



Opinion Science Podcast

Hosted by Andy Luttrell

Language of Opinion with Matt Rocklage

February 15th, 2021

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

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

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

Andy Luttrell:

I've always been transfixed by online reviews. I go down rabbit holes when I'm comparing products or looking for a new restaurant. I remember once seeing an Amazon review that said they loved the product and recommended it 100% and they gave it four out of five stars. And there's a restaurant in my neighborhood with mostly five star reviews, except one person, who said he called 911 because the owner is very cheap. And the number of times I've seen people say that they would give something zero stars if they could... It's such a trope of negative reviews. We've got to get past that one.

But I think it's no accident that I study opinions and am mesmerized by online reviews, because what are reviews but expressions of opinion? And sure, you could just look at the average star rating in a pinch, but when you read the reviews you realize how varied people's perspectives can be. One five star review might be a long story about every moment this person was inside a restaurant for their birthday, and another five star review of the same restaurant might be a detailed breakdown of the culinary expertise on display. For most of what we say and write the words we use are revealing, but this may be especially true when we're expressing our opinions.

You're listening to Opinion Science, the show about our opinions, where they come from, and how they change. I'm Andy Luttrell, and this week I'm happy to talk to my friend Matt Rocklage. He's an Assistant Professor of Marketing at the University of Massachusetts, Boston. He and I went to grad school together, so I've known him for at least 10 years, which might be an Opinion Science record, actually. In our conversation, we talk a lot about a tool that Matt developed called the Evaluative Lexicon. Just to make sure we're all on the same page, you can think of this as a funky kind of dictionary. In a normal dictionary, you look up a word to discover what it means. But in Matt's dictionary, when you look up the word you get numbers that quantify how much that word is related to positive or negative opinions, and how much that word conveys a sense of emotion in that opinion.

So, like if you were to say to me, "I adore your podcast," I could look up adore in Matt's dictionary and it would tell me that it's a word that conveys an extremely positive emotion-based opinion of my show. And thank you, by the way. But if you were to say, "I thought Toy Story 2 was okay," I could look up the word okay in Matt's dictionary and it might tell me that it's a word that conveys a pretty neutral, emotionless opinion of the movie. So, I talked to Matt about how we use language

to express opinions, how we look to language to learn about other people's opinions, and why emotion-based opinions are so potent.

You know, the starting point for a lot of the recent work you've done is the Evaluative Lexicon origin story, right? That first paper, where you introduced this cool new way to measure attitudes in a way that sort of harkens back to some of the ways that people have done it before, but in other senses of it is totally brand new. So, I wondered if just to start you could describe one, the mouthful that is Evaluative Lexicon, if we give some sense of meaning to that word, and also where the germ of it came from. Where did it start from?

Matt Rocklage:

Yeah. It really started from the realization, I think this was about my third year of grad school, so I was about halfway through grad school, and I needed to pick up, basically out of practical concerns, a new line of research. And you know, it came from the realization that we tend to have way too many words to describe our opinions. Why do we have so much redundancy in our language, especially English I think, for expressing why we like or dislike something? Why say something's amazing, fantastic, perfect, great, good, wow, that's a lot of words. And so, we thought there must be something subtle, or sometimes not so subtle, that these words are really emphasizing or conveying, and being attitudes researchers, my advisor in grad school, Russ Fazio, being attitudes researchers, we thought, "Well, surely this must give some insight into someone's opinion beyond just whether they're positive or negative."

And we wanted to figure out what we could get out of language, and so we thought about, "Okay, positivity, negativity, sure." How extremely positive or negative someone is, so the word good is positive, but great, even more extreme. But then we thought, "Okay, above and beyond that, can we say something about what kind of information someone's opinion is based on?" So, if I say something is helpful, that sounds like I've considered some property about that object, whereas the word amazing or enjoyable, that seems to convey I am having some sort of feeling-based reactions toward that.

And so, yeah, it really just stemmed from the idea that we have so many words, and why do we have these words, and what can we get out of these words to help us understand people's opinions?

Andy Luttrell:

Is there like a biopic moment where you had this insight? It seems like the kind of thing that would be like, "Okay, I've been in grad school for this many years, studying whether people like or dislike things." You did a bunch of work on negativity and how people respond to that. And then sort of I just imagine you're like on Amazon, and you're reading a review, and you're like, "Gosh, these words are so much different than this other review that makes exactly the same conclusion."

Matt Rocklage:

That's not far off, because I was a Yelper back in the day.

Andy Luttrell:

Oh, that's right.

Matt Rocklage:

I was one of those people, so in Columbus, Ohio, and many other larger cities, they have groups of Yelpers who, they write many Yelp reviews, and then they join this group. It's called the Yelp Elite. Not to sound too ridiculous, but it's a little-

Andy Luttrell:

Or elitist.

Matt Rocklage:

Exactly. So, they have these groups in your city, and they have get together where they put on events, and just get people to try out new restaurants, et cetera. And so, I was really deeply involved in writing reviews and so that is a part of where this all came from. In fact, originally we hadn't planned on using the Evaluative Lexicon tool to measure natural language, which now in retrospect, that seems pretty obvious. But we were just gonna give people a list of words and have them select the words that best describe their opinion, but with the just explosion of words online, well, there's a huge untapped dataset, basically, of people naturalistically telling people about their attitudes and I was one of those people on Yelp.

And so, that connection to me became very apparent when we started to actually get into the mechanics of the word. It's like, "Well, wait a minute. Of course, we could give them a checklist, but why don't we just grab data from people naturally expressing their opinions and see what we can predict from that?"

Andy Luttrell:

So, maybe by way of giving a little bit of an undercover look at the Evaluative Lexicon, how might you use it to understand those natural texts, right? So, what does it mean that there is a lexicon, and then what do you do with it once you have it?

Matt Rocklage:

Yep. That's a great question. So, the idea being that each word that we use: perfect, good, great, terrible, hates, et cetera, those all imply some level of positivity and emotion, right? And so, what we can do is create a very large list is what we ultimately did, create a very large list of words that we use to express our opinions, and then we can say, "Okay, each one of these words implies something. Can we quantify what it implies?"

And so, we simply just got a really large set of external participants, individuals, and asked them basically if you were to use this word, the word perfect, how positive would you be on a scale from not at all positive to very positive, or very negative to very positive? And then we asked another set of people, "Okay, if you were to use this word, how emotional would you be in your opinion?" Not at all emotional to very emotional. And once we get a large group of people, we can say, "Okay, now some words have different meanings to other people, different people, but in general, people tend to say amazing is more emotional than the word superior." And we can use the underlying ratings that these participants gave us and say, "Ah, when you use the word amazing, I can look that word up in the Evaluative Lexicon and give that a number. I can give it a number in terms of how positive people believe it is and I can give that a number in terms of how emotional it is."

And then I can just put those numbers in place of the word in texts when I come across those words. So, now perfect has quantities, numbers underlying them, and now amazing has numbers underlying them, and we can average those together and say, “Oh, on average, this person seems to be saying that they’re about this much emotion involved in their opinion and this much positivity.”

And so, that’s the mechanics underlying it.

Andy Luttrell:

And it obviously is saying something, right? Because you can cross-reference it with... The nice thing about those Yelp reviews is people not only are writing out their opinion, but they’re scoring it on the kinds of like bread and butter kind of measures that we usually use, and you can, I think, see that those cohere together pretty well.

Matt Rocklage:

Absolutely. So, star ratings, those are the final ratings that people give their restaurant that they go to, or give the Amazon product that they’re reviewing, and we can use their language to predict that star rating. So, it’s nice that we can quantify the language using that approach and then check it. Well, did we do a reasonable job or are we just totally off by predicting that final star rating. Then we can start to ask interesting questions about like, “Okay, well, this person’s more emotional in their language. What does that mean for their star rating?” Those were some of our first questions.

And then we can go beyond that to say like, “Okay, well, this person’s now expressing a mixed opinion. They’re both positive and negative. Their negative side is quite emotional, whereas their positive side, they’re talking about ooh, this product is helpful, beneficial, but it broke, and it was just awful because of that.” How do they form their final judgment, their final star rating when you have conflict like that? And so, these became kind of a second question, and we proceeded from there.

Andy Luttrell:

Would you go as far as to say that the Lexicon is a better measure of opinion than your kind of standard survey questions? It reminds me just a little bit of like the back and forth with neuroscience like 10 years ago, where it was like, “Wow, look, we can find out a person’s whether they like or dislike something by looking at their brain.” And you go, “Or we could save thousands of dollars and just ask them a question.” Is there any... But some of their arguments are like, “No, but we actually learn more by looking at the brain.” Do you think we actually learn more by looking at words? Or is it just looking from another angle?

Matt Rocklage:

So, putting on my responsible science hat, it’s a tool that we can use to enhance our understanding about some things, but it has its own limitations, right? So, I wouldn’t go as far to say it’s a better measure, but it certainly simplifies the measure of attitude emotionality as we call it. The emotion that people base their opinions on. It’s simplified, in many ways. But because it relies on language, language is complex, right? So, even as simple as asking you to check the words off, select the

words that best describe your opinion, well, in what context are you doing that? Are you gonna communicate that opinion to someone else? Well, now I have self-presentational concerns. And so, I might give you a different opinion than I hold personally. That concern can be lessened based on if you're in a lab and you're not being told that you're gonna express this to anybody. Well, then it doesn't matter. I'm just gonna give you my personal opinion.

But in other contexts, it could really matter, so if I'm a Yelp reviewer, I have to be helpful too, right? I can't just tell you how emotional I am. That gives you some information, but I really want to tell you about this restaurant. I want to tell you something about this product that I am using. And in order to do that, I'm gonna alter my language to go beyond my opinion because I have to tell you something helpful.

So, it has its limitations, but I also think those are interesting opportunities. How does context change your opinion when you express it? And then later, does actually having written about that opinion change the underlying attitude? So, it provides benefits, but opportunities and challenges as well.

Andy Luttrell:

The social desirability thing is interesting. So, one of the problems with surveys, like you said, is that people can just lie. They go, "I know I'm not supposed to like this, or I know I'm not supposed to hate this, so I'll give the answer that you think I want." And I'm curious if you use the words people choose, is there still something that they're selecting that betrays something about their true feelings, right? So, yes, maybe they're lying on the valence dimension. Maybe they're saying something, "Oh no, this is bad." But are they choosing those lies in a biased way?

This is... I mean, we are way off, but it just struck me that... Are there things about words that tell you more under those sorts of contexts?

Matt Rocklage:

Yeah. I'm not sure. I'm not sure how exactly to answer that. So, I'm trying to think... This gets a little bit off, so we can correct course if it's not the way you want to go for this, but one paper that I wrote with Derek Rucker and Loran Nordgren, and this was published in 2018, was how people go about persuading each other and the language they use to do that. So, obviously they have their own opinion, but now they're being asked to persuade someone else to purchase a product or to go to a restaurant, and so they're obviously going to be changing something about that language to do so.

And this idea actually didn't come from, "Oh, how do people go about persuading each other?" The original idea came from, "Hey, if someone writes a recommendation letter, how do we know whether they're being accurate in their representation of the person they're recommending or if they're just trying to persuade you to want to take their student, or want to hire this worker because they want you to have this person, because it'll make them look good or whatnot?" So, the idea being can we see how people shift their language when they're not fully representing their opinion and they're doing so in a motivated way? They're trying to get you to do something.

And so, what we find there is that people have the lay belief that I ought to shift my language toward emotion. Try to get people to think that, “Oh, there is a feeling that was generated based on this worker who I’m asking you to hire, or my student who I want to get hired by this committee.” But then it shifted away from hiring toward just general persuasion and what do people do when they persuade others? So, that kind of gets at your question about how do people shift or what is the underlying attitude versus how I’m gonna present it, but maybe a little bit different than what you were thinking.

Andy Luttrell:

The difference is that there people are probably being consistent in whether it’s good or bad, right? I’m trying to persuade you of what I think. But maybe... Do we know that people who don’t already have a more emotional attitude will still turn to emotion? To go like, “Okay, it’s not emotional for me, but I’m gonna manufacture some emotion to get my point across.”

Matt Rocklage:

Yes. So, we asked people in a couple different ways. In one experiment, we asked everyone to think of a restaurant that you like. Doesn’t need to be your favorite, but just one that you liked. And then you’re gonna tell someone else about that. And then we paid one half the participants, we randomly assigned them to... We told them, “You’ll receive a dollar every time someone selects your restaurant as the one to go to.” So, they’re being incentivized to persuade. And in the other half of the participants, we didn’t give them any sort of added instruction. We said, “Just tell someone about the positive aspects of this restaurant.”

And there we see, okay, with an existing object, the restaurant, people shift their language toward more emotion when they’re incentivized to persuade. We also just randomly assigned people to a product they’d never heard of and we just gave them a short couple sentences, maybe three or four sentences about the product and said, “Okay, now please write a review to someone.” And in one condition we incentivized them to persuade. In the other condition, we didn’t give them any added instructions. We just said, “Just tell about the positive aspects of this product.” So, there they didn’t have any background of that product. These ranged from GPS units, to movies, to books, et cetera, so yeah, even without an existing experience with the object, they still imbue it with this more emotional language.

Andy Luttrell:

The emotion thing is a good pivot point, I think, just because that has been a lot of where you’ve gone with this. And was it always part of the plan in terms of the Lexicon? So, like when I see these words convey positivity or negativity, got it. Of course. How extreme. Sure. Of course. We’ve always measured extremity. And then there’s other column for emotion that to me wasn’t like a, “Obviously that’s part of this enterprise.” And so, I’m curious. Was that part of it from the beginning? Like, “Of course, we need to account for emotion.” Or it was like, “Oh, one of the things that’s moving here is the emotion that’s conveyed.”

Matt Rocklage:

We definitely had that in mind from the beginning and this goes back to the history of attitudes, where you can base your opinion or attitude on many different sources of information. And then three major kind of overarching categories being your behavior, what have you done in the past?

Oh, well, I ate macaroni and cheese prior to this, so I guess I must like it, right? So, I'm using my behavior to understand my opinion. That's one source of information. Another source of information is cognition. As we define it in the attitudes literature, when you base your opinion on cognitive information, it's like, "What are your beliefs about the attributes and properties of that object out there?" So, I'm looking at my laptop right now. I would say, "Well, it's pretty reliable. Long-lasting battery. Yeah, I think it's pretty useful." So, I've come to my opinion based on the properties and attributes of the object. And then finally, the third facet that we talk about is the extent to which you base it on a feeling or emotional reaction, and we define that attitude emotionality as, yeah, the extent to which you base your opinion on your feeling, emotion-based reactions.

And so, yes, it was always part of the Evaluative Lexicon to consider the basis of the opinion. We have a hard time doing that with behavior. Don't really know how we do that with behavior. But we did concentrate on emotion right away with cognition there in the background, kind of just sitting there. We haven't pursued really investigating the cognitive aspects of attitudes. We focus really on dimension of low to high emotion, but there are some times that we talk about the cognition, as well.

Andy Luttrell:

And so, what bonus do we get from knowing how much emotion is in the words that people are choosing to express their opinions? Why does that give us any more than just knowing that you like this?

Matt Rocklage:

Right. Yeah. So, this is one of those times that I say, "I'm glad you asked that, Andy. Let me tell you..." So, for a long time you would talk about cognitive basis versus emotional basis and like which one is stronger, which opinions last longer, these kinds of things. And there's been a real focus on thoughtfulness, right? Work by your advisor, Rich Petty, on need for cognition. So, how much do I enjoy thinking about things? How much do I pursue activities where I can think deeply about something? And so, people high in the need for cognition, or people who think more about their attitudes in general tend to have fairly strong opinions. They last longer. They're more resistant to persuasion, et cetera.

Well, you know, you would think, "Well, if I thought something was amazing versus just superior, that also could have some sort of strength related to it, right?" And so, this really comes from the history of attitudes where we think about being thoughtful about attitudes, and that fits well with our scientific approach to things. When you think hard about something, it's gonna stick around. But basing your opinion on emotional reactions, that sounds like that could be strong too. And there's some evidence of that in the '80s and '90s, but my hope was to kind of rejuvenate that line of work more recently and talk about the strength of basing your opinion on an emotional reaction.

And so, we find for instance, and we can talk about any of this if you want, we talk about attitudes based on emotion, you know some work that we've done on this, tend to last longer. They don't fade as quickly. Your positivity sticks around for longer time if you base your opinion more on emotional reaction. And that doesn't say anything about whether being thoughtful or not leads to

a long-lasting attitude. It seems to, but we're also saying like, "Well, but so could an emotional reaction." Right? The more emotional you are toward this thing, it could stick around for longer.

Andy Luttrell:

But this is the part... You just acknowledged that we've done some of this together, but I'm gonna keep my interviewer hat on and ask you the questions, which is why? What is it? Do you have any sense? And I get the sense that this is still an in flux part of the idea of emotionality, but what is it about basing my opinions on my feelings that gives those opinions strength?

Matt Rocklage:

Yeah. So, I'll give you a little history and then I'll jump into what I think the process could be, or the possible process is. So, historically, you asked someone, "Well, does this thing make you angry? Does it make you excited? Does it make you happy?" And that would be ways of measuring how much emotion they have in their opinion. But those also differ in how extreme they are, like if I'm happy, or I'm angry, that means I'm also very extreme in my opinion, right? But conceptually, we can differentiate those, and empirically we have, where I can think something's really positive, so I think something is excellent, or perfect even, and that has relative emotion, but not as much emotion as fantastic, amazing, exciting, right?

So, what we wanted to do is say like, "Okay, we know that people can be more or less positive. Let's hold that constant and let's say above and beyond just how positive someone is, what if it's based on feelings?" And this gets back to the original theory about attitude basis. What kind of information are you basing your opinion on? Emotion? Or this kind of more thoughtful approach? So, we wanted to go beyond simply how positive or negative someone is and say, "Okay, what are you basing your opinion on?" And that's important for the process because if I have a feeling that accompanies my reaction, that really gives me some interesting and powerful information, right?

I'm actually having some sort of feeling when I experience that thing, or when I think about that thing, and that again provides me with information that I can use to make decisions. So, I sometimes use the example if I'm shopping for a movie and I see the movie and I say, "Well, that's a really superior movie. I think it's a great movie." Okay, yeah. I'm probably gonna purchase that movie. But what if I say, "Oh, that was a really enjoyable movie. It was really exciting, and it really was amazing." Okay, well, now you added more feeling to that and it's that feeling that gives me information beyond just being equally positive toward it.

So, it's based around the idea that feelings give me information, which has a long history in psychology, but we're bringing it to attitudes stored in memory. When I bring those attitudes to mind, that can really drive my opinion.

Andy Luttrell:

It strikes me too that it's harder to counter argue my feelings. Or you know, like, "Yes, I'm getting information from my feelings, but they're mine." It's undeniable they're happening. I'm feeling this way. Whereas with information, you go, "Oh, even if I came up with a reason to accept this idea, I'm still just trying to approximate the truth that's out in the world and anyone could come around and convince me that something else is true." But my feelings, you go, "It makes me happy

when I watch this movie.” You can tell me a million times it’s a garbage movie, but if I still laugh whenever I watch it, that... I can’t shake it. That’s just part of it.

Matt Rocklage:

Yeah, so the way that we’ve talked about this, so Russ has some theorizing on this, and then we’ve kind of built on it since his original theorizing, but there’s something special about a feeling, right? Like you say, it’s particularly undeniable, right? It provides me with a very clear sense of my opinion that potentially my more thought-based or more thoughtful, if you want to call it that, reactions don’t necessarily give me. Now, whether or not I should be basing my opinion on my feelings, you could maybe reverse the effects of strength, so like, “Oh, if I tell you your feelings are illegitimate in this case, you might reverse it.” But everything being equal, why would I mistrust my feelings? I’m just gonna go with them unless there’s other information to overturn my use of my feelings.

And so, over time, people probably come to believe that my feelings are a valid source of information and I can use those to make decisions.

Andy Luttrell:

Are there any other examples other than that they last longer that sort of highlights to you that emotion packs a powerful punch in opinions?

Matt Rocklage:

Yeah, so let me think of some good examples. So, one of my current favorite lines of research that I’m working on right now, it goes a little bit beyond just personally-held opinions and just talks about emotion as a general signal. And what I mean by that is, so imagine you want to predict the success of a movie in the future, or you want to predict the success of books in the future, what kind of information can we rely on to understand what’s gonna be successful in the future?

And so, I have some research showing that in fact star ratings for online reviews, for instance, you’d think that oh, that would probably reflect the success of a book or a movie in general in terms of how many books are sold, or in terms of revenue that a movie made. In fact, though, star ratings are very inconsistent in their predictive ability of success. And in fact, sometimes negatively predictive, so if this movie got a five star average from reviewers from let’s say IMDB.com, or RottenTomatoes.com, or Metacritic, in fact that can predict doing worse at the box office.

And so, what we wanted to talk about as well, if emotion acts as this kind of general signal that something important has happened or impactful has happened, perhaps we can aggregate emotionality at like a population level and use that to predict the success, future success of movies and books. And so, we brought this idea of a signal and the information that emotion can give us and said, “Well, if we aggregate these opinions, that’ll also give us like a normative idea of the impacts of these movies or books.” And so, yeah, we find that taking the first 30 reviews for a movie online predicts in the future how much revenue that movie will go on to make, whereas the star rating is very inconsistent and even sometimes negatively predictive. And so, again, just taking that idea of information that we can get about emotion and as a signal, and use it to predict something strong, right? Something that happens that’s impactful in the future.

And we can talk about what we think the mechanisms are there. We haven't nailed those down. But nevertheless, that's another indication that oh, emotion seems powerful, and that there's something special about it.

Andy Luttrell:

It seems... Like, I'm trying to put myself in the shoes of someone, like let's say a movie gets just panned and people are like, "I hated it. It made me mad." There's something just intriguing about it where you go, "I want to know what that feeling was." That's just... Maybe it's unusual, but maybe that's not what it is, whereas if someone just said like, "It just wasn't good. It just was poorly written." I'd go, "Oh, I don't need to see that." But if someone said, "It made me so disgusted that they would even spend a dime on this." I'd go, "I kind of... I want to know what that was about," in the same way that someone who said, "I just have never felt elation like that at a movie," I'd go, "I want to feel that too."

But if you go, "Oh, it was... They nailed it from every inch of the script was perfect." I'd go, "Oh, good. Cool. I'm glad that's out there."

Matt Rocklage:

Well, that's interesting, Andy. I think that might give an insight into your psychology. So, right, I thought this movie was awful, terrible, the worst movie I've ever seen. Andy's like, "I wanna see that movie. That's the movie that I want."

Andy Luttrell:

Check it out.

Matt Rocklage:

Yeah. No, but it's interesting because we... In this line of research, we couldn't really investigate that, and that was part of the problem that we were trying to address, that in fact, upwards of 80% or more of online reviews are positive, which is kind of surprising because everybody talks about how negative online stuff is and how angry people get, et cetera, et cetera. In fact, online reviews are overwhelmingly positive. Half of all reviews on Amazon I think are five star reviews, right?

And okay, sure, maybe products are really all that great, but I don't know if all products are really that great. So, in any case, we weren't able to investigate for instance if I say something is disgusting, awful, terrible, about a movie, what that means for the future success of that movie. There's just not as much data on that end of things.

Andy Luttrell:

I have a question that there's a very good chance you don't know the answer to, which is about how Amazon works. I have noticed that when you go onto a product page on Amazon and it lists all the reviews, it could be something that's like four and a half stars on average, thousands of reviews, everything looks great, but the first page of reviews that they display are all one star reviews. Which, I go, "What?" That doesn't make any sense for anyone, right? It's misleading to the consumer. It's dissuading people from buying something that might otherwise be good. I just...

I can't wrap my head around it and I just didn't know if maybe you had, in weaseling around these areas of the web, know what's going on there.

Matt Rocklage:

That's interesting. Yeah. I had that same experience that you're talking about where you say, "Oh, this product looked great. I should just go ahead and buy it right now." And then you're like, "Well, I guess I should read the reviews just to make sure, because the star rating's high, but what do people say?" Oh, and then you jump in and the first page is like one star review, one star, and they're like people talking about, "Thought this was gonna be great, terrible." So, in terms of the algorithm that Amazon uses, it's proprietary because they don't want people to try to game the system, right? They don't want people to get their review at the top if it's not helpful or if it's being bought in some way.

So, if I had to take a guess, and maybe you were thinking the same thing, you know, it's about diagnosticity. I'm giving you information that is very different from what other people are giving you, and yet it seems really important, right? So, if 80% of online reviews are positive and I give you a very well written, thought out, negative review, that could make it to the top of the website or that product page because it's giving me information that is at odds but still seems important.

Andy Luttrell:

That's what I wondered. If it's somehow built on helpfulness ratings and people like, "Oh, thank God someone had the bravery to say this." But almost to your... The positivity bias that you just talked about, and then the other world that you've soaked your feet in, which is the negativity bias and things that are related to it, I wonder if you'd see the same thing if the roles were reversed, where if overall most of the reviews were negative, like 80% are one, two star reviews, and then someone has the bravery to come along and say, "I loved it! It was really great, and I've used it for years," people would just go, "That person's crazy." Or they would discount it in a way that they're not discounting the negative reviews for mostly positive stuff.

Matt Rocklage:

Yeah, so I'm happy to talk about that in terms of valence, but I could even bring that to emotion, because I have some research on that. But what we find is, okay, I was telling you earlier that people turn to emotion to persuade, right? So, people believe emotion should be persuasive. And then we investigate, "Well, is that the case?" Does emotion written in these online reviews actually act as more persuasion than not? And we find that it depends, and so for instance, if I am reviewing a movie, or a fiction novel, or something that makes you want to feel something, that's why you purchase the product, then people should go to town using emotion. The more emotion that people use to describe those books and movies, the more helpful those helpfulness ratings on Amazon, was this review helpful or not, those helpfulness ratings increase.

In fact, we don't even find an attenuating effect where you can just really go to town on emotion. There's no diminishing returns on that. So, that gives us an indication that greater emotion, yeah, you should use it. It seems to be more persuasive.

However, if you're describing a product where that's at odds with why you use the product or how other people are describing the product, such as an electric shaver, or a book on how to program

computers and using Python, people start to look at you like you're very strange. Wait a minute. You're telling me how enjoyable, exciting, and fun it was to use this programming book? Okay, I guess I could see that, but so I don't know anything about you, you're an anonymous reviewer, I'm gonna disregard that review and move onto the next review.

So, in that case, in fact, emotion backfires, because you lose trust in the reviewer. You don't know this person and you don't know much about the product, but that doesn't fit with what you think about the product. This person's giving me weird information. And so, there, emotion, yes. Backfires and leads to less persuasion.

To your original point about valence, if you're giving me information that's at odds with other people, in some cases if it seems really thought out like we're talking about, that could be very persuasive. If, on the other hand, you just seem like an oddball, I don't know whether to trust you or not, and so I'm gonna move onto the next review. You might be right. You could be very right. But if it doesn't seem well reasoned and you're just kind of shouting how emotionally evocative this product was, I'm gonna move on and say, "No, that was not helpful."

Andy Luttrell:

Was the outcome you looked at always helpfulness of the review? Or did you look at actually minds changed less when it was an emotional pitch for a not emotional product?

Matt Rocklage:

Yeah. We had a couple different outcomes that we looked at. So, the primary one was helpfulness, because this all stemmed from the realization that, "Wow, we have a lot of online data. There happens to be this helpfulness ratings. Can we identify situations where people are more or less helped by content?" So, that was where we started, but then we said, "Okay, does it stop just with helpfulness? Or does it have implications for people's opinions of the product itself and their willingness to want to purchase that product?"

And so, I think the most direct evidence that we have in terms of importance from a marketing perspective, and for those of you who couldn't see, I used air quotes for importance there. In terms of a marketing perspective is what do people ultimately choose? So, we did a raffle where we said, "Okay, we're gonna show you these reviews and you get to choose. You get to choose a review for this, or you get to choose this microwave, or you can choose the blender." And we wanted to see if I used different language to describe them, would you actually stay away from the product that was reviewed emotionally if that's at odds with the nature of the product? So, would I take a blender if someone said amazing? Well, yeah, okay, a blender might be amazing, and fun, and exciting, but that seems weird. I'm gonna disregard it.

And in fact, we showed that they're less likely to enter the raffle for the product that's at odds with the emotional language, right? So, I'm less likely to enter the raffle for a blender if it's described with exciting and fantastic, because it just seems to be... I don't trust that this is actually a great blender. I'll go with the microwave, which was described with less emotional language.

Andy Luttrell:

It reminded me, I think a lot about this episode with Alex Coppock earlier this last year, where he looks at political ads and has... kind of comes to the conclusion that like, sure, people might like ads that agree with their predispositions or that sort of match their expectation, but that's irrelevant to whether people actually are convinced, right? So, Democrats might actually be convinced by Republican ads, even though they say, "I don't like them." And so, that's why I wondered if it was like, yeah, everything in the world feels right when emotion gets used for emotional things. I find that more helpful. And it seems wrong when it's for a book about programming, but nevertheless, at the end of the day if you're excited about this programming book, I might buy it. Even if it doesn't seem like you know what you're talking about, I still am now curious.

Matt Rocklage:

And I think that is a really important point, because I think there are context effects. In an anonymous online review setting where I'm trying to discern whether the product is good and whether I should trust the reviewer, I have to do that simultaneously, context becomes really important. I don't know this person. I don't know the product. I need information that I think is trustworthy. Whereas for... So, if for example... Jonah Berger, he has research showing that more emotional New York Times articles are shared more. So, there you're seeing a spread of information and he finds that emotional information spreads more even for more mundane topics, like... I don't remember, I would say finance or something like that.

Those are still shared more, and the context is important there because in mine, the online reviews again don't know whether to trust this person or not. New York Times, though, that's a very trustworthy source. So, I'm willing to believe that this thing happens in this emotional way, and I find that very engaging and interesting, so I'm gonna spread that information to others. So, there, with commercials we get into another realm of context, like okay, I know that this advertiser is trying to convince me of something. Are they going to the lengths required to convince me this ShamWow is really amazing and fantastic? Which is partly why infomercials are a thing, because I need to convince you that really this thing is amazing, and so I'll go ahead and feature it in ways that really try to convince you of the emotion-evoking nature of this product.

Andy Luttrell:

I thought to sort of wind down we could go back just a little bit to the idea of language as a tool for understanding opinion. And almost just give you an opportunity to talk about how to do that well and how to do it poorly. And so, as I was thinking this morning before we got online, I was remembering I was talking to a colleague and I was interested in doing some sort of linguistic analysis of something, and they were like, "Oh yeah, I want to do this project where we look at the words people use on these forums in this very niche field." And I go, "Oh, to what end?" And he was like, "Oh, just to get the sentiment scores and just to know what they were saying." And I was like, "Oh, okay."

And so, I wondered like in the grand scheme, where do we get a lot of mileage out of looking at language, and where are we maybe lured in by something that ultimately is challenging to learn much from?

Matt Rocklage:

That is a huge question and a very excellent question, I might add. I think I could talk about this for a long time because ultimately I didn't plan on getting into the science of language measurement, but if you want to measure language, you better get good at understanding the limitations and the benefits of the measurement process. So, in terms of doing it good versus bad, a lot of early language research was like a fishing expedition, right? So, we've got this text. Well, what can it tell us? And that has its place, right? Descriptive knowledge is very important.

But what that also led to was a literature just kind of riddled with inconsistent measures, so you're like, "Oh, this person came up with a measure, this person came up with a measure, I'll see which one works and I'll use that one." You see that a lot. And that's a very bad way to do it, because in psychology, and in marketing, and in many, many, many other disciplines, if you're not measuring the same thing as someone else, you're talking past each other, right? And that's really what the history of language analysis was. It's like, "I'm gonna measure this thing and choose the thing that works. You might call it one thing. I might call it another thing. It might be made up of different words, but it doesn't matter. It's a thing and I want to publish this thing."

And so, I think the new generation of language analysis needs to be more construct driven, right? So, it needs to be driven by someone who comes in and says, "No, this is a measure that we should all try to use, and I'm happy to be swayed one way or another based on if we find that something is a more accurate or more valid measure of something." But we need to agree at least on what constitutes a good measure of some certain construct in language and then we rally around that construct and then use it to measure things across different domains.

Andy Luttrell:

Can I pause you there for a second? Is there any example you could give of what you mean by people disagreeing about the measure or just picking the measure that works? Or-

Matt Rocklage:

Oh, yeah. I don't want to step on any toes, because-

Andy Luttrell:

Or even make something up. I mean, just to sort of be concrete about what would it mean for people to say they're studying the same thing, but not actually be?

Matt Rocklage:

Right. So, because it's easy for me, I'll use emotion. The original thing that I talked about, it's like, "Oh, sometimes people are very positive and very negative." And what's happened in the literature measuring sentiment, it's called, or attitudes, or opinions, whatever word you want to use, what happens is that people sometimes call that emotion, right? Just how positive or negative someone is, because oh, if they're more extremely positive, that must mean there's some emotion there. But yet we know that extremity of someone's opinion is different from emotion. And so, people will measure emotion by saying like, "Oh, this person's more positive or more negative. That must mean emotion." Whereas it's like, "No, no, no. Emotion has feelings behind it, and these may or may not have feelings. Are you really measuring emotion?"

So, that's just one easy, straightforward example of like, well, there's disagreement, or at least people talking past each other willingly because they can get more publications if they don't try to address this question, where it would be very helpful to agree on what constitutes a good measure of X, Y, or Z, and then proceed from there and test hypotheses instead of doing these fishing expeditions where I just test a bunch of different measures and see which one gives me something sensible that I can publish.

Andy Luttrell:

I wonder if a lot of that is due to a lot of the language analysis coming out of other fields, like computer science, where in psychology it's just the, "Well, that's what we do. You define your variable. You make sure that your measure is an accurate reflection of that variable, and then you ask bigger questions." Whereas if your point is like, "Oh, can I turn this giant backlog of words into numbers?" Yeah, let's do whatever efficiently does it in a way that seems kind of interesting.

Matt Rocklage:

Yeah. It's... No, I wouldn't put this on computer scientists. That's the way that they tend to think. How can we enhance predictive ability of these language measures? And so, they don't really think about the underlying construct, but psychologists are equally blameworthy in this case. So, early measures basically just took the kitchen sink and said, "Well, let's see what differs." And then what happens is that these findings get into the literature and then if you submit a new paper, well, now you have to reference the old paper and say, "Well, we used the same construct there. We have to use it again here." And then they kind of just snowball where it's like no one really validated this initial measure. You were just doing a fishing expedition. And now it's made its way into the literature and is established, but should it be established? No one's really looked at the underlying construct of what you're actually measuring.

So, no, I wouldn't put this all on computer scientists. I would say it comes from where language studies appeared in the beginning.

Andy Luttrell:

Is there any other thing that you wanted to... Because I cut you off when you were talking about the right and wrong way to do language work for the example. I just wanted to give you the chance to go back to that if you were going somewhere else.

Matt Rocklage:

Yeah. If I were to sit on... Sit on. If I were to step onto my soap box, it would be language analysis I feel like needs to move from a fishing expedition of descriptions of what the text is about to more hypothesis-driven questions around really strongly validated language analysis tools. And so, I feel like that needs to be the next step. I don't know if that will be the next step, because people are reluctant to put in that time and effort to validate these kinds of measures. It's very difficult to do. Time consuming. But my sense is that should probably be the next step if we're gonna get the most out of these language measures.

Andy Luttrell:

So, those are the big next goals for science as a whole, but for you specifically, as we wrap up, what sorts of directions are things moving in terms of when you look at words, and emotion, and opinion, what is still that itch that you're trying to scratch?

Matt Rocklage:

Yeah. So, my niche within marketing is a lot of marketing computer science, and even within psychology to a lesser extent, has really been about measuring positivity and negativity. Is your language positive or is it negative? Okay. But I think we can also move beyond just positivity and negativity, so one way we can do that is emotion, but I'm also developing a tool that we're calling the Certainty Lexicon. We're getting very creative by just instead of calling it Evaluative Lexicon, Certainty Lexicon. And it's the idea, well, I can be more or less certain about my opinion as well, and work that you've done and others shows that that belief about my attitude, whether it's I'm very certain about it or not, has real important consequences.

And so, the next frontier and something I've been working on for the last couple years is measuring, again, these factors outside of valence and positivity that really inform us about someone's attitude. Because that's what computer science has really been focused on, just valence, and marketing as well, so I'm hoping to expand horizons and try to introduce the attitudes literature to these disciplines who may or may not want anything to do with me but trying to do that anyway.

Andy Luttrell:

Very cool. Well, I will look forward to all the things yet to come. Thanks so much for taking the time to talk about this.

Matt Rocklage:

Yeah. Thanks, Andy.

Andy Luttrell:

All right, that'll do it for this episode of Opinion Science. Thanks a bunch to Matt Rocklage for coming on to talk about his work. Check out the show notes for a link to his website and links to the research that we talked about. You'll also find a link to some software that Matt made that lets you analyze text using his Evaluative Lexicon. By the way, there was some clicking in the audio throughout the interview. Sorry about that if it bugs you. I have no idea where it came from and I couldn't figure out how to do anything about it. Just the trials and tribulations of podcasting.

All right, go ahead and do all the things. Subscribe to Opinion Science. Leave a review. Tell the world. And follow me on social media @OpinionSciPod. Okay, and that's all for now, so I'll see you in a couple weeks for more Opinion Science. Bye-bye!