Trustless or Trustworthy:

Value Change and User Engagement in Digital Security

Mohandas Schuyler Towne

Arizona State University

Abstract

This paper reviews the recent history of social science research into value change in communities, and uses those ideas to critique recent work in security engineering that seeks to engage users in the use of security and privacy technologies. The differences between "trustless" and human-centric security engineering are explored, and examples are provided of security research into modifying user engagement with privacy and security tools. Concepts such as asset mapping, social influence, and perceptual control theory are briefly explained and applied to example security regimes. In an attempt to bridge ideas from traditional information security research with ideas from social science, current scholarship on the security regime of the organizers of the Panama Papers is used as an example of successful technological and social influence on security behavior. Finally, recommendations for future research are made, with a particular focus on the open source Stethoscope tool released by Netflix's security team in early 2017.

Trustless or Trustworthy:

Value Change and User Engagement in Digital Security

The state of digital security education and planning is in flux. Many security developers have focused on creating technologies that route around any need for trust and tend to imagine the users of their technology as the weakest link in maintaining the security of their systems (McGregor, 2017, p. 505). However, recent research, advocacy, and advice on securing users has pushed back on that idea, suggesting that if users are engaging in insecure behaviors, it may be the result of the poorly designed and inconvenient technologies with which they are forced to engage (Jeong, 2017; Franceschi-Bicchierai, 2015; Towne, 2015). Perhaps the most interesting aspect of the renegotiation of terms between security engineers and users is the recent influence of sociology and activism on issues of digital privacy and security.

This paper will unpack recent thinking from technologists on how to effect change in user security behaviors, explore the recent history of social science research into value change in communities, and bridge those two ideas using current scholarship on the security regime of the organizers of the Panama Papers as an example of a technological and sociological triumph of security engineering. It will conclude with recommendations for future research and point out potential opportunities for sociologists to collaborate with security engineers.

What the Security Researchers Have to Say

With various motivations, from the threat of state surveillance to the difficulty of affecting user behavior, there has always been a community of security engineers and researchers that have focused explicitly on routing around the need to trust any individual or central authority. Without question, the technologies they have introduced to the world, including encrypted messaging (PGP, Signal), secure public ledgers (blockchain), and anonymizing networks (TOR, I2P), are invaluable tools for a wide range of end users, however, they are also notoriously difficult for many people to use. This is particularly frustrating in light of the large body of research that points to a very low engagement with security technology among the general population. (Alsaleh, Alomar, and Alarifi, 2017, pp. 1-2; Gkioulos, Wangen, Katsikas, Kavallieratos, and Kotzanikolaou, 2017, pp. 5-7; Ketelaar and vanBalen, 2018, p. 176).

Gkioulos, et al. (2017), conducted a large survey of three sample populations to establish a baseline of security awareness and behaviors of digital natives, defined by the authors as persons who grew up between 1987 and 1997 (Gkioulos, et al., 2017, p. 3). Theoretically, those persons grew up with the idea of perpetual connection to a network as normal (Gkioulos, et al., 2017, p. 1). Taken in aggregate, the responses paint a picture of users who value convenience, ignore the details of privacy controls, and are not overly concerned about the security of their devices. However, if we follow the arc of the sample set of information security students, defined by the authors as their "high competency" sample (Gkioulos, et al. 2017, p. 3), we find that the group the authors consider the best informed and most capable to teach and influence the behavior of others, only significantly outperformed the medium and low competency groups in a few areas, and dramatically underperformed in others. While they were well ahead when it came to password selection (Gkioulos, et al., 2017, p. 10.), in questions that involved using their devices in dangerous ways, such as jailbreaking, remaining perpetually logged in to services, and reporting the loss or theft of a device to authorities, the highest competency group were far more likely to engage in the riskiest behavior (Gkioulos, et al., 2017, pp. 7-8). Similarly, while the high competency group were more likely to use a screen lock on their mobile devices, they were also more likely to prefer biometric authentication, which is typically considered less secure than the less convenient long PIN (Gkioulos, et. Al., p. 9).

While the authors of the paper didn't offer a unifying reason for those behaviors, it is reasonable to imagine that while the high competency cohort may be more comfortable with the basic tools of security, that knowledge could also make them less confident about the efficacy of those tools, or perhaps simply more fatalistic about the nature of security breaches. Within the security community, the topic of breach fatigue is perennial, but this is just one flavor of what Manuel Castells, writing in *Communication Power* (2009), labeled "Scandal Fatigue" (pp. 253-254). Castells describes a situation in which any one scandal can raise alarm and alter the behavior of those potentially impacted by the news, but unending scandals inure the audience to each individual scandal, while simultaneously lowering their faith in the institution that births each new scandalous event. While there is little academic investigation of breach fatigue, there has been a great deal of popular and industry writing on the subject (Walker, 2015; Hu, 2014; Lystrup, 2016), with the general consensus being that it exists and that it is difficult to overcome the inertia of the twin attitudes of apathy and doubt.

In trying to both understand user behavior and come up with novel ways to move users to more secure behaviors, there has been some compelling work by security researchers that borrows from the social sciences. Foltz, Newkirk, and Schwager (2016) applied the theory of planned behavior to users making changes in the privacy settings of their social networks. Alsaleh et al. (2017), conducted a broad study into perceptions of security risks and possible methods of persuasion. Of particular note was their suggestion that being alerted to how their better-secured peers are engaging with privacy controls may help influence the target user to change (Alsaleh, et al. 2017, p. 28). Those security researchers that have been building tools to take advantage of this research have largely focused on peer pressure or scare tactics to influence users into better engaging with security tools.

Starting with the idea of fear appeals, as described by protection motivation theory, and noting that interactive "strength meters" were known to improve password quality when users were selecting passwords, Vance, Eargle, Ouimet, and Straub (2013) created a new password selection experience that used what they called "interactive fear appeals" (p. 2989). An interface analyzed the password chosen by a user and reported back to the user how many hours, days, or years it would take to crack their password. Vance, et al. found that by adding their fear appeal to what was effectively an interactive strength meter, they were able to improve average password complexity (2013, p. 2997). Keeping the interactivity, but opting for a peer-pressure approach, Facebook filed a patent in 2016 for a system that would identify privacy settings a target user hasn't turned on, then alert the target to users they are connected to who *have* made use of that security setting, in an effort to encourage the target user to do the same (Global IP News, 2016).

These are interesting approaches, but they are not yet moving the needle on secure user behaviors, though perhaps future surveys will reveal they are more effective than they currently appear. To better understand why these simple fear and peer pressure appeals don't appear to be making a significant impact, we can return to the idea of fatigue. Writing for Cisco's Umbrella Blog, Owen Lystrup (2016) goes a step further than his peers and suggests that we have moved beyond breach fatigue into breach acceptance (para. 1). That tracks both with the affect of the high competency users studied by Gkioulos et al., and in private conversations with well informed, competent, but ultimately jaded security technologists. There is a belief that breaches and leaks are inevitable, and that rather than protect against it, we should either gird ourselves for the pain, or change our behavior so that we don't use our technology for anything that we wouldn't want broadcast to the world. The latter proposition is particularly troubling if we are at all interested in having robust security in online spaces.

The Physical/Digital Security Divide

In 2013 at a dinner with several experts in information security the question of public reaction to security breaches came up. Starting with the premise that one of the core principles that keeps us safe in our private physical spaces is collective moral outrage at the transgression of those spaces, the guests tried to imagine what it would take to finally create that same collective emotional response to transgressions against our digital spaces (personal communication, June 1st, 2013). Unsurprisingly, the most compelling answers came from medical device researchers who suggested that the first malicious, pre-meditated medical device hacking that led to a death with be the trigger for social change. No one there guessed that the moment of transformative change would come from the mass release of tens-of-thousands of private photographs of both celebrities and non-celebrities just a year later.

Alice Marwick (2016) framed what was alternately called *CelebGate* or *The Fappening*, a massive leaked archive of nude or semi-nude photos of celebrities, as an issue of gendered privacy, which she defines as "a type of privacy violation that is more likely to happen to one gender due to structural inequality" (p. 178). Within that framework, she describes non-consensual sharing of sexual images as a form of sexual violence, and connected them to the mass photo leak to examine the role of gender in privacy discourse (Marwick, 2016, p. 179). Marwick analyzed active reddit users' responses to the leaks, pulling out several major themes, from plain unfiltered desire for more, to practical concerns about the appearance of child pornography limiting further leaks, and then moralizing justifications that placed blame on celebrity or female promiscuity. (2016, pp. 183-186). Unfortunately, that all reads as par for the course for those comfortable with victim blaming both in sexual violence, and information security. However, Marwick also discusses the prevailing counter-narrative in the popular press that the victims of that crime were, in fact, *victims of a crime* (2016, p. 187).

That narrative was solidified by Jennifer Lawrence, one of the victims of the leaks, who used the phrase "sex crime" early and often when speaking publicly about the event. Though she may not have been the first to use those words, her celebrity and willingness to speak publicly and furiously about the events enshrined the idea of the leaks as a sex crime in the public discourse. The explicit framing of even the viewing of the leaked photos as a sex crime immediately moved into other domains. Surprisingly savvy commentaries like James Kosur's efficient description of the event as "...a sex crime. More specifically, the Fappening is dozens of sex crimes, committed en masse, with a socially engaged audience." (2014, para. 1), were unrelenting in their approach. The public attention, framing, and social pressure even seemed to have a very direct impact on the individuals who were viewing and sharing the images on Reddit. Writing for the Daily Beast (2014), Marlow Stern revisited Reddit's sub-communities related to the Fappening. While Marwick's paper focused on the immediate discourse in the week following the initial releases (2014, p. 181), Stern followed up on the conversation about a month later and found several users who had enthusiastically engaged in the sharing and viewing of nude photos of women now apparently quite honestly engaging with the idea that they may, in fact, have been perpetuating their victimization (2014). Further, they in turn attempted to engage their community with those same ideas.

In her analysis of the leaks, Adrienne Massanari (2017) framed her critique on the idea of a "toxic technoculture", defined as networked public communities who assemble around a particular topic of interest, but whose defining feature includes the harassment of non-members of their community (p. 333). She pointed out that Reddit's corporate interests appeared to explicitly outweigh any moral imperative (2017, p. 363), a critique affirmed by Marwick (2016, p. 177). Massanari used the example of Reddit administrators' willingness to put extra work and human capital into maintaining access to the website during a massive uptick in users accessing the celebrity nude photos on the "Fappening" subreddit, receiving much higher user payments during that window, and only shutting down access to the subreddit a week later, when forced to confront the existence of underage photos in the leak archives (2017, p. 336). Massanari made the case that Reddit's technology, moderation, and lack of uniform action nurtured some truly abhorrent communal behaviors (pp. 336-341).

Yet, the public discourse still remained primarily focused on the criminal and victim aspects of the event, and even some of those users participating in gendered privacy violations inculcated within the toxic technoculture that Reddit profited by, precisely the most unlikely targets for social transformation, publicly accepted some responsibility and engaged their peers in publicly reckoning with their role in perpetuating a sex crime (Stern, 2014).

This reframing of public opinion is not only an important step toward harmonizing security and privacy perceptions of the physical world with those of digital spaces, but it also provides the groundwork for considering future efforts to exert similar positive influence. Nicole Ozer (2012) laid out the broad strokes and future direction of the privacy movement, and Lee Ann Banasazak and Heather Ondercin (2016) used their study of the women's movement to challenge other social scientists to study changes in public opinion as an outcome of social movements.

Ozer's work quite neatly predicted an event like the Fappening as a watershed in the privacy movement. She pointed to early observations that the privacy movement needed to better spread knowledge of basic concepts of data collection and connect those concepts to individual's lives to establish broad concern about encroaching privacy violations (Ozer, 2012, pp. 217-219). Ozer covered some factors that had previously limited the establishment of a social movement around privacy, such as simply under-valuing privacy and a general lack of technical sophistication (2012, pp. 220-230). She gave similar treatment to emerging trends (circa 2012) in the privacy social movement, examining the influence of technology-focused media, the technical savvy of both commentators and audiences, and new ideas of personal branding (Ozer, 2012, pp. 231-250).

Public perceptions aren't universal, however, and nailing down exactly what the final result of big public privacy events like the Fappening may be is difficult at best. While public discourse was certainly moved in a positive direction, the last few months of 2014 saw other major leaks, including several that were targeted at non-celebrities. A follow-on event referred to as *The Snappening*, in which approximately 90,000 images of snapchat users leaked onto the open internet, further complicated the dynamics of digital privacy discourse (Brandom, 2014). A shift in attitudes about whether it is ok to endlessly re-victimize persons whose private photographs are distributed non-consensually may not do much to offset the recurring apathy and doubt that accompanies breach fatigue.

Ruminating on these ideas in *Alone Together: Why we expect more from technology and less from each other* (2011), Sherry Turkle describes the paralyzing nature of the uncertainty around the "rules" of privacy in a networked society, and the certainty that someday our data will get out there (pp. 253-254). She describes a younger generation that are still concerned about privacy, but who do not know what is acceptable, or where the line between an innocent prank and a crime might be. While the generation above them are very publicly failing to sort all of this out, they have been left in limbo. With no clear social guidelines, few comprehensible or evenly enforced legal guidelines, and an unending barrage of news stories about major data breaches, the fatigue, acceptance, or paralysis around security behaviors may have reached its nadir.

Unfortunately, there is strong evidence to suggest that privacy concerns, and an inability to trust the tools that are provided to mitigate those concerns, dramatically impacts our willingness to connect to other people online and reduces group cohesion (Pan, Wan, Fan, Liu, & Archer, 2017, pp. 152-153). That reveals the issues of trust and trustlessness in digital security as an existential problem. *If people cannot trust the technology that mediates their connection to one another, they are much less likely to connect to one another at all.*

Pan, et al. do offer us a possible way forward, however. They found that increased group trust leads to reduced privacy concerns and increased social cohesion (Pan, et al. 2017, pp. 158-159), and pointed out that there has been little attention paid to what they call "collaborative privacy" (Pan, et al. 2017, p. 153), which is effectively the ability to trust that those persons with whom you are connected will value your privacy and security, and to be trusted by them to do the same. In this way people are transcending or subverting trustless security architectures, that fittingly have lost the trust of most users, in order to establish direct relationships of trust with one another.

What the Sociologists Have to Say

What Pan, et al. referred to as collaborative privacy appears to be a form of what Castells would describe as setting values in a network. Essentially, if you want to alter a target audience's behavior, you need to encode a new value in the network within which that audience locates itself (Castells, 2009, p. 29). In a small group of friends who connect online, this could just be that each of them values their own privacy. If enough of them make it clear that privacy is a shared value then they may collectively value privacy and will both engage with privacy tools and behaviours that support the privacy of their peers, and encourage others within their network to do the same. While this may seem too small to effect massive change, Castells also explains that what is valued in one network will shape the values of all sub-networks (2009, p. 27). If we go a step above our previous example and look at the platform that the friends use to connect with one another, we might find that privacy isn't a central value of the larger network of the platform, and thus, the tools available to the group to maintain their privacy are either non-existent, difficult to use, or ineffective. Even if members of the group individually cared about their privacy, the larger network of the social platform they are on makes it too difficult to engage with privacy, and thus imposs or ineffective to establish privacy as a value within their sub-network.

Castells frames establishing the values of a network as an essential form of power, and insofar as it can shape the behavior of a network’s individual members, his framework aligns with some core ideas in identity theory. Writing in their seminal book *Identity Theory* (2009), Peter Burke and Jan Stets describe the perceptual control model, which states that "people control their perceptions, not their behaviors" (p. 29). That means that people will take action to return their perceptions to a reference point. Burke and Stets use the example of a person being cold (2009, p. 29). Their perception is that they are cold, and so they take an action, perhaps turning up the heat, in order to return their perception to the reference point of a comfortable temperature. Identity theory also locates the self as both individual and social in its nature. We know ourselves as referenced against other people, so our self-perception is inherently informed by our group perception (Burke, et al. 2009, p. 10). Thus, many of the reference points for our perceptions are informed by the norms of our society.

Taken together, identity theory, the perceptual control model, and the power of defining value in a network, gives us a more effective lens than peer pressure or fear with which to affect change in a community. Castells also offers us a lever with which to influence the network. He makes it clear that his theory does not seek to identify individuals who wield social power, but instead to understand how that power is created through the influences of multiple actors, and networks of actors. There are two recent studies that illustrate how this can be done, and how that information can be used to create change.

Asset Mapping and Intervention

Susan Jakes, Annie Hardison-Moody, Sarah Bowen, and John Blevins (2015) used *asset mapping* as a tool to investigate two communities in North Carolina, with a particular focus on access to healthy food and safe spaces to be active. Asset mapping is a community-dialog centered approach to surface valuable persons and resources extant within a community, rather than beginning a process of change from a perspective of the community's needs and deficits (Jakes, et al. 2015, p. 392). However, as modified by these researchers, asset mapping can also be used to identify the shared values of a community (Jakes, et al. 2015, pp. 392-394). The results of their study suggest that conducting the asset mapping exercise can provide the information necessary to create impactful change that reflects the values of the community.

The other study focuses on influencing a small random selection of members of a network, and uses them to seed value change throughout their networks. Elizabeth Levy Paluck, Hana Shepherd, and Peter M. Aronow (2016) conducted an experiment to influence the individual behaviors of randomly selected students, then study whether their new behavior had an influence on their social networks, and ultimately, if that influence reached the total population of their schools. Their work brought together three forms of behavior change and provided a framework for how one might create mass change from individualized interventions (Paluck et al. 2016, p. 566). The authors focused on conflict avoidance, with prior research that suggested students accept high levels of conflict not because they are comfortable with it, but because it is perceived as acceptable or even typical (Paluck et al. 2016, p. 567). Paluck et al. focused on changing the perceptions of individual students to find conflict less acceptable and to be comfortable expressing that opinion publicly (2016, p. 567). The results suggested that their interventions may have accounted for a 30% reduction in reported conflict behaviors at the treated schools (Paluck et al. 2016, p. 369).

Trust as a Network Value

In her critique of the current vogue of the *blockchain*—distributed databases that act as reliable public ledgers—Rachel O'Dwyer (2016) expressed concern with their current uncritical adoption and development (p. 1). She was particularly suspicious of the idea that we can replace social processes with interesting technology (O'Dwyer, 2016, pp. 2). That idea rises from the blockchain's stated goal of being "trustless", as in, by using the blockchain, we are able to conduct exchange, whether commercial, informational, or even social, without any need to actually trust each other, or any organizing or governing social/political body (O'Dwyer, 2016, p. 2), a prospect that O'Dwyer finds dubious.

The idea to route around the need for trust is deeply engrained in information security which has adopted President Ronald Reagan's favorite Russian proverb, "trust, but verify", as a central tenet. A quick google search for "trust, but verify" in slides from talks at information security conferences turns up thousands of results. Some researchers, however, have chafed at trust getting top billing, and have proposed alternates like "Verify and never trust" (Banafa, 2014), or "Don't Trust, Verify", which was the slogan of the 2015 Scaling Bitcoin conference (YourBTCC, 2015).

In the introduction to *Alone Together* (2009) Turkle tells her audience that what drew her to MIT was that she was noticing a change in language (p. ix). Terminology from computing had started to permeate the culture, and provided people with not just new words, but new ideas with which to reassess themselves. Turkle had studied a similar shift in the cultural use of terminology from the psychoanalytical tradition to reflect on the self. With that in mind, it is important to really think about what it means to encode the idea of trustlessness into our language and how that may bleed into our values.

While the "Zero Trust" model has become the gold standard for many security researchers (Banafa, 2014), there was a time when security was sought through the active development of community. Josh Boyd (2002) established that the core problem in early eCommerce was trust, both between a user and a website and between users and each other, and made clear that an eCommerce company's main job was to give users reasons to both *trust and be trustworthy* (p. 2). Boyd detailed how risk and trust came together for eBay's users and described eBay's solution to perceived risk as a rhetorical framing of the eBay community of users (2002, p. 6). Despite the demonstrable success of eBay's community-centric security model, their security regime eventually followed the trajectory of their competitors. Boyd warned that external tools, such as insurance and escrow schemes would ultimately reduce trust between users, as they were explicitly implemented in order to obviate the need for trust (2002, p. 12).

Where Trust, Community, and Security Meet

McGregor, Watkins, Al-Ameen, Caine, and Roesner (2017) wrote about the remarkable, and remarkably successful security regime employed by the organizers of the Panama Papers, including the perceptions and behavior of the otherwise unaffiliated investigative reporters who were given access to the documents. The authors detailed the security infrastructure of the project as organized by the International Consortium of Investigative Journalists (ICIJ), both the community-building aspects and the technical tools created for the journalists (pp. 505-506). Via surveys of the participating journalists and interviews with project leaders, the authors investigated the high level of successful security engagement. Much attention was paid to the tools created by ICIJ technologists (McGregor, et al. 2017, pp. 506, 509-512), but particularly compelling were the attitudes of the participants. The reporter-participants found the use of encrypted messengers, email, etc. to be relatively easy to get started with, and even easier to maintain (McGregor, et al. 2017, pp. 506, 512-513). The authors pointed to the sense of purpose and community within the participant group as a significant contributor toward the perception of ease in dealing with traditionally complex technologies.

We can look to the perceptual control model and value setting in a network for a possible answer to why the reporters involved in the Panama Papers self-reported an ease of use of security technologies that are routinely criticized for their difficulty. As part of that cohort, and with both the technical, and rhetorical guidance of the ICIJ, the perceptual reference point for the importance and ease of using security software was set very high. Additionally, knowing the reporters would be collaborating with one another and that any individual could greatly decrease the efficacy of the reporting of the whole by leaking information, the organizers actively, and regularly reasserted the core values of both security and community (McGregor, et al. 2017, p. 506). They cultivated group cohesion and trust. In this way, the ICIJ introduced its cohort of reporters to communication tools designed to be trustless, but maintained their efficacy and use by establishing trust and trust-worthiness as core values within their network of reporters and technical workers.

Opportunities for Future Research

There are people and organizations who are currently reorienting their security practice and development around respect for the actual behavior of their users. Netflix's development of Stethoscope is an interesting example. Stethoscope is a tool used internally at Netflix to track the status of devices employees connect to their internal network. Stethoscope scans the devices for important security features, alerts users to any issues, and provides them actionable options for fixing those issues. However, Stethoscope does not force users to take those actions. This platform is interesting both because they have had some internal success with Stethoscope and user-respecting rhetorical frameworks, but they have also open sourced their technology. If other organizations see value in Stethoscope and choose to implement it, there may be many interesting research opportunities with myriad networked communities. Even more exciting is that the developers point out that the tools they create reflect the values of the people who designed them (Kriss, J. & White, A. 2017), thus, to use Stethoscope most effectively, organizations may be led to a position of trusting their users.

Whether thinking about Stethoscope particularly, or other security regimes designed to work with, and trust users, there are interesting opportunities for a researcher to modify, contribute to, and ultimately study the interaction of members of the network with each other and the technology. Possible modifications could include conducting an asset mapping study with the creators or implementers of the technology to identify current physical, digital, and personnel assets, and to get a sense of the current values of the network. The development team could also work on influencing individual users who could act as advocates not only for use of the technology, but also for the values that underlie the technology. While Stethoscope tracks longitudinal trends in security engagement, with the addition of user surveys or interviews, a social scientist could additionally track the longitudinal maintenance of the goal values of the network.

Conclusion

A generation of security developers insisting that we need to find ways to route around the need for trust has affected not only security design, but may also have affected our culture as a whole. Security regimes that leverage high quality technological solutions while also encoding trust and trustworthiness as core values may find greater success than regimes which focus solely on trustlessness. It is unlikely that ubiquitous digital security can come from technology alone, but it may be found at the intersection of technological and social change. That does not mean that we do not need robust security technologies. In fact, they are perhaps even more important in light of the social consequences discussed in this paper. With the emotional fatigue and social disengagement that comes from a barrage of constant privacy and security breaches, strong, easily accessible and usable security tools will provide the secure spaces within which the more delicate, iterative processes of value change can occur.

References

Alsaleh, M., Alomar, N., & Alarifi, A. (2017). Smartphone users: Understanding how security mechanisms are perceived and new persuasive methods. *PLoS ONE, 12*(3). doi:10.1371/journal.pone.0173284.

Banafa, A. (2014). *Verify and never trust: The Zero Trust model of information security.* [Slides] Retrieved from: https://www.slideshare.net/professorbanafa/zero-trust-33999348

Banaszak, L., & Ondercin, H. (2016). Public opinion as a movement outcome: The case of the U.S. women's movement. *Mobilization: An International Quarterly, 21*(3), 361–378. doi:10.17813/1086-671X-21-3-361.

Boyd, J. (2002). In community we trust: Online security communication at eBay. *Journal of Computer-Mediated Communication, 7*(3). doi:10.111/j.1083-6101.2002.tb00147.x.

Brandom, R. (2014). *A third-party Snapchat client has leaked tens of thousands of user photos.* [Blog Post]. Retrieved from https://www.theverge.com/2014/10/10/6956725/third-party-snapchat-app-leaks-13gb-of-user-photos

Burke, P. & Stets, J. (2009). *Identity Theory.* Oxford, United Kingdom: Oxford University Press.

Business Insights: Global (2016, May 19th). *U.S. patent and trademark office releases Facebook's patent application for sytems and methods for increasing security sensitivity based on social influence.* Global IP News.

Castells, M. (2009). *Communication Power.* Oxford, United Kingdom: Oxford University Press.

Foltz, C., Newkirk, H., & Schwager, P. (2016). An empirical investigation of factors that influence individual behavior toward changing social networking security settings.

Journal of Theoretical and Applied Electronic Commerce Research 11(2), 1–15. doi:10.4067/S0718-18762016000200002.

Franceschi-Bicchierai, L. (2015, September 2nd). *Even the inventor of PGP doesn't use PGP: Security is hard.* [Blog post]. Retrieved from https://motherboard.vice.com/en_us/article/vvbw9a/even-the-inventor-of-pgp-doesnt-use-pgp#

Gkioulos, V., Wangen, G., Katsikas, S., Kavallieratos, G., & Kotzanikolaou, P. (2017) Security awareness of the digital natives. *Information, 8*(42). doi:10.3390/info8020042.

Hu, E. (2014). *I feel nothing: The home depot hack and data breach fatigue.* [Blog post]. Retrieved from: https://www.npr.org/sections/alltechconsidered/2014/09/03/345539074/i-feel-nothing-the-home-depot-hack-and-data-breach-fatigue

Jakes, S., Hardison-Moody, A., Bowen, S., & Blevins, J. (2015). Engaging community change: The critical role of values in asset mapping. *Community Development, 46*(4), 392–406. doi:10.1080/15575330.2015.1064146.

Ketelaar, P., & van Balen, M. (2018). The smartphone as your follower: The role of smartphone literacy in the relation between privacy concerns, attitude and behaviour towards phone-embedded tracking. *Computers in Human Behavior, 78*, 174–182. Advanced online publication. doi:10.1016/j.chb.2017.09.034.

Kosur, J. (2014, September 2nd). *The Fappening: When sex crimes become easily consumable entertainment.* [Blog Post]. Retrived from https://www.business2community.com/social-buzz/fappening-sex-crimes-become-easily-consumable-entertainment-0995479#s2BJjoC86Q4VevKA.97

Kriss, J. & White, A. (2017, February 21st). *Introducing Netflix Stethoscope*. [Blog post]. Retrieved from https://medium.com/netflix-techblog/introducing-netflix-stethoscope-5f3c392368e3

Lystrup, O. (2016). *From breach fatigue to breach acceptance.* [Blog post]. Retrieved from: https://umbrella.cisco.com/blog/2016/04/14/from-breach-fatigue-to-breach-acceptance/

Marwick, A. (2016). Scandal or sex crime? Ethical implications of the celebrity nude photo leaks. *Ethics and Information Technology, 19*(3), 177–191. doi:10.1007/s10676-017-9431-7

Massanari, A. (2017). #Gamergate and The Fappening: How Reddit's algorithm, governance, and culture support toxic technocultures. *new media & society, 19*(3), 329–346. doi:10.1177/1461444815608807.

McGregor, S. Watkins, E. A., Al-Ameen, M. N., Caine, K., & Roesner, F. (2017). *When the weakest link is strong: Secure collaboration in the case of the panama papers.* Submitted for publication.

O'Dwyer, R. (2016). *Blockchains and their Pitfalls*. Manuscript submitted for publication. Retrieved from: https://www.academia.edu/23524276/Blockchains_and_Their_Pitfalls

Ozer, N. (2012). Putting online privacy above the fold: Building social movement and creating corporate change. *New York University Review of Law & Social Change, 36,* 215–282.

Paluck, E., Shepherd, H., & Aronow, P. (2016). Changing climates of conflict: A social network experiment in 56 schools. *Proceedings of the National Academy of Sciences of the United States Of America, 113*(3), 566–571. doi:10.7910/DVN/29199.

Pan, Y., Wan, Y., Fan, J., Liu, B., & Archer, N. (2017). Raising the cohesion and vitality of online communities by reducing privacy concerns. *International Journal of Electronic Commerce, 21*(2), 151–183. doi:10.1080/10864415.2016.1234281.

sarahjeong. (2017, November 22nd). *I've now done enough stories where PGP keys feature as one aspect of identity verification to where I can say for certain: ...everyone is bad at PGP.* [Twitter post]. Retrieved from https://twitter.com/sarahjeong/status/933427775632846848

Stern, M. (2014, October 4th). '*The Fappening' perpetuators have a J. Law come-to-Jesus moment and 'cower with shame'*. [Blog Post]. Retrieved from https://www.thedailybeast.com/the-fappening-perpetuators-have-a-jlaw-come-to-jesus-moment-and-cower-with-shame

Towne, M. (2015, June). *Selling security in a post-lock society.* Talk presented at RVASec 2015, Richmond, VA.

Turkle, S. (2011). *Alone Together: Why we expect more from technology and less from each other.* New York City, NY: Basic Books.

Vance, A., Eargle, D., Ouimet, K., & Straub, D. (2013). Enhancing password security through interactive fear appeals: A web-based field experiment. *46th Hawaii International Conference on System Sciences*, 2988–2997. doi:10.1109/HICSS.2013.196.

Walker, Z. (2015). *Data breach fatigue, have we already reached our limit?* [Blog post]. Retrieved from: http://info.rippleshot.com/blog/data-breach-fatigue

YourBTCC. (2015, December 5th). *We're at #ScalingBitcoin and the talks have started. Don't trust. Verify*. [Twitter post] Retrieved from https://twitter.com/yourbtcc/status/673351271588458496