You're listening to Opinion Science, the show about our opinions, where they come from, and how they change. I'm Andy Luttrell. This week, our tour guide aboard our social time machine is Tessa Charlesworth. She's currently a postdoctoral research fellow at Harvard University, and she

uses data that have just been sitting around, holding untold treasures, waiting for someone curious enough to find them. Records of implicit bias scores, long-out-of-print books. Data that tell us how people think—and thought—about other people—their prejudices and stereotypes. Tessa shares what she's found in these data and what it means.

Andy Luttrell:
So, a lot of times, and in the work that I've done even on this, when we talk about change in people's opinions or reactions, we're often talking about one person changing their opinion on this topic over time, right? So, today I think this. Next week, I might think something different. But that's not quite the level at which you've been looking at change. So, just to kind of get things started, could you talk about what level at which you've been looking at change and why that might be a valuable way of thinking about change in these kinds of opinions?

Tessa Charlesworth:
Absolutely. Yeah. And I, like you, came from the experimental psychology tradition of thinking about change as something we induce in the lab, and we can measure across two time points, but there's so much more of how individuals are nested in our society, right? We're also nested in what we would call sort of our mezzo level, so our local communities, like our parents, our teachers, our friends, as well as in our macro level, and our macro level includes the media we consume, the books that we read, the kinds of billboards we see on our drives into school. That macro level is really often thought about as the most pernicious and persistent part of our society, right?

So, like when you think about a billboard on a drive into school, it's gonna take sometimes years, decades for that billboard to change, and so it's really interesting to ask how do the kinds of attitudes and beliefs that we might hold in a similar way at an individual level, or even amongst our local community, how do they shift at this broadest level of our society? And if we can see shifts at that broad macro level of our society, then we might be making real traction in actual individual level attitude changes. While of course it's kind of an empirical question that I'm excited to continue to pursue of like what is the interaction between these levels, but so far within the psychology tradition, really there's only ever been a focus on the macro or the mezzo and only now that we have these kinds of big data methods are we starting to also think about the role of the macro in producing change.

Andy Luttrell:
Yeah. We were talking a little bit before about the difference in traditions, so like my impression is psychologists are generally like, "Our whole deal is we say we're interested in the individual and how they're shaped by their surroundings." And so, not to out you, but are you a psychologist still? Or why are you still a psychologist if this is the perspective you take?

Tessa Charlesworth:
Yeah. I think I'm still a psychologist but I'm also more just like a social scientist, I guess I would say, because the things that I'm interested in at the macro level are the bread and butter of psychologists. So, I'm interested in representations, and attitudes, and beliefs, and obviously they take a slightly different form when we're thinking about them at this broadest level of our society, but there's still, again, that kind of basic format of there is a group. For instance, men-women, and

there are attributes. Competent, incompetent. Warm, cold. And I'm interested in the relationships between those concepts.

And that's really a psychological representation, a psychological phenomenon. It's just that it's now scaled up in a way that we don't usually think about.

Andy Luttrell:
It seems like you're riding the wave or were part of creating the wave of like… This is a new direction within the field of social psychologists, taking this more macro level perspective. Probably because it's afforded by the data we have access to, and so maybe that's a reasonable place to transition, so what actually are you doing with macro level change? One of the things that you've done is to look at, and we can define this, but implicit reactions to different social groups and how those have changed over time.

So, maybe just to set a foundation for that, could you talk about what… At the broadest level possible and without dipping into too many controversies, what the distinction between implicit and explicit stuff is, and how you're able to track those implicit things over time.

Tessa Charlesworth: Yeah. At the very broadest level, again, not dipping any toe or even looking over at the mountain of far controversies, I often think about explicit in terms of the way we measure it. So, explicit is something that's self-reported, something that's consciously accessible and controlled, so if I ask you how much do you like Russians, for instance, you can tell me on a survey how much you like Russians. We typically think about that as your explicit or endorsed attitude. In contrast, there are other sort of ways that we might measure your mental representations, and those are slightly more implicit, slightly more indirect, and so that would be through something, for instance, called the implicit association test, which is essentially a method using reaction time or speed for you to make paired associations. Typically, on a computer.

So, you can think of how much easier it would be for you to sort concepts like insect and bad, and flowers and good, than it would for you to make the opposite associations of flowers and bad, and insects and good. Heaven forbid. And that implicit association test is actually something that, again, is sort of the basic way that we've most often measured these more automatic or less consciously controllable alternative representations of the attitudes that we hold.

Now, what's really cool about those two methods is we actually now have really granular data across 15 years at this point, collected through this massive website called Project Implicit online, and back in 2002 they basically put up all of these implicit association tests, or IATs, online for 14 different social group topics to track how these people in our society, millions and millions of people in the U.S. society, are changing in terms of how quickly they associate concepts like good and straight, and bad and gay, or good and young, and bad and old, for instance. And so, that data set, although it hadn't been used when I sort of came onto the scene in my PhD, hadn't been used to track change over time, because of its longevity and because of the scope of people that have come, has been a really useful way of starting to get a sense of how do these average trends in society shift over time?

Andy Luttrell:
So, was that always part of the reason to put it online, was to track change?

Tessa Charlesworth:
It's a good question. So, my advisor, Mahzarin Banaji, was one of the people who put it up originally, and she now has the memory that that was always the intention, but at the same time you have to wonder. You know, I think honestly they originally put it up for more of like an educational demonstration website. It's still called a demonstration website. It's only now that we're using it for research purposes. And so, I think they had no idea how it was gonna blow up and just become this massive source of information, but also data.

Andy Luttrell:
I love the story. It was originally at Yale, right? And the servers crashed within a few days of putting this thing out there.

Tessa Charlesworth:
Yes.

Andy Luttrell:
And so, that also suggests that they had no concept that anyone was actually in large numbers gonna come to this. But the numbers are huge. What do these numbers look like? How many people do you have data from over time?

Tessa Charlesworth:
Well, from just the six tests that I look at most commonly, so like age, race, sexual orientation, we had 20 million last year. Not everyone completes the test and not everyone gets a data point in our final papers, but 20 million at least started and made it halfway through. However, there are 14 tests, so by my sort of back of the envelope calculations, we're probably at least at 30 million people, and it's exploding in the sense of it's just like an exponential growth in interest. Especially following kind of like large scale social events that increase understanding or increase interest in understanding about implicit bias or attitudes.

So, we saw a huge spike, there were 100,000 people who took the test immediately following George Floyd's murder.

Andy Luttrell:
Wow.

Tessa Charlesworth:
So, it's clearly something that people are interested in, and we just have this kind of massive appetite for understanding.

Andy Luttrell:
So, one of the things that that means, though, is that the people who are taking the test also might have changed over time. And so, this is one of the things I know that you've sort of… There are ways to get at this, but it is a curious question of if… Well, let me pause on that first and let's talk

about what changes you've actually seen. So, essentially what you have access to are like a number for each person who's taken this test, and that number represents how easily they associate some social group with good versus a different, opposing, let's call it social group as bad. And so, you just have a bunch of those numbers, millions of those numbers over different amounts of time. What are you generally finding? That they are changing? That they're not changing? They're changing a little bit? Just kind of paint a picture of is this a thing that changes.

Because to set the stage, most people would say these are not the things that change, right? These are our gut, innate, unconscious reactions to the world around us, and once we have formed them through our childhood experiences and what we've learned, those define us forever. So, the fact that I'm setting up that strong case probably gives listeners a sense of where this is going, but what do you actually find?

Tessa Charlesworth:
Yes. That was definitely the strong case that we came onto the landscape with. It was one of the few places where I feel like I've actually changed my advisor's mind, which has been great. Yeah, so we definitely came in with the assumption that everything was gonna be stable. That was true for some biases, so we did find stability for implicit age, disability, and more or less body weight bias, as well, so those trends over time, basically when we aggregate across those millions of data points in each month since 2007, all the way now through to the end of 2020, we find that those trends for age, disability, and body weight are pretty much flat.

When we use sort of like economic forecasting models of trying to predict how those trends might shift in the future, and how long it would take for them to touch neutrality, turns out it would take about 200 years for age bias to ever touch neutrality. So, we're very solidly flat there.

Andy Luttrell:
Meaning people generally automatically associate elderly with bad, right? Is the valence one, right? And that just doesn't seem to be budging.

Tessa Charlesworth:
That doesn't seem to be budging at all on our IATs or on these implicit measures of attitudes. Interestingly, and I'll get to the other attitudes in just a second, but I think it's worth pausing to remember that contrast between implicit and explicit where we actually are finding shifts on explicit biases for all three of those topics. So, in general for all three, age, disability, and body weight, there's been a drop in sort of anti-elderly, or anti-disabled, or anti-fat attitude, so people are much less willing on average to say that they don't like those groups anymore explicitly. No change implicitly, though.

However, there is a different story when it comes to implicit race, sexuality, and skin tone attitudes. So, the implicit race and skin tone attitudes have changed in a very similar way, which lends confidence to us in the results, because they should. They're toward similar topics. And over 14 years, so from 2007 to 2020, they dropped by about 25 to 26%. So, it's not a huge change. We're on average moving about a quarter. But it is a meaningful change when we compare it to the kind of stability that we saw with age and disability, for instance.

Now, by far the most sort of for us shocking change, but also in some sense it's expected from previous work, is the change that we see in implicit sexuality attitudes, which have dropped by 78%. I should edit that. 78% explicitly, 64% implicitly. So, implicit sexuality attitudes have dropped by about 64%. So, again, just a massive, unprecedented change coming from this tradition where we were just expecting six flat lines across all biases.

Andy Luttrell:
Do you have a way to say whether that is a constant downturn since 2007 or a more rapid downturn after a period of stability?

Tessa Charlesworth:
It depends on the topic, so for most of those changing attitudes, so implicit race, skin tone, and sexuality, we actually do see a fair amount of stability, sort of a slow slope downwards until about 2012, and then it starts to speed up. Now, again, what's interesting about this kind of large scale aggregate macro data is it has tons of benefits, right? We can look at psychology at scale as it happens in the real world. We can bring in this new perspective, this new level of analysis. It also has a lot of drawbacks. It doesn't give us that same kind of experimental control that we want and that we get from measuring the same person over time, so why it happened in 2012 that we had a kind of like inflection point of greater increase is… We can speculate on, we can do some quasi-experimental tests, but ultimately it's more open to armchair conversations than to any data that we have.

Andy Luttrell:
I mean, is there speculation about what's driving these changes? Because it's not just like, "Oh, everyone decided to change their mind today." That's not a super satisfying explanation. Because yeah, that's what you're faced with, like all you can say is they're changing. But we want to know what's going on.

Tessa Charlesworth:
We want to know why. Absolutely.

Andy Luttrell:
And why are these changing and those are not changing? So, what's sort of the state of the thinking in terms of why this might be happening?

Tessa Charlesworth:
Yeah. There are so many hypotheses, so many speculations. The biggest one, and I think the one that we now have the most evidence for, is the broad contrast between race, sexuality, and skin tone attitudes, and why those are changing versus age, disability, and body weight attitudes and why those are not changing. And the key dividing line basically seems to be something that I talk about as social prioritization or the degree to which we have made it a priority as a society, as legislators, as policymakers, to tackle and intervene on those biases. So, think about the same sex marriage movement, George Floyd, Black Lives Matter. Those are all directed towards race and sexual orientation. We don't have any of those same kinds of changes and landmark legislation, increases in media rhetoric, really not much change at all in media representation for age, disability, and certainly not for body weight.

So, that sort of difference in frequency or difference in attention seems to be a really key dividing line between those two factors. Then, when we look just within race and sexuality, there are different distinctions, because sexuality is changing so much more than race. And so, we can then look above and beyond just this broad distinction to a little bit more granularity between these social groups, and one of those dividing lines there is that sexual orientation is concealable in the sense of you can meet someone, not know their sexual orientation, become a really close friend, really like them, go for dinner with them, and then at that dinner they reveal that they're gay. You have to do something that we, as psychologists, again love talking about. You have cognitive dissonance. You have to resolve the inconsistency between your belief about their group, not liking gay people, and your belief about this particular person, really loving this new friend that you've made.

And that can only happen, that kind of resolution can only happen for sexual orientation. There's no way to form a close relationship with someone who's Black and not know their social identity before you form that close relationship. And so, that kind of, again, psychological process of cognitive dissonance, and reconciliation in favor of moving your attitudes about the group rather than moving your attitudes about your friend, might be one of the… Again, a key dividing line between the more slow changing race and skin tone attitudes and the much more fast changing sexual orientation attitudes.

Andy Luttrell:
There are groups also, though, for which the… What do they call it? The arrow of time? There's a saying. The arch of-

Tessa Charlesworth:
The arc of time. Yes.

Andy Luttrell:
Yeah. The arc of time bends toward progress. But it would seem that that's not necessarily the case in the window of time that you've looked at. So, just to fill out the story, what are other groups that you've looked at that tell a slightly different story about change?

Tessa Charlesworth:
In the sense of those six?

Andy Luttrell:
I had the impression that some you find that the biases have strengthened over time.

Tessa Charlesworth:
Yeah, so body weight is an interesting case. In recent years, so from about 2010 to 2020, they've been flat, which… We can talk about why. We have, again, a whole bunch of speculations and some evidence that support those speculations, as well. But before 2010, they were one of the few implicit biases that have been increasing, and from 2004 to 2010, they increased by 44%. So, again, just compared to this idea of stability, this idea that our society doesn't influence these implicit attitudes, we're seeing that we can have change towards neutrality for sexual orientation, we can

have increases in some periods of time. So, for about six years, 2004 to 2010, increases in implicit body weight attitudes, or we can have the stability in age and disability, for instance.

But again, body weight is really interesting because those early increases haven't persisted. So, it's also an interesting case, getting back to one of your earlier questions of is it always kind of a linear arc towards progress or arc away from progress. No. Turns out that even within a period of 14, 15 years, we can actually see groups change directions, level off on their trends. We can capture the kind of dynamics that happen between patterns of change. And one of the reasons we think that there was this kind of leveling off for body weight bias is that those early increases might have been driven by the really quickly increasing rhetoric about the body weight or obesity epidemic and these kinds of negative representations that were happening around health, and weight, and all of those problems, as well as the sharpest increases that we've seen in just the representation more broadly of overweight people in America.

So, from 2004 to 2010, almost those exact years, we had the sharpest increases in rises in obesity. Those trends have leveled off as has the kind of inflammatory concern about the obesity epidemic. So, it could be, again, open empirical questions that we need to test quasi-experimentally, but it could be that some of that leveling off is just because we've kind of stabilized as a society and aren't as fearful, if you will, about this rise or this wave of obesity.

Andy Luttrell:
One of the things that this is making me think is the other value of the perspective you've taken is you can look to history, like these data have just been sitting there. I remember when I talked to Adam Mastroianni for this, and he's done this work on shifts in public opinion over time and how we're wrong about thinking about how much they've changed. His point was just like how incredible is it that we can go back in time? Because these data exist, thank goodness that we can go back and check ourselves and go, "Oh, it wasn't always this way," or, "We weren't always as forward thinking as we are now," or things just like… Yeah, it's just an interesting thing where in no other way could we go, "What was going on 10 years ago?" We go, "It's just gone. That time is gone."

Tessa Charlesworth:
Yeah.

Andy Luttrell:
But the fact that a couple people put some tests on a website means that now, oh, we can sort of at least open up a story about changes in these sorts of things. Okay, so let's circle back to the question of is it potentially challenging to this story that presumably more indifferent people are taking this test, right? So, the strong version of the criticism would be sure, it looks like society has changed, but all you can really say is that the people who come to this website and take the test are evincing different reactions to these groups. And so, maybe it's just the people who come to the test in the last 10 years are just different people with different values than the people who were taking it in the first 10 years. How legitimate a concern is that? How have you tried to deal with that? What's the sort of reaction against that strong criticism?

Tessa Charlesworth:
Thankfully, as I said, I'm a social scientist, so I get to draw on all sorts of methods, including sociologists, political scientists, and economists who have been thinking about these kinds of problems for decades, and they have really great methods, including a method called raking and weighting where you can essentially rake through the data and find similar kinds of people across all your time points, and then reweight your data so that you basically have the same kind of synthetic sample that make up your final average at time one and at time 168, basically.

And so, we use that approach here even, because we're looking at those averages, to essentially reweight the data so that it's always the same proportions of liberals, and Republicans, and men, and women, and gay, and straight, and young, and old, and all those things. So, all the numbers that I've mentioned so far control for those kinds of population demographic changes in our sample.

What is good, however, is when we actually just look at the raw amount of change in our sample, there's much less than we were concerned about. The biggest change is that the sample has become slightly younger. A large part of that is probably because the website is being used more and more in classrooms and in colleges rather than in just sort of word of mouth and businesses. But again, we can control for that, so it's not a huge concern when it comes to our average trend.

I will say one more thing on this point, though, because I think it's an important qualification of some of the average trends that we talked about, which is that these kinds of sociological models also have a really great way of separating out different sources of that change. So, in addition to just controlling for it, you can actually then try and say, "Okay, well, if we didn't control for it, how much of it is coming from this demographic change? How much of it is coming from a true average trend change?" And these kinds of models, one of the… The biggest one is called an age period cohort model, where you essentially try and partial out how much of the change that you're observing is due to your sample getting older? That's the age effect. How much of it is due to a period effect, which is what we kind of want to be capturing, kind of this general movement of all of society, which would suggest that everyone's kind of changing their minds. And how much of it is due to cohort replacement? So, the baby boomers getting replaced by the millennials, getting replaced by the Gen Zers.

For sexual orientation attitudes, that big drop, 60%-plus drop, that is mostly due to a period effect. Pretty much every cross section of demographics that we've looked at is changing towards neutrality. Liberals, conservatives, old, young, men, women, and at similar rates. For race, however, race and skin tone, it's a little bit more of an interaction between cohort replacement and period effects, and that shows up as faster change amongst younger cohorts. So, although all cohorts are changing, even baby boomers have changed their minds a little bit, millennials are changing much faster. And I think that's an important thing to remember, especially when it comes to trying to identify the source or identify the why of this change, because it tells us that whatever the source of that change is, it's most likely affecting millennials more than it's affecting baby boomers.

And that brings us back to this kind of conversation of it's probably things like protests, changes in media representation, changes in the rhetoric through social media, because those are just more likely to be invested in and participated in amongst younger people than amongst older people.

Andy Luttrell:
What was the last? Oh, cohort replacement. Is that not as much… I mean, it's maybe a short enough time span that you can't see that as easily, but is that… Can you parse the data that way, too?

Tessa Charlesworth:
Yeah. Yeah. So, definitely… So, there's a little bit of cohort replacement going on with race in that basically the cohort by period interaction essentially means it's a little bit of both. The average trend is partly because baby boomers are being replaced with millennials who at baseline just have lower attitudes. But there's also a period effect which shows that even baby boomers are changing, too. So, it's not just that we have this stable group, and this stable group down here, and the demographics are changing, but rather that there is a little bit of change of both. It's just that the slope for the millennials is faster than the slope for the baby boomers.

Andy Luttrell:
Yeah, so my concern then, which is now resolved-

Tessa Charlesworth:
Good.

Andy Luttrell:
… is like ideally you would want to be able to look at this kind of macro level change and draw some inferences about individual level change, right? And so, I say ideally, I'm speaking for myself. And it sounds like based on all of that is hopefully giving us that conclusion that it does seem like these macro level changes are the product of a bunch of individual level changes, and not simply the fact that the people of 2020's Project Implicit respondents are just different people than those who were there in 2007, but actually even though you probably have not nearly the number of the same person twice, or even you wouldn't even know that. We could probably infer that this is tapping into people individually having different reactions now than they did before.

Tessa Charlesworth:
Absolutely. And I think the wealth of evidence, because this is not the only paper that's studying these kinds of changes in public opinion, and the few sort of cross-panel surveys that have done this for race attitudes, for instance, they find very similar rates of change that we do. And that tells us we're not just some weird, unique data of Project Implicit that's this big, massive convenient sample. It's really we're mapping onto similar kinds of changes in these representative panel surveys, as well, telling us that probably what we're seeing is at least partly, especially for race, at least partly due to individual level change, although partly due to cohort replacement, but for sexual orientation I think we can safely say lots of individual people are changing their mind.

Andy Luttrell:
Nice. So, luckily also it's not simply… We don't know this area only through the people who visited this website or the handful of super expensive panel studies. There's another way to get at

society level change, which is to look at these cultural artifacts over time. So, could you talk a little bit about where do you look for clues beyond a website that's been asking people to play an implicit association game for the last 20 years? How else can we find evidence and go into that social time machine to get a sense of these sorts of changes over time?

Tessa Charlesworth:
If you think about one of the hallmarks of what makes us human and what makes us social humans, one of the key things is language and how we talk to each other, communicate to each other, write to each other, and the reason why that's often touted as sort of a social hallmark of humans is because we've been doing it for a very, very long time. You mentioned Adam Mastroianni and his remark of how great it is that we can look back in time. He was referring to the kind of like social surveys that we have that go back until the mid-1900s, like how cool, we can look at the 1950s and say, "Look at how much we've changed since the 1950s about gender rights."

Turns out with language we can go back 200 years or more, so we can actually look back to the 1800s and say, "Now look at how much we've changed in terms of the way we represent or think about men and women or old and young." And so, what we've been doing is… This has been done in the digital humanities, and by historians looking at archives, and papyrus scripts, and all those things for decades. That's their bread and butter. To quantify it, though, at scale, and with the kind of flexibility that we might need to look at multiple groups and compare across targets, that really requires really intensive computational power.

And so, what's been exciting is in the past basically five to 10 years or so… Yeah, 2014 was almost 10 years ago. That's wild to think. There've been this new class of models in the area of natural language processing in computer science that use word embeddings, which essentially take all of that massive text, like billions of words from the language of the internet, or these historical books, and then convert those words into numeric representations. So, each word get a long string of numbers, about 300 numbers, that positions it in this high dimensional space of meaning, so that words that are closer together in meaning, like bread and butter, which I've said a lot today, bread and butter will be closer together in that space than say bread and cactus, for instance.

And once we do that, it means that we now have a way of quantifying all of the word meaning in 1800, in 1810, 1820, and so on and so forth, and then comparing it statistically across time. So, we can see how the vector or the string of numbers that represents the meaning of the word old, or the meaning of the group concept old, has shifted in its kind of trait associates from 1800 through to now.

Andy Luttrell:
And so, just to even unpack that a little more, in what books are we talking about? Where are these words living that tell us anything?

Tessa Charlesworth:
Yeah. There are a number of different corpora of books or texts. The biggest one is Google Books, so when you go online and just search Google Books, you'll get all sorts of different texts, the Bible, old historical children's books, those kinds of things, and it's basically a massive dump of all of that language. The Google Books corpus, which is the one that we rely on most, although

we replicate it with some other corpora, is thought to represent about 5% of all books ever published. And 5% seems like kind of a small amount, but when you think that that actually turns out to be about 500 million books, we're pretty good. We're getting a good sample of text over time.

Andy Luttrell:
Does that shift with time, though? My guess is that there are probably more of them in the last 20 years than 200 years ago.

Tessa Charlesworth:
Absolutely. Absolutely. Yeah. And another key point, and this actually relates to some of the things we were talking about in terms of measuring opinion shifts over time, too. It's not only the sample size that's changing, but the composition of the sample is also changing. So, you might have more fiction texts in early years and more non-fiction texts in later years, or also the sample demographics of the authors themselves are changing, so writing and publishing was obviously usually a pursuit of the wealthy in the 1800s, and now, today, it's much more accessible, and open, and public.

So, some of those are things that we just can't control for because of the sheer scope of the data that we're looking at. We don't have meta data on every single book that's in Google Books, so we can't perfectly control like we could with our opinion surveys. But there are some things that we can control, so we replicate all of our analyses in corpora that are perfectly genre balanced. So, same proportion of fiction and non-fiction texts over time to make sure that it's not entirely due to sample change.

Andy Luttrell:
So, you're raking and weighting on genre rather than demographics of people.

Tessa Charlesworth:
Exactly. Exactly.

Andy Luttrell:
I take it that the words bread and butter are not exactly what you're looking for in these books. Maybe they are. Maybe that's the paper that you're working on that I don't know about. But when you are looking into these hundreds of millions of books, what are the kinds of words you're looking for? And what would it mean for them to change in meaning over time?

Tessa Charlesworth:
I'm particularly interested in social groups and how social groups are changing in their association with different kinds of trait adjectives. So, good, bad, warm, cold, competent, incompetent. The way that we look at change, though, is not only by requiring us to look at the kind of key dimensions of meaning of stereotypes like warmth and competence, but we actually allow it because we have the flexibility of language to go entirely bottom-up. So, we can literally just say, "Here's a vector. It represents the meaning of the group old. It's using a whole bunch of synonyms for group old, like old, and elderly, and grandmother, and grandfather, and so on." Then we ask, "Which words, of all possible trait words, are most associated with that group in 1800?" And then

we do the same thing in 1810, in 1820, and 1830, and so on. And basically to see change, we can look at it both qualitatively in the sense of looking at what those top trait associates are just in terms of their meaning, and the general gist that we get out of that qualitative content. And we can also look at it quantitatively, so we can say, "What is the valence of these traits? What is the positivity or negativity of these traits in 1800 versus 2000," for instance.

Andy Luttrell:
And so, have those changed?

Tessa Charlesworth:
For some and not for others. So, it's a very similar story to the one I was telling before, and it reminds us of the importance of looking across multiple groups, because for some groups, like gender for instance, we actually see very little change. So, women in the 1800s were represented in very similar ways to women in the 1990s. And you know, that includes top trait associates like feminine, charming, warm, caring, those kinds of things. There's a little bit of a qualitative shift that we haven't interpreted too much, whereas in the 1990s there's a little bit more competence-like adjectives, so they start to get words that are slightly more negative too, like artificial, and abusive, and those kinds of things, which might reflect… I don't know. The 1990s idea of women are starting to gain a little bit more power in society, so they need to be pushed down in-

Andy Luttrell:
How dare they?

Tessa Charlesworth:
How dare they? Exactly. But in general, gender is very stable. Age is also quite stable. There are some shifts. Again, interesting qualitative shifts that align with what we know about these groups and their histories. So, actually the representation of young is a really interesting case study, where in the 1800s young was mostly associated with physical vigor, and positivity, and activity, and all of those things. Around the 1970s to 1980s was the introducing of this idea of the adolescent life stage, or in other words, the stereotype of the teenager. And around the 1980s and 1990s we start to see these new words, like antisocial, aggressive, impulsive. They start to come in alongside still holding onto some of that activity and vigor stereotype. So, we get this new addition of a slightly more negative stereotype, as well. But in general, these kinds of groups, age, gender, even social class have remained quite stable over time.

The groups that have changed much more in their content on average as well as in their valence are our kind of like ethnicity and racial groups. So, Black, White, Irish, Hispanic, Native American have generally changed quite a bit more. And what is interesting for us from a sort of long standing social psychology perspective, social psychology theory, is that those are the groups that should be changing more because they're thought to be a little bit more culturally constructed or a little bit more "arbitrary." So, as Jim Sidanius would say, these are groups that we didn't have in our long evolutionary past, so they're much more likely to be malleable based on what we need, what function we need them to serve in our society today. And we see that kind of shift and turnover happen across time.

Andy Luttrell:
We're constantly trying to figure out our relationship with these kinds of categories, whereas in the language of social dominance, right? The meaning of age and the meaning of gender ought to be fixed based on whatever the reason is.

Tessa Charlesworth:
Exactly.

Andy Luttrell:
Whereas these other ones are just like, "Oh, these are meaningful only by virtue of the fact that we've given them meaning," and so we're constantly trying to figure out what that meaning is. And what's cool about that, so again, you have specifically stereotype associations, but you also have the valence of them, right? Which you could ask, is the valence stable even if the specific traits are swapping in and out? Or is it just that like, "Oh, we're just rolling dice every day and it's just changing without any real underlying goal of maintaining inequities and those sorts of things."

Tessa Charlesworth:
The answer is what you said first. Valence is much more stable. The traits themselves are shifting out over time. So, we see something, for instance, like a top trait associate of Black American in 1800 is lazy. In 1990, it's helpless. It gets counted as a shift in stereotype content. It gets counted as a new trait. But we can clearly see that a little bit of the semantic meaning of that trait is preserved, and also, clearly, its valence, both are still quite negative. So, actually we see much more stability in these kinds of latent dimensions of the meaning of those stereotypes, even if there's more shift in the kind of manifest or surface level content.

And I think to me, as someone who's ultimately interested in understanding change, but also persistence and why it seems to be the case that on many of our papers we're seeing these large scale social shifts in opinions, or in representations of groups, and yet we get a lot of pushback of people saying, "Hang on. There's still so much discrimination happening. There's still so many persistent hierarchies and divisions in society." And I think this additional analysis of really contrasting how you can have more surface level change that can be picked up and is probably meaningful in its own right, right? Being able to shift between traits that we associate with a group is clearly saying something about the processes of change.

But the fact that we have this kind of undercurrent of latent stability or latent differences across groups and how they're ranked according to their positivity and negativity tells us one place that stereotypes and hierarchies may be enduring.

Andy Luttrell:
What I do love about this project is… And a lot of these kinds of studies that have come out, is just how lucky it turned out to be that it was possible, right? So, the reason this exists is because the Google Books people were doing their thing, and for whatever reason logging a bunch of books over time. Also, some other folks were like, "Let's just get adjectives in the world and see whether people think of them as positive or negative." And then someone like you comes along and says, "Wait a minute. These two things could go together to answer a brand new question that is new." And so, just that's very exciting as how science works, but it also makes me curious, like I get why

you might look at Project Implicit data over time, right? As you say, you worked with someone who apparently, since the beginning, expected this project to happen, and you pulled it off. But where did the notion of this come from? Why look to books and the way that these groups are construed in them?

Tessa Charlesworth:
I started… So, I started grad school in 2016, and at the exact same time Aylin Caliskan, who's now a Professor of Computer Science at University of Washington, but at the time was a postdoc at Princeton and computer scientist, published this paper in Science called The Word Embeddings Association Test. And no one from the IAT world, no psychologist who developed the IAT knew that this paper was coming out. She independently had just thought, "Hey, we have these word embeddings. Let's look at the associations between groups and trait adjectives as a computer scientist." And it was reading that paper that immediately, working on IAT and working on IAT change, was like, "Hang on. If someone can create an analog to the IAT but in text, couldn't we do that across time, as well?" And so, it was really actually the first paper that we published with that kind of method. I know longer use the word embedding association test, which we can talk about for other reasons, but the first paper that we used to show the validity of that method more rigorously was looking at differences across children and adults in terms of their gender biases in their natural language.

So, how do children's books and children's TV shows differ from adult books and adult TV shows in terms of the degree of they associate men with science and women with arts, for instance. And so, that was a fun first project in the sense of being able to validate the method, but it also left a very sort of unsatisfactory conclusion about change, because we had corpora that varied across time as well as across age of the respondents, so younger versus older. And it was really in the concluding sentences, as so often we do in these concluding sentences of papers, where we say, "Future work should explore blah, blah, blah, blah, blah." And I said I think something, a kind of throwaway sentence of future work should explore continuously tracking these same groups within the same kinds of texts over time to provide more rigorous quantification of change.

And it was like in rereading that final sentence it's like, "I should probably do that myself." So, that's where the project came from. It was kind of this slow evolution, being inspired by Aylin's work, and how we might do that across longer historical time scales.

Andy Luttrell:
It seems like this is only just beginning, right? To look at text as a way to understand how people think about stuff. Where does this seem to be taking you now?

Tessa Charlesworth:
For me, the main place I'm going with this is actually to pause a little bit and reconsider some of the assumptions that we have about what we're measuring. So, I think we all got really excited when natural language processing got onto the scene, and the expansion of these methods, just so cool, and it also meant that we started sort of overinterpreting what we might be measuring. And one of the key overinterpretations I think that's happened is we often call these measures of our language and our language biases implicit biases because they're measured indirectly. Turns out

we have no quantitative evidence about whether they're more implicit, whether they're more explicit. They obviously have features of both of these kinds of things.

We also don't know the degree to which these kinds of collective representations that are shared through language, and explicitly expressed and edited through language, map on at all to our actual endorsed or opinions that are collected through surveys or through IATs over time. So, really we need to have a better kind of triangulation across methods to figure out how all these pieces fit together. How do attitudes relate to language? How do implicit attitudes and explicit attitudes relate to language? And perhaps even most importantly, how do real world behaviors and real world outcomes relate to both of those things?

So, how do attitudes plus language help us explain why some behaviors are really static and persistent across time? For me, it's a little bit more of this kind of like conceptual weaving that needs to happen to understand and really validate the use of these methods for broader social science.

Andy Luttrell:
Nice. Okay, that makes a lot of sense. In terms of just the notion of change, one thing I wanted to sort of get as a reflection point is just in doing this, you've interrogated this question of change a bunch, in a lot of ways, for most of your time in this world. And I'm curious, what is surprising to you about the changes that you see? What did you expect? How has this sort of changed your relationship with just thinking about how the world, what progress means, what change means? I just want to get a sense of it just seems like this is the kind of work that is sort of… shakes preconceived notions away of like, "Oh, this is what's been happening." So, I just want to get your take. How surprising have these findings been to you and how much have they affected the way you look at the world and the biases and opinions people express?

Tessa Charlesworth:
I definitely think… So, I grew up sort of with this social activist mindset. It was like me, my mom, and my sister, and we were all little gender feminist people and stuff like that, and so part of that social activist mindset almost requires a constant belief that people don't change, and you're always fighting this idea of we need to make things better, and it's not good enough yet, and we've made no progress, and we need to keep pushing. And I think starting to see some of these massive trends and starting to really pay attention to just how much change has happened… You know, we have another review paper of how much change has happened for women in science, for instance. There are, of course, still a lot of persistent discrimination, but there are huge shifts in the number of women coming into secondary education, the number of women in psychology, and biology, and even computer science, and I think I just… I never paid attention to those before. It was almost my own sort of cognitive filtering about evidence that change is something that can't happen.

So, I think as a really general takeaway point, it's been a perfect reminder for me as an individual that change has already happened, and therefore that change continues to be possible. And the key push for us… You know, I still have a little bit of that feminist activist fire inside of me. The key push is now no longer we need to rail against the man, and the top, and all those things, but really we need to learn from those successes of the past, bottle them up, and continue to implement them. Not only for the attitudes that we're most interested in, like sexual orientation and race, but also

transporting those positive lessons of change over to those attitudes that we're seeing that are really stagnant, as well. So, age and disability need some of that special sauce that we've been giving to sexual orientation and to race.

So, again, it's like change is possible, change needs to be spread, always coming at it from this more sort of promotion focus or positivistic perspective.

Andy Luttrell:
That's great. That was a very inspirational way to end, so thank you so much for taking the time to talk about this. This was all super interesting.

Tessa Charlesworth:
Yeah. Super fun. Thank you.

Andy Luttrell:
Alrighty that'll do it for another episode of Opinion Science. Thank you so much to Tessa Charlesworth for sharing her work. You can find more about her and the research we talked about by visiting the links in this episode's webpage.

By the way, Tessa's PhD advisor came up a couple times—Mahzarin Banaji. So I'm obligated to direct you to a favorite episode from the Opinion Science archives: Episode 16 with Mahzarin Banaji herself.

And of course, if you want to dig into all the past episodes or just keep up with new ones, subscribe to Opinion Science wherever you get podcasts. You can also go to OpinionSciencePodcast.com for all past episodes, links to various resources, transcripts of episodes…I think there's a picture of me on there in case you want to shatter this radio illusion. I also tweet out new episodes on Twitter--@OpinionSciPod if that's something you're into.

Alrighty, that's it for now. See you back in a couple weeks for more Opinion Science. Buh bye…