Social constructionism in security protocols

Simon N. Foley Vivien M. Rooney

This is joint work between myself and my co-author, Vivien Rooney. I'm a computer scientist, and Vivien's an applied psychologist. We're interested in understanding how humans experience working with security protocols. And, when I use the word "security protocol", I mean it in the most general sense: a set of rules that people and machines are supposed to follow.

As an overview of what we mean by social construction for security, let's start by looking at the Equifax data breach from 2016. The problem was that they had not updated their copy of the software package Struts, and as a consequence, a vulnerability in the software was exploited and there was a large data breach. When the breach was considered by the US House of Representatives committee on oversight and government reform, the Equifax CEO testified that "the human error was that the individual who's responsible for communicating in the organisation to apply the patch, did not." The individual involved testified that "To assert that a senior vice president in the organisation should be forwarding vulnerability alert information to developers three or four layers down in the organisation on every alert just doesn't hold water, doesn't make any sense. If that's the process that the company has to rely on, then that's a problem".

Here we have security controls in place and two different views. Whether it's their actual view or a contrived view, it suits my purposes as an example. One person's view was that it was human error and that it was this person's fault. The other view was that, "No, it's not my fault, it's stupid controls."

As technical people, we think, "I could observe this system. I could study the controls that are in place for software update. I could study what the humans are doing, and out of this, I would probably have found this vulnerability." But there's another dimension that goes beyond just looking at, or observing, what people are doing, and that is, how do we make sense of the humans themselves and how they're experiencing this technology. We might be tempted to look at these specific controls from the designer's point of view; their view of how the security works is a valid reality. And equally for the end-user, their view of what is happening is also a valid reality.

Let's look at an indepth example of this. Last year we conducted a study of people who worked as network defenders in Security Operations Centres (SOCs) and Security Incident Response Teams (CSIRTs). We wanted to understand how those individuals working as front line defenders were experiencing working within the system.

Within these systems there are many rules and procedures that the defenders are expected to follow. Among other things, we wanted to understand whether they actually follow these procedures. Are they experiencing things that, as individuals, might cause problems following procedures? I'll talk a bit more about the details in the second half of the talk. To give you a you a sense, here's an extract from a semi-structured interview that Vivien carried out with one of the defenders in the study. The question was, what happens when you're in a crisis situation and you see there's a possible attack or threat, and you want to find out more information?

There's rules that one is supposed to follow. Are they followed all the time? You can see in the transcript of the discussion with the interviewer, that the participant in the interview is coming to understand their interpretation of what they actually do on the ground. The participant says: "Oh everyday we're really aware that, we know if there's an issue, the quickest way is to call the person who is doing it, who's working an issue." They're talking about the issue, but then, that's not the procedure. It emerges that "I think in the organisation, the procedure is one thing, but maybe we're doing something different."

From the interview, "Well, this is informal", "Yes, and they encourage it","And is it acceptable?" And the answer is, "Well, most of the time, it's acceptable. There's a way to do it." And you can see they're sort of hedging it, so in a sense, they're coming to an understanding of what it is they're actually doing on the ground. I'll talk a little bit about the psychology behind this later and how that can affect those working in SOCs and CSIRTs, in our experience.

Through the interview the participant is constructing an understanding of their experience of working in the SOC or CSIRT. As they explain and describe their own experience with systems and technology, they're constructing its meaning to themselves. This does not come through checkboxes or prepared survey questionnaires, but through dialogue. In this way both the interviewer and the participant come to understand things that either might not have realised beforehand. It's jointly constructed understandings that form the basis of a shared reality of the participant. In this reality, most of the time they follow procedures, but sometimes they're not. And, of course, the meaning of their experience of working with this technology will change over time.

Frank Stajano Would it be too cynical to say it's acceptable not to follow the strict procedure, so long as nothing bad happens? But, if something bad happens, then it's your fault. That seems to be the case of when people say, "well, following procedure will take forever", so everybody knows that we can take shortcuts, but we know that we are not supposed to". And in fact, that goes to the point that, when people want to go on strike without going on strike, they would say "I will not strictly follow the procedure" and they're screwed just as badly.

Simon Foley That is the point. We learned that individuals who work in these situations have different social identities. That they have the identity of a

person who works in the organisation, someone who follows procedures, because that is good. There are good reasons why you need to follow procedures, even if it is just to cover yourself. But equally, they have an identity within their team and, what we saw in the interviews was that, within a team, people have a sense of camaraderie. There's a crisis. We need to defend against this crisis, and these procedures can become at a remove from the crisis: we know they're there, but we're not really following them right now. However, we then need to be mindful about how we report the crisis, because, if in my report I write "Frank didn't follow the procedure, and Frank contacted somebody outside the organisation" then that may get Frank into trouble, which might threaten the identity that I see myself having within this team.

There's another social identity which is that, as an individual defender, I have a network of contacts; I'm part of a global community of defenders and we have a common worthy goal; I identify myself with this community and this is important to me. What I'd like to be able to do is to phone one of my contacts in the community and say, "We're witnessing this traffic coming into our system. We think it's part of a threat. Are you seeing anything like this?".

There's definite procedures and guidelines in this case. For example, NIST Special Publication 800-150 recommends that there be a Memorandum of Understanding in place between the defender's organisation and that of their contact. Defenders are supposed to follow the MoU which precisely defines the kind of information that they are allowed share with their contact. But in a crisis, you're not going to be paying to much attention to that. Also, when you look at the defenders as individuals, they liken themselves to being firefighters defending, fighting the good fight as part of a global community. This is important to them and is part of this identity.

This is something that can influence whether they follow the rules. The defenders have different social identities "I'm a person who works for the organisation", "I'm a person who works within a team", "I'm a person who's part of a global community". And each identity may treat rules differently. Defenders construct and manage these different realities of themselves. This can result in cognitive dissonance and can produce psychological stress on the individual, which can then influence their performance within the group.

This was one result in our wider study and was one of the sources of tension that we found within SOCs and CSIRTs, this issue of having multiple social identities and the resulting cognitive dissonance. We could ask, on the one hand, "how can we help people who work in these situations deal with these issues". But then, on the other hand, sometimes you can't fix these issues. When we talk about security and when we talk about people following rules, the lawyers say, "You must follow these rules". But the reality of the situation is that as individuals, they have many more complicated things going on that they're having to deal with. So, maybe from the company's point of view, maybe they need to change what they mean by security.

Jovan Powar I agree. I found some of the thrust of what you're saying is that, if we're going to reason about these systems that we need to understand the constructive meaning of them, the shared reality of the people who actually use them. But, there's an experimental problem there, isn't there, because when you're doing this interview, you are guiding things, and they're coming to you, like you said, they're coming to the realisation of how systems work, but that is a joint reconstruction driven by what you've imprinted to the systems. So, if we're really trying to understand it, how do you get that constructive understanding from the people on their own that is useful.

Simon Foley There's two points here to consider. One is that I don't do these semi-structured interviews because I'm pretty sure that I would lead the person into technical explanations of things. My co-author is an Applied Psychologist. She conducted the interviews and as a Qualitative Researcher she is careful how she prepares for, and conducts, the semi-structured interviews, so as to avoid leading the person in their answers. The second point is that as a research method, this is social constructionism. The interpretation that the defender is forming is an interpretation from their shared experience of working. But equally, it's an interpretation based on the dialogue that they have with the person who's doing the interview. So, it is possible that you could have another interviewer of this person and they could come up with a slightly different theory. And that's the nature of social construction.

Jovan Powar When you're talking about their interaction with the security procedures, even asking them questions about it on its own might be a bias that you're introducing because if you're looking at. For example, a failure, and you're asking them about all the procedures. If they didn't have an understanding of the procedure then—if the reality was nil or very different—you're making them form an understanding of it at the point of interview, which is not the same as what they originally understood.

Simon Foley Let's go back to the transcript fragment from the interview. I can't recall the original question that lead to this discussion in the interview, but it would not have been an explicit "do you share information outside the organisation". It might be along the lines, "Okay, so you work in a SOC and your job is to deal with threats, what is that like?" and encourage them to expand on that. It is difficult to do this kind of interviewing without introducing bias, but it can be done. Nevertheless, you do make a good point.

That's the jointly constructed understanding and humans change their mind all the time. They reinterpret and come up with different explanations about what they're doing and what they think about them; they change their mind and what they do over time. Except of course for Vulcans who never change their mind. If security operations centre were run by Vulcans, then of course they'd follow procedures and fully understand their experience of it. What we

want to do is to build understandings of the human understanding of our human experience of security systems

We use Qualitative Research methods to do this. In particular, Vivien likes to use the Grounded Theory. It's socially constructed grounded theory, used to help to uncover what's happening among people. In this study it was cybersecurity defenders. Having used the method to identify what's happening, she studies the psychology of the human in the phenomena. What are they feeling? What are they thinking? What are they experiencing? Qualitative techniques are good for discovering unknown knowns. Finding out about the things you didn't realise you knew. We can map this activity into the model that Giampaolo Bella and Lizzy Coles-Kemp proposed in 2012. Our work on the human experience would map into the communal and personal levels.

We're using qualitative research as a means to uncover the human experience. This is not human behaviour nor are we trying to model the behaviour of these individuals. We're looking at what people's emotional responses are. How are they feeling about things, sensory awareness, physical environments, and situations[1]. Is it hot. Is it cold? Are they upset? Are they tired? Are they really annoyed about something? How does that influence their experience. Their desires, their choices, they're very ambitious. They want to do things. They want to do good and use their intellectual reasoning.

This is that what is happening in the security protocol, the security system, along with its users, is a social construction. It is by engaging with the individuals involved that we can come to an understanding of what's happening.

Michael Millian We have a very interesting anecdote from the hospital down the road that our department does some work with. For security of the computers that the doctors use, somebody put in some software that used the webcam on the computer to make sure there was a human in front of it. A human walks away and the screen locks after five seconds. But this interrupted the workflow of the doctors who would need to change their gloves every time they walked away from the computer. So, they discovered that if they put a large coffee cup over the webcam, the screen would just stay on. It's very interesting how often the people designing the security protocols don't interface with the users who are going to be using the security protocols. And then these two groups of people just end up being at odds with each other.

Simon Foley That's a good example of how people work around ill-considered design. There's lots of great examples of poor design in Don Normans book on the design of everyday things. While not specifically about design for computer security, it's about good user-centred design. User centred security promotes the idea that we should not blame the user; in a sense we should "blame" the designer because their design is ill-considered. Many times this is true. But I'm not so sure it should always apply. Read the transcript of the dialog with the

[1]This is nicely depicted in the painting *The Dance of Life* (1900) from Munch's project *The Frieze of Life*.

defender. I know "blame" is not a good word to use, but in a sense you might like to blame the defender in this case (about threat information sharing). It's not the designer's fault. Equally, it's not the defender's fault.

Jim Blythe I'd like to push back a bit on that. I don't think that security makes you always blame the designer and not the user. We should also be mindful of Angela Sasse's paper, the user is not the enemy. But what we have is a kind of mismatch between what the user expected and what the designer's got. And it might be one fault or another, the notion of fault here is a tricky one. That's why the notion of blame itself is tricky. And the other thing I'd like to point out about this example is that we're talking about defenders who are not your typical end-user. But I want to say that end users have entered into a social contract with the technology they're using, which is that, as we've said before, we shouldn't overburden them with what it's doing. These defenders are professionals who should have had some process which may not exist, which could be the organisation's fault.

Simon Foley My co-author would say, "You should not talk about blame." I'm trying to be contentious to push people to think about, the tension that goes on between these two views. And you're quite right that in this case, that defenders are not your conventional users, and we're not thinking about it in terms of typical end-user/consumer computer interaction.

Lydia Kraus I'm thinking that maybe it's not about blaming one side or the other side, and maybe it's a communication problem in the end. Perhaps the challenge is to find a constructive way to satisfy both sides.

Simon Foley I agree. I'm trying to capture human transgression in the security protocol. We may usually think of human transgression as error. Perhaps, we should re-think it: transgression becomes a normal part of security. At the moment I don't have a concrete answer about how this might be adequately modelled in the protocol. At what point is a tolerated transgression no longer tolerated?

We want to build a model of human experience so that we can then study how it interacts with the conventional models that we would use for describing a security protocol. At the moment I'm using a Bayesian Network to represent the relationship between attributes representing human experience and the system state and actions carried out by those humans. The Bayesian Network is elicited as part of the Grounded Theory analysis. Analysis may then provide some insight about the security of the system. Perhaps an increase in cognitive dissonance has become evident in the system: defenders are enacting multiple identities arising from a recent increase in external sharing. A remediation might be to help the individuals by providing them with training, or perhaps by strengthening team identity.

Our focus to date has been on identifying the phenomena and providing a psychological understanding. We haven't considered remediation to a great extent. In our study of the SOCs and CSIRTs we identified positive things and negative things about the experience of defenders; things that one could change and things that could not be changed; things that should change, things that should not change. With limited resources, you'd spend them on the things that you should change. The example in the paper is about how threat information is shared amongst cyber defenders.

Jonathan Anderson I think you see some cognitive dissonance in regulated professions. Engineers, doctors, lawyers, et cetera, where you've got these multiple identities that you're active in sometimes and make decisions. While analysts in a SOC might not exactly have the same kinds of credentials, I think that's true for the people who do critical things, like pilots and ship's captains. Is there anything that we can learn from those kind of environments that helps understand this better?

Simon Foley That's a good point. In some of the regulated professions, such as air-pilots, they tend not to have a culture of blame. When there is an incident, the culture is to be open and report and learn. They're not inclined to shoot the messenger. It would be interesting to explore what this culture is in the case of network defenders.

Jim Blythe You mentioned there were three components and one was a socio-technical, and that was really good. I noticed that the notion of teamwork comes in a lot from the description that you gave. But your focus on socio-technical is for individuals. I know that a number of psychological groups look at defenders, explicitly studying the teams and their construction. What you think about that?

Simon Foley It's not something that we had analysed in depth. While there is some mention of teams in the data, our focus has been more on the individual, and for example how they think about themselves and their relation to the team and social identity.

Fabio Massacci Do you consider how the strong time pressure affects these people working in the SOC? They get zillions of messages and and some point they have to take a decision, they are expected to, they have to do something.

Simon Foley In our study there were five different themes and the most interesting one was the theme around tensions. One of the issues within this we identified was the use of intuition by defenders. Within the SOC, they're not supposed to use their intuition; they are expected to follow procedures. However, sometimes they avoided following procedures as they slowed things

down; they use their intuition in order to get something done. Dealing with these multiple realities, the cognitive dissonance can result in stress.