Welcome and thank you for standing by. At this time, all participants are in a listen-only mode. During the Q&A session you may press star one on your touchtone phone if you would like to ask a question. Today's conference is being recorded. If you have any objections, you may disconnect at that dose at this time. I would now like to turn the conference over to Dr. Jim Kurose. You may begin.

Thank you very much. So welcome, everybody, to the first of our 2015 distinguished lecturers. For those few I haven't had the pleasure to meet yet and I will meet you soon, my name is Jim Kurose, the new AD for CISE. This is day 18 for me. Every day almost everything I do is new. Yesterday congressional testimony that was new, this is really a joy coming to the first distinguished lecture that we have and the opportunity to introduce our speaker. Before I introduce our speaker I want to remind people that we got a great lineup of people coming in to participate in the Distinguished Lecture Series. The next lecture will be given by Lucy Sanders. And that will be on Wednesday, February 11, 11:00 a.m., room 1235. That's just a look ahead. It's my great pleasure to be able to introduce Professor Dr. Robert Kraut from Carnegie Mellon. He's the Herbert Simon Professor of human computer interaction in the school of computer science. And also in the temper school of business at Carnegie Mellon University. He served as visiting faculty member at Facebook and a visiting fellow at the University of Canterbury in New Zealand. He's worked at Bellcore and Bell Labs and has positions that Princeton, Cornell and then -- Pennsylvania. He's got a broad range of interesting design and his recent research has focused on the analysis and design of online communities. He'll be speaking about that today. He's the author of a number of books. His most recent is the 2012 book building successful online communities, evidence-based social design and is the author of many, many research articles and chapters. He's terms -- serves on a number of editorial boards. He serves on the telecommunications Board of the National Research Council. And he's a fellow of the ACM and also of the American psychological Society. Thank you, Robert for honoring us with your presence and thank you very much for presenting on this Distinguished Lecture Series. Please join me in welcoming --

---

[Applause]

It's my pleasure to be here and I also feel I have to give back. the research that you're going to hear with NSF sponsored -- most of the research that I've done has had NSF funding. What I'm going to do today is to talk about how support it gets generated. Some of it is in from social dynamic consequences in health support groups. So online support communities are places where people gather to share experiences. People often with some physical or mental disease. They share experiences, ask questions, give each other emotional support, provide advice about what people should be doing. They are heavily used including one of the most heavy uses is among people with cancer. And there is some kind of clinical trial evidence. Not a lot of it. That participation in online support groups has beneficial effects on people's mental health and health quality of life. But there's little understood about how the support that gets generated in these groups actually gets produced. Since it's a social dynamic that produces it, where volunteers are helping other people, and there's little knowledge about what is the path by which the support that's produced, the various resources that are produced in these help support groups, have their effects? Most research is focused on whether the clinical outcome, whether participation in these groups makes people healthier and happier. And less research on how the support is produced and its effects on things like how people -- how long people stick around. What I'm going to be doing today is asking, what is it that people who are trying to get support for these groups, which is why they join, what do they do to elicit support? to the providers who are giving them support, are they able to diagnose what it is that people are seeking and give them the right kind of support? Are the seekers satisfied with the kinds of support they get? and those receiving support change commitment in the group? And then in some research that I'm just starting, we'll be looking at what are the particular kinds of support or resources that are exchanged in these groups that change health quality of life as well? So overall research goals is how do these conversations produce support? What are the consequences? I'm going to be talking about commitments to the group and satisfaction with particular support of exchange and then in order to answer these research questions, we need to have a way to measure the resources that are exchanged in these groups at scale, because there's too much happening to be able to use conventional social science methods of having undergraduates read them and code them manually. So we'll be looking at ways to automatically measure what goes on in these groups. So the research agenda is looking at the support that generated this middle box, getting emotional support or getting informational support. Part of the research is going to be asking what is it that people do in order to get the support. What is it that the seekers or investors do conversationally in order to get the support? And then we'll be asking what consequence does the support they receive have on their satisfaction of particular exchanges and their commitment? In this talk, I will be talking about the consequences for their well-being. So let me just give you an illustration of the kinds of support that we're talking about. So the prior research literature focuses on primarily two different kinds of support. Emotional support which provides understanding, encouragement, sympathy, affirmation for care and from one person to the other. So people getting this presumably feel that others care for their welfare. And an example would be from the cancer support group we've been looking at, Julie, you've had such a difficult role yet you still managed to do well in school. I'm impressed you. Because and test wishes to you. Human raters judge this as having 5.7 on a seven-point scale of emotional support, but no information is being exchanged. People are just saying we're concerned about your welfare. On the other hand, the next example, informational support which provides advice, referrals, knowledge, information, an example would be the extra nodal extensions occur when tumors extend through the wall of the lymph node. This is noted on pathology reports. The main isn't very significant and isn't used in assessing the cancer stage. One person is telling another how to interpret her radiation results. And here it's just the facts. and no warmth is being exchanged. No understanding or caring. and I'll mention this later but essentially these two kinds of supports are uncorrelated. Many messages have both some messages -- some messages are people talking about their vacations or their dogs and they don't have anything to do with the illness or they have neither. And some messages have one or the other so the correlation is about a minus .17. There's a negative correlation but it's not a very big one. and so what I'll be first asking because it's the core methodological question is how do you measure support at scale? and in this talk I'll be focusing on being able to measure it at scale as a basis of scientific understanding. What produces it and what consequences it has. Essentially a quantitative observational study. But you can imagine that that same way of measuring support automatically can then be used as a way of doing dynamic interventions in people's communication, nudging them to get the support they seem to be looking for, but your machine learning models seem to say if they are looking for emotional support. Somebody is giving them information so this is meant don't understand. Women are telling you what problem -- men give them the information and what effect they are looking for is emotional support. If you see that automatically that's happening, you might be able to repost that message to someone who in the past has the same condition and is emotional support provider. You can think about dynamic interventions as well. So then we'll be asking what participants due to elicit support and then finally downstream what consequence of receiving the support has on satisfaction and commitment. The research site that we've been looking at is breast cancer.org. It's the largest disease specific support community in the United States, maybe the world. Nonprofit dedicated to providing reliable, complete up-to-date information about breast cancer. Their mission statement, they focus on information. But an awful lot of this emotional support among people who recently diagnosed or have had it for a while and want to give back to the community also produce that sort of thing. It has about -- when we did this research about 150,000 registered users, when we did this research, 1.5 million posts organized in about 70,000 threads. And those are in turn nested within different forms, generally organized around disease type. Breast cancer, pancreatic cancer, colorectal cancer, or types of disease or particular demographics. 'S so mothers with young kids, might be different from 55 and 60-year-olds. Or particular problems that people with rest cancer have, like finding Whigs or breast reconstruction. So they deal with particular topics. and what I'm going to be describing here is probably a very standard -- probably -- people understand it reasonably well. Standard machine learning set of steps to measure this sort of thing at scale. So first, you have human judges whousing human understanding, assess what is happening in some of these messages. So in particular, they are making judgments about the amount of emotional or informational support that is in a sample of messages. Then you take those messages and represent them into a set of features that machines can understand. Topics and particular words and grammatical forms, then use machine work -- learning algorithms, some variation of regression, more sophisticated than I do as a social scientist but machine learning algorithms that produce the best correlation between some combination of those features and the human judgments in order to get a model that predicts the human judgments. You've got that elevated, you apply that model to the rest of the data. And then use that imputed emotional support or informational support to the rest of the research. So in order to get the human judgments, we use mechanical trip workers. So a lot of research shows that you can get naive judges -- if you get enough of them -- all of their errors cancel out. And you can get enough of these naive judges to make judgments that are as good as trained research assistants are experts in the field. So we had 10 mechanical trip workers. Amazon service where you can hire short-term contract workers for very little money, below minimum wage although we actually pay minimum wage. We had 10 workers each -- 1000 messages each, weighted by 10 different workers, rated different messages. And they were answering the question, how much emotional support does this message contain? I showed you some examples. How much informational support does this message contain? They make highly reliable judgments. So the workers agree with each other about how much emotional or informational support. Intraclass correlations are very high. They are reasonably accurate in the sense that they correlate with what expert judges have judged. So correlations, .7 or .75 between the average of the turf workers and expert judges for a subset of messages that we had judged by both of the naive judges and the expert judges. and as you'll see later, the judgments that predict human behavior reasonably well. So they are useful for understanding what goes on in the group. So they have in that sense, they are useful. They're valid in that sense. We then represent these messages as a vector or a set of textual features and we use three different kinds of textual features. We used to some general-purpose dictionaries. This is the dictionaries that Jimmy Pennebaker and his colleagues at University of Texas have generated, that are words that have high psychological value. Then we used topic modeling to find a set of domain specific dictionaries in this case, cancer specific dictionaries. And then we have a miscellaneous set of other features. So we've got a total of around 60 features that are being used to predict these judgments of emotional or informational support. So I can go over more of this in the question and answer part if you've got questions. But to just get through some of the technical details more quickly, we used these loop dictionaries, linguistic and word count dictionaries. These are generic dictionaries that include dictionaries of first-person singular pronouns or I dictionaries or first person first-person plural, we dictionaries, words that don't come-- or negative emotion words. Words having to do with anger or sadness or anxiety. and other relevant topics from the Pennebaker dictionary. Cognitive mechanisms, people's discussions of time, religion and death for example were used. Then we did LDA topic modeling wearwe feed something, 40,000 randomly selected messages into something that's like a big factor analysis that is an unsupervised learning technique that finds what words cluster together that you can use to predict the words that are going to be in a particular message. And you get dictionaries out or topics out that make sense and are reasonably homogeneous. So I'll show you some examples but they are things like -- that are illness specific. Dictionaries having to do with the things you might do before you get diagnosed. Like biopsies. Or things having to do with treatment plans. Or things having to do with post surgery problems or things having to do with hair loss because of the radiation and chemo et cetera. Then we created a set of ad hoc features. Things like word count and -- sentence count and a measure of complexity which would be word count% -- per sentence, terms like not and no, part of speech tags, and then a rule that we called and advised tag. Please do or an example. If you could, if you would be -- another rule, rules that allowed us to extract questions. Question marks, does anybody know? That would be an example. Indirect questions, things -- words having to do with states as opposed to actions. Dictionaries from sentiment, subjectivity dictionaries that are weak, positive and negative sentiment. And then a set of medical terms that we got from FDA dictionaries. So those would be 60 features in a regression equation predicting the average emotional support or informational support that the TRC workers thought was in a message. He would be some examples of these topic dictionaries. So pre-diagnosis would be things like total appointment wait and back. Those are the most informative words associated with that topic. Or surgery would be breast surgeon, mastectomy, chemo and radiation topics. Chemoradiation treatment, family, dictionary would be things like mom, children, age, younger, talking about their family circumstance. So automatically measuring the amount of emotional support and informational support in a message is almost as good as the correlation between any two people. So the correlations between a measure of accuracy, between the machine measurement of emotional support and human measure of emotional support is .8-ish. Similarly the correlation between the machine measurement of informational support and human judgment of informational support is .85. And the features that predict whether it is emotional or informational support also makes sense. So when people are talking about emotional support, they talk about us and we. They talk about first-person plural pronouns. and they use a spiritual terms. Pray for you, they use God. They use the dictionary, the LDA dictionary about adjusting to your diagnosis. Or the LDA dictionary having to do with emotional reaction. But they don't objectify people. They don't talk about others. They don't say she and he. Or they don't use the vice rule. They don't give advice to somebody. And they don't talk about drug terms. Which you might think would be more informational. On the other hand, informational posts, words that are judged as more informational, or longer sentences. So more complexity. They don't have subjectivity terms in them. They do have the advice rule. They don't talk about God and prayer. They don't talk about we. They do talk about surgical problems. They talk about treatment plans. They talk about diet, they talk about surgery. So there's in addition to being accurate in the sense that it models what humans do, these dictionaries make conceptual sense. Okay? So all of that was to say we can actually measure whether someone is getting informational support or emotional support reasonably accurately and a scale. Now we can apply those measurements to the 1.5 million messages and the data set and we can ask, what is it that people do in order to get others to provide them emotional support? Or others to provide them informational support? There's reason to think those aren't going to be the samestrategies or conversational moves in order to get those two kinds of resources directed toward you. So the prior research, there's suggestions from more qualitative work that would suggest when people want information, they ask for it directly. And I think the reason they do that is because it's a focused way of getting what they want and there's no implications for their face. Either the face of the person giving them the information or failing to give them information, or save their own face if they don't get it. So if somebody doesn't get them information it probably means they don't know. And it doesn't mean that I'm an unworthy person. If somebody's failing to give me useful information about what I have. On the other hand when people are trying to get informational and emotional support, they do this indirectly. And I think for these face-saving reasons. So first of all, emotional support is probably less valuable if you have to ask for it. So information is valued -- valuable whether you just find it on the street or ask the reference librarian or somebody in this community. On the other hand if you have to ask somebody, think about this if you have a spouse or significant other, if you have to ask whether they love you, then them telling you that they love you is devalued as a sign of their affection. It's the same thing here. If you have to ask that we want emotional support, probably the asking devalues what you get in response. In addition if you don't get emotional support -- anybody can provide it. They don't need to know the information. They have to say I hope the best for you. That has an implication for your face. I'm not a worthy person. If I seem to be asking for emotional support, and I'm not getting it, it probably means there's something wrong with me that no one has given it to me. As a result, my guess is, people do face-saving moves where they can deny that they're really asking for the support. So that would be a reason why you might expect that asking for informational support is direct, asking question, while getting emotional support or seeking emotional support is going to be indirect. You can deny you were asking for it in the first place. So in order to understand what it is that people do in order to get support, we look at characteristics of the thread is starting messages and we looked at five large dimensions. I won't go into the details of how we measure these things at scale but it uses the same having Distinguished Lecture Series -- TRC workers, textual features to predict their judgments and then applying those models to all of -- and we tried to get negative emotional self-disclosure which is a high-level concept, and positive emotional self-disclosure, negative informational self-disclosure,positive informational self-disclosure and question asking. Then we created machine learning models for this. So here's an example of an emotional self-disclosure. Somebody saying I'm giving up in treatment for the past two weeks. I've been crying and praying and meditating and talking to my love ones about declining more treatment. I know it's a huge step and my oncologist warns me if I continue on without it, my liver comes so degraded. She's struggling with this major decision about the side effects are too much. Should she top stop the treatment knowing that she's going to die and she won't be able to reverse her decision? She's describing her feelings about making this decision. Or she is describing an external event in her life that went badly. This is the complications that were associated with getting a port installed in her chest for medications and drainage and that sort of thing. She's telling you for all the bad things that have happened, but it's objectified. Events in the world that happened. We also did the same thing for a positive self-disclosure revealing good thoughts and feelings that you had. So my radiologist just called, told me the pathology report came back benign and I can't believe it. I'm so relieved. She's expressing her elation that she's in remission. Or positive informational. Events in the world. This woman is talking about the wonderful Mother's Day session she had with her kids who did very nice things for her. She's describing events in the world. Then we also measured at scale question asking. So here's somebody who is asking questions. We got reasonably not as high but reasonably good accuracy for our machine learning models that measured these things. And then we applied to those -- if you think about these five characteristics as independent variables. These are what people do in their thread starting messages and then we ask as the dependent variable, what's the consequence? How do others respond to them? What you see is there's a different pattern of what you say in order to get informational support and what you say in order to get emotional support. for getting informational support, the strongest predictor is respect. And simply asking a direct question or indirect question about something that humans would recognize as an explicit question. As you also talk about negative events in the world that you are generally asking about, or positive and not talk about positive events in the world and you're not talking about your thoughts and feelings, these are objective eyes to things and when people do those three things, the next message they received tends to be informational support. On the other hand, getting emotional support is associated with self-disclosure. So negative self-disclosure and positive self-disclosure both are associated with receiving emotional support. Talking about negative things that have happened in your life is associated with receiving emotional support. This one. .18. and not explicitly asking questions is associated with getting emotional support. So people seem to be doing different things, in order to get information on emotional support. To make sure that we're not simply finding the effects of behavioral mimicry, we also had humans judge the speaker's intent in their first message, whether they seem to be seeking informational support or seem to be seeking emotional support. And most of the effects of these language features have their effects by changing the perceptions that judges have that someone is seeking informational support. So mediation is occurring through intention. When other people and communities see that you are seeking informational support, they provide it. When they see you are seeking emotional support, they provided. That's the reason you see relationships between low-level language features and the kinds of support that's provided. Now we can ask, what's the downstream consequences? I'll talk about two June consequences. One is satisfaction. with a particular exchange. That's kind of short-term, whether someone seemed to express that they got some value out of the interaction that just occurred. And then the other is going to be much longer term, whether people who are exposed to different amounts of informational or emotional support stick around the community longer. So you might think that as a measure of aggregated satisfaction over the long-term. We can ask, since most people are coming to these communities in order to receive informational and emotional support, the straightforward prediction is if they get it, they are satisfied. Somewhat more subtle prediction, and one that there is mixed evidence in the psychological literature is people are most satisfied when they get particular kinds of support that they are seeking. So going back to [Indiscernible -- low volume] don't understand, probably are less appreciative of the good advice that your husband gives you. If what you really wanted was a hug. So if there's a mismatch between what you were really trying to elicit by telling the horror story about this horrible day that you had and what the kids did and how the dishwasher broke down and horrible time you had at work, and your husband tells you how to fix the dishwasher, you're not going to get as much benefit from that as if you were asking about how to fix the dishwasher. So that's kind of the logic here. So essentially we're predicting interaction between what you were seeking and what you get. Leading to the satisfaction that you actually derived from the support. We had, in order to measure satisfaction, we again did this from the data. We would show TRC workers, the first post in an exchange and then the thread starters second post. And not show the intervening content which would be the communities replied to the first post. So here is -- can anybody explain to me what does it mean for the dots on x-ray to be clustered? I have many flex according to the results. Within this Pac-Man shaped thing. And there's some community replies. We don't show that anybody. And then the person who asked about the flex says, thank you so much for your thoughts. This website is fantastic. Everyone here is delightful. And then she gives more information about her state. People in this case would say okay, that person seems to be really satisfied with the exchange that she's got up to this point. We asked TRCers to judge overall how satisfied is the person? Overall how satisfied was the health information she received? How satisfied -- how much did this person's level of distress change from the beginning to the end of the interaction? And how much did she seem to express connection to the rest of the people in the community? Those things clustered together so we have a reasonable measure of satisfaction. Dependent measure is this composite measure of satisfaction. The independent variables are the number of replies that people receive because lots of research says that merely being talked to is itself -- leads to satisfaction. And then what we were especially interested in was the amount of informational and emotional support they received averaged per message, holding constant the number of replies that they got and then finally the interaction between those two. Between the type of support they sought and the support they received. And because some of the same factors might predict how satisfied you are with the reply, with community communication with you, with whether you actually give us data, whether you have a second post that would be analyzed. We used a Heckman correction because we didn't have all of the data. There were a third of the people when they asked for either emotional or informational support didn't reply to the community about whether they were satisfied or not. So we used the statistical procedure to deal with that missing data. and in general, just getting resources is valuable. So the more replies that you got, the more satisfied is your expression of satisfaction in your follow-up comments. the more emotional support he received, the more satisfied you seem to be. and the more informational support you receive, the more satisfied you are. So just getting resources seems to be a good thing. On the other hand, there is some evidence for matching. That is, for informational support but not for emotional support, if you are asking for informational support and got it, you are extra happy. If you're asking for informational support and someone is trying to give you hugs, emotional support, your less satisfied with that emotional support than what you would otherwise be. On the other hand if you were asking for emotional support, it doesn't matter what comes back. You are satisfied. Presumably because the mere fact that people are responding to you is a sign of their concern, whether they are responding to you by explicitly saying we're thinking about you, praying for you, hope everything is for the best, or whether they are giving you advice about reconstructive surgery, it's a suggestion that they are investing in and they care about. So just to give you this more graphically, how much informational support people receive, if they were asking for informational support, the more informational support they receive, the more satisfied they are. On the other hand, if they were not seeking informational support, then they are less satisfied with that informational support. Similarly, the x-axis here, emotional support they got. If they were asking for information, then getting emotional support doesn't up a lot in terms of how they are satisfied. But if they were not asking for informational support, or if they were not asking for informational support and getting emotional support, it makes them substantially satisfied. So matching only occurs for informational support, not for emotional support. Okay? And then the last outcome is, how long do people stick around? in every online community, there's high dropout. the average number of edits that -- new editors make in Wikipedia is one. the community we looked at here, you lose something like 50% of the people are never seen again after two weeks. So high dropout. And in clinical trials that have to do with online behavior, they talk about the law of attrition for any sort of Internet-based interventions, where 50% dropout of treatment that people have gone through a lot of trouble to get into, is typical. So they are not taking the treatment. Half of them are more dropouts. So many people leave before they can get the benefits. And the question we're asking here is, what's the consequence of receiving informational or emotional support on keeping people around? So we're doing this using survival analysis, asking does the amount of support that an individual is exposed to predict the length of that person's subsequent participation in the community after getting episode of support? and because we don't actually in this data set have data about what people read, we are making the assumption that if somebody posts into a thread, that they read everything before their post in the thread. And in another analysis that I want to talk them that I will talk about, everything in a thread. We are assuming if you posted into the thread you're reading all of the messages and we're saying your exposed to the average amount of informational support or the average amount of emotional support in that thread that you yourself posted to. You might be the starter but you might be third or fourth person but we will make the claim that you are exposed to all of that support. So here would be kind of a thread starter. I'm new to this forum. I did a mistake to me two months ago. I'm now suffering from -- telling her cancer story. She has for people who are responding to her and we're assuming that every member of this thread, both the person who started it and every other member are exposed to the amount of support that is on average in that thread. Except for their own particular message. Subject that out. So we assume that all five participants read all five messages. Clearly not the only people who are exposed. There's lots of lurkers who don't themselves post. We don't have any data about them. But this is data just on people who are active posters and a new data set we're working with the American Cancer Society, we'll have behavior as well so we can test this in more detail and the analysis is survival analysis where the event is staying in or dropping out of the group and not dying. and we've done this is a build. the biggest predictor of people sticking around is if there are parts of threads where there's more participation overall, there's more activity or there's more reading, then they are more likely to stick around. So that's comparing this line that is the second from the bottom that says, average message count of support. Somebody seeing atypical thread with a typical amount of support. And we compare that to this line above here, that is a high message count and average support. So after 90 days, if you are in threads that are a standard deviation, have a standard deviation more activity you go from about 15% likelihood of being in that group and sticking around to 35%. Doubling the likelihood of people sticking around if there's more activity in the group. And then at every level, the more emotional support that colors those messages, the more likely people are to stick around. So at the average support volume, you are 33% increase in survival. And at high support volume, if you get lots more emotional support, standard deviation more emotional support then you are 50% more likely [Indiscernible -- low volume] so emotional messages, more messages and more emotional support per message keeps people around. On the other hand, and maybe this is surprising, maybe not, getting information is associated with people leaving sooner than you would have expected. It could be that the information isn't very good. It could be like a telephone effect. Once you get the telephone number, some people may remember what telephone books are. Once you get a telephone number from a telephone book, it's not like a novel. You don't read the rest of it. You've got the information you want and so you can leave. and it could be that there's a difference in the kind of person who is seeking information and emotional support. People seeking emotional support might be more community oriented and are sticking around for that reason as well so we don't have a good causal story. We got some hypotheses. But getting emotional support keeps you around or is associated with staying around longer. Getting more informational support has a small but reliable effect of getting people to dropout more. We talked about this. And then the future work that we're doing is looking at measures of well-being. We've got with the American Cancer Society, longitudinal surveys where we asked people about their depression and health quality of life, self-efficacy over time and we can ask what were they are exposed to and what they do in these groups that are associated with changes in these psychosocial measures? and we've got the data back. We've got a new data set. We have to do the same machine learning in this new data set I described here in order to do this analysis. And then once we have data about what is in the groups that seems to be associated with improvements in health quality of life? We can try those nudging experiences -- experiments to make sure people are more likely to be exposed to things that in the correlational data is associated with improvements in health quality of life. Just for reasons of time, I will stop at this point. Just get to a place where you can find out more information. I can take questions or comments from the audience and then when that is quiet, we can go to the remote audience.

Thank you. At this time we will begin the question and answer session. If you would like to ask a question from the phone lines, please press star one. You will need to unmute your line and record your first and last name clearly when prompted. to withdraw your question, please press star 2. If you would like to ask a question from the phone lines, please press star one and record your name. One moment for our first question please.

I think you should press all of those buttons anyway.

The information and intelligent systems division. the idea occurred to me since number of responses and especially providing the emotional support since -- seems to be the driving Association. I'm not going to stay does not going to say causality. What comes to mind is something like Eliza. Would you provided all these nice emotional messages for using something like Eliza? Generating them? Might be easier to do than provide useful information.

So I'm going to extrapolate from other research that I've done, which is to look at what keeps people around when they join a Wikipedia project, which are these subgroups that are designed to curate particular topics. And some of the same phenomena occur there. So getting more people to talk to you, getting old-timers to talk to you and be supportive of you is -- increases your likelihood of sticking around. On the other hand if it's a template munication, it drives people away. So if they are perceiving or associated with people even more quickly. So even if it's a human, that takes a pre-formulated message and passes it on to one of these Wikipedia new editors, that person is more likely to leave than if they got no communication with existing members of the community. So based on that extrapolating, I would say Eliza would drive people away.

My name is Peter McLetchie. I'm visiting from the human competition astute. First of all thank you very much. I think it's really worthwhile, interesting questions you're asking and investigating. So what I noticed in general about observational studies is the challenge one has these indirect methods. And that when one uses indirect methods, there's sometimes the risk of analytic tautologies arising. And in this case what I wonder about in particular is in and assessing attend or intentionality among the original posters, that you are using the Amazon TRCers to do that. And the only difference between the Amazon TRCers that I can think of and the folks who are responding is that the Amazon TRCers are performing this task overly whereas the responders are doing it implicitly. in a social medium. So it leaves the question of, do we truly know the intentionality of the original posters? and is there a less objective way of assessing that? So what I was going to offer is maybe a different way of framing your analysis and wonder what you think about this which is to say that we can assess the intentionality of the poster by actually using the objective measure of their commitment to the system, over time? And their level of engagement with the system over time? and then the other related question is I remember you had a nice example of someone asking what your Amazon TRCers thought was an informational question but I saw a lot of embedded emotional words. I'm worried, concerned, I don't know what this means. And having participated in forums like this, myself, I recall some responders tend to respond almost always with emotional responses. And some almost always with informational responses. So I wonder if it would also be useful to do some kind of agent modeling to ferret out those kinds of effects?

We've done some -- are not in his environment. We've done some agent-based modeling in order to predict people's continued commitment of making assumptions that people have some stability in the kind of behavior in this environment. You can -- from modeling individual decisions -- decision-making about whether they are offering information or whether they are posting or not, you can get good estimates of population level parameters. The power law distribution of contribution, and average length of time people stick around and things like that. I think we could probably delve more deeply -- we have a much better model of what it is that people seem to be trying to get and how they are expressing those intense that we could do a much more detailed agent-based model of the participants in this group or in these kinds of groups as we study them much more and we can feel the details of the rules. So I think that's a way that we would like to go. Actually had a proposal rejected from the interdisciplinary social behavioral science one because a letter of collaboration was too long. But maybe somebody will want to find that to be able to do that agent-based model. I guess I didn't understand -- maybe we can talk afterwards,predicting support by engagement or predicting intention by engagement. I think we'll have to talk about that. I'm not sure there would be enough time.

Question to both. What instructions were given to the TRCers? In your first example, when they put informational content at 4.2,it was perfectly informational. Why it was so low at 4.2.

We gave TRCers essentially that phrase. How much informational support is in this message? and as some previous training examples, gave them examples of -- one example of something that experts judged as informational support. in general, the mean of the emotional support and the meaning of the informational support measures were different. So TRCers thought there was on average more emotional support in messages than informational support. I don't know if -- I don't really have a good answer to why a particular place -- what's the absolute value on the scale and whether that represents that TRCers were inaccurate in some, as God sees it, sense.

Related question. You talked about the agreement you got between learning techniques and the TRCers. What was the inter-TRC agreement?

So the average for any two TRCers at random or for any two experts in the small set of data we have for experts was around 0.8 as well. So --

About the same? Okay.

Exactly. Then we didn't have -- the measure we have for the agreement among humans is the average judgment. So when the intraclass correlation is how much signal areas in the average of the 10 judgments. It's not looking at pairs of the reliability of the scale associated with 10 judges.

Is there anybody in the remote audience?

Online? Any last questions?

We have no questions from the phone lines at this time.

Thank you again very much.

[Applause] and this does conclude today's conference. You may disconnect your audio light at this time. At 10 -- again, this does conclude today's conference. You may disconnect your audio lines at this time. Thank you.

[event concluded]