Corporate Transparency and Consumer Trust in Household Digital Assistants

A Research Paper submitted to the Department of Engineering and Society

Presented to the Faculty of the School of Engineering and Applied Science

University of Virginia • Charlottesville, Virginia

In Partial Fulfillment of the Requirements for the Degree

Bachelor of Science, School of Engineering

Audrey Swart

Spring 2024

On my honor as a University Student, I have neither given nor received unauthorized aid on this assignment as defined by the Honor Guidelines for Thesis-Related Assignments

Advisor

Caitlin D. Wylie, Department of Engineering and Society

Introduction

Have you ever had a conversation about something and then seen it in an online advertisement? You might have thought it was just a coincidence, but you also might have wondered if something more was going on: was your privacy violated? If you're part of the 35% of American households that own at least one smart speaker or digital assistant device (National Public Media, 2022), it could be the culprit. It's impossible to know everything about how our data is processed by companies because they are not required or inclined to provide that information to the public. They are, however, inclined to provide transparency in other ways that serve their image.

This investigation of trust and transparency in the context of household digital assistants, also known as smart speakers, will be based on the theoretical framework known as the Social Construction of Technology, or SCOT. Under SCOT, it is thought that human action shapes technology (Pinch & Bijker, 1987), in contrast to technological determinism which states the opposite (STS Infrastructures). I will specifically draw from Klein and Kleinman's critique of Bijker's original work and their addition to the theory of SCOT that a technology cannot be understood without understanding how that technology is embedded in its social context (Klein & Kleinman, 2002, p. 34). My goal is to argue that big technology companies use selective transparency as a tactic to gain consumer trust. The issues I will discuss are relevant to multiple companies, such as Apple and Google, but this analysis will have a focus on Amazon and their line of "Echo" smart speaker products enabled with "Alexa" virtual assistant technology. If we can better understand why we place trust in digital assistants, we will be better able to interact with companies in a way that protects our privacy and allows us to be more fully informed about how our data is used.

Transparency

As consumers, we tend to feel like the more transparent something is, the more trustworthy it is. If the operating procedures of an organization or the inner workings of a piece of their technology are made visible to the public, and they are determined to be sound, it's generally assumed that people then place trust in that organization and are justified in doing so. Others disagree. Mike Ananny, a professor of communication at USC, and Kate Crawford, an internationally acclaimed researcher in the realm of technology and society, argue in their 2018 paper that one of the limitations of transparency as an ideal is that it does not necessarily build trust. I would agree with them, and I would also go a step further to say that it does not necessarily *warrant* trust, either. Simply being able to look inside the black box of a technology is not enough to capture the nuances of how that technology is already established in the world and what relationships it has with various social groups (Ananny & Crawford, 2018, p. 982).

The information that is made transparent and available also has to be comprehensible by the people to which it is presented. This creates an extra layer of processing in which experts have to describe information in a simplified way that leaves out the more complex elements. In this regard, I think Thi Nguyen, a philosophy professor at the University of Utah specializing in how technology shapes social values, is correct in his 2021 paper when he states that experts might tailor their explanations to be suitable for public consumption to the point where their presented justifications are different from their actual justifications. In the context of the topic of this paper though, my agreement with Nyugen ends there. He goes on to argue that this demand for transparency can actually undermine the expertise of professionals, becoming a force that affects their actions rather than only their justifications, and that this phenomenon is a form of surveillance that can cause harm (Nguyen, 2021, p. 340-342). However, if we define surveillance as the power to enact change through continuous observation (Ananny & Crawford, 2018, p. 975), then in reality, in the context of large technology corporations such Apple, Google, and Amazon, the power dynamics in place in society protect these entities from being truly surveilled by the public. A powerful company will be able to control what information they make available and the way in which it is presented in order to minimize negative perceptions and consequences. This disconnect from power is another limitation of transparency described by Ananny and Crawford, as “the power to shame is ineffective against those with the power to endure visibility” (p. 984).

Transparency can be broadly defined as the sharing of information. In the realm of digital assistants, it can be thought of as going in two directions between two relevant social groups involved in the construction of this technology, the producer and the consumer (Klein & Kleinman, 2002, pp. 40-41). These groups have different levels of power. Transparency from the producer to the consumer entails providing information about how the product works, what the company’s motivations are, and in what ways consumer information is used. The level of transparency is controlled by the producer. Transparency from the consumer to the producer, while it would be controlled by the consumer in other contexts, is *also* controlled by the producer when it comes to digital assistants due to the nature of the product. By design, the digital assistant is processing user data every time it is activated to carry out a task or answer a question. The amount of data and the exact extent to which it is used are unknowable to those outside the producer’s organization. According to Amazon, the only data that is processed by Alexa is what the consumer says directly following the activation phrase, and the only way that it is processed is to carry out the task at hand (and occasionally to anonymously improve the service) (Amazon, n.d.). But this information is coming from the producer. For example, user data from Alexa interactions are supposedly stored until the user wants to delete them, but it is currently unknown how long the data actually remain recoverable on the servers.

The motivations of a corporation are not always as pure as they may lead the consumer to believe. On the surface, the motivation of Amazon's release of Alexa was to improve people's lives and provide an intuitive way for people to interact with technology, but the deeper motivating factor, as with most corporations, is profit. And profit depends on the collection of user data; Amazon and similar companies have indeed been known to collect user data to personalize advertisements (Neville, 2020). Amazon's website (Amazon, n.d.) is reasonably transparent about how the physical technology works, but not so much in regard to their broader motives or how the public's data is really used. In this way, the level of transparency present from the consumer to producer is controlled by the producer, and selective transparency is used to foster trust.

Trust and Anthropomorphism

There are two possible ways to define trust in the context of this paper. One would be to say that trust is the act of making oneself vulnerable to others (Feltman, 2009). The other is to say that trust means putting faith in someone to create a desirable outcome (Foehr & Germelmann, 2020), or in other words, believing that bad things will not happen. Depending on which definition of trust is used, discussing a consumer's trust in a company like Amazon implies different feelings from the consumer. If they are agreeing to make themselves vulnerable, then they know Amazon might use their data for company purposes. If they are under the impression that this will not happen, that is a very different kind of trust. I will be focusing on the latter for this investigation.

It is common practice for the manufacturers of technologies like digital assistants to provide documentation to consumers with clear statements about how the product works, largely to ease potential fear from the public about data misuse. For example, Amazon's website (Amazon, n.d.) discusses how the physical technology and processing algorithms behind Alexa work together to complete consumer requests. They also emphasize that there are tangible cues (such as buttons, lights, and sounds) built into Alexa devices that indicate to users when the device is and is not recording data. The aspect of giving users the ability to see and hear when a smart speaker is active helps to create a sense of control over the technology. This is a concept explored by Foehr and Germelmann, two scholars in the field of marketing. Their research dives more deeply into the concept of trust and how it is formed by consumers in relation to smart voice-interaction technologies. They found that anthropomorphization increases feelings of trust, and that consumers are more likely to anthropomorphize smart speaker technology when there is an absence of clear cues to determine when the device is active (Foehr & Germelmann, 2020, p. 183). Since there *are* clear cues with smart speaker devices from Amazon, as well as with those from many other brands, this would mean that users are more likely to view them simply as pieces of technology rather than human-like assistants.

However, the other factor coming into play is the producer's inherent anthropomorphization of smart speakers. This is through their names, voices, natural language programmed responses, and even embedded in the phrase "wake word" (p. 183). Research has shown that female voices are generally perceived to be warmer than male voices (Stern, 2017). Additionally, many believe that digital assistants use female voices in order to make the service come across as helpful rather than commanding (Grattan, 2016) and so people feel more control over the technology. Being able to refer to a device by a name and then hearing a female voice in a conversational tone facilitates users assigning human-like traits to their smart speakers, thus increasing feelings of trust. Foehr and Germelmann briefly mention that companies have the potential to use anthropomorphism to their advantage for data collection, but they fall short of more deeply considering the implications and the reasons behind a company's choice to build a product with these features.

The question arises then of whether or not Amazon intentionally uses anthropomorphization as a tool along with their selective transparency. Fetterolf and Hertog, researchers in the fields of technology and social science, seem to think the answer is yes. They imply that Amazon is aware that the power imbalance between them and the public lends itself to wariness and dislike from the public, and Amazon therefore actively encourages users to think of Alexa as "human-like and relatable" (Fetterolf & Hertog, 2023, p. 11). Their study found that Alexa users tend to manage distrust in Amazon by "separating the [voice assistant] from the company through anthropomorphism" (p. 1). People assign their anxieties to the broad company entity of Amazon rather than the specific product of the smart speaker, and therefore feel comfortable enough to continue to use said product (p. 4). This encourages Amazon to continue to develop and emphasize the anthropomorphic qualities of Alexa.

Data Use and Targeted Advertising

Amazon's website presents the following frequently asked question: "*I'm not talking to Alexa and am having a conversation at home near my device. Is Alexa still listening and recording everything I say?*" The answer is as follows: "*No. Alexa is a part of your life only when you ask Alexa to be. By default, Alexa begins listening after your Echo device detects the wake word, so Alexa does not listen to your personal conversations*" (Amazon, n.d.). This is a statement that is seemingly transparent with consumers, reassuring them that Alexa only listens to statements that immediately follow the wake word, but it is actually a false statement. Until the microphone detects the wake word, it is "in an inert state… allowing the microphone to passively 'listen' for a key word without recording or transmitting information" (Gray, 2016, p. 6). So while it's true that Alexa is not always *recording* everything a consumer says, it is false that Alexa is not always *listening*. Amazon's website does not adequately address this distinction, nor do they ever refer to their smart speakers as "always on" devices (Gray, 2016). This would be the most concise and accurate phrase to use but Amazon avoids it because it tends to create unease. Instead, they explain that "the device detects the wake word by identifying acoustic patterns that match the wake word" (Amazon, n.d.), while skipping over the fact that the process allowing it to identify the acoustic patterns is constant passive listening.

Additionally, there have been known instances of smart speaker devices having false positive activations, meaning they start recording data when nobody has said the activation phrase. Amazon's website mentions this but does not quantify any related data, instead choosing to simply use the word "infrequent." In 2022, cybersecurity researchers at Murray State University carried out an experiment over the course of eight weeks in an apartment with three roommates. In the apartment, they placed four Amazon smart speaker devices which were all different generations of the Echo. They found that there were 225 false positive activations (Combs et al, 2022). Dividing this number by the amount of time and the amount of devices, that is an average of approximately one false positive per device per day. While this could be considered infrequent, it is undeniably quite significant and concerning from a privacy perspective.

A paper published by computer science researchers from UC Davis uncovered that it is very likely that Amazon uses data specifically from Alexa devices to infer consumer interests and show them targeted ads online (Iqbal et al, 2023). The researchers simulated personas with different interests interacting with smart speakers and came to their conclusions based on the presence of "statistically significant differences in the online targeted advertising" between the different types of personas (p. 570). The study also found that Amazon's data practices are not clearly disclosed to users, stating that at the time of their research they did not find any information from Amazon on Alexa Echo interaction data for ad targeting (pp. 569, 579). Interestingly, one researcher pointed out that in between the paper's preprint in 2022 and official release in 2023, Amazon updated their privacy statement to include that information (Heath, 2023; Iqbal et al, 2023, p. 579). I would argue that the information is still not adequately explicit. Upon searching through nearly a dozen links to different pages on the Amazon website, users can find a page that says Amazon shows you interest-based ads that use information from your "interactions with Amazon sites, content, or services." On a completely different page, they state that you provide them information when you "talk to or otherwise interact with our Alexa Voice service" (Amazon, n.d.). It is only through piecing together these two snippets of information that consumers can realize that they are being shown targeted ads based on what they say to their Alexa devices. This statement is never seen unfragmented anywhere on the site. By avoiding emphasis on this statement and thus avoiding the perception that using consumer data is a core aspect of the company, Amazon fosters consumer trust. This serves to demonstrate the truth and importance of the argument that transparency does not necessarily warrant trust; even though Amazon is transparent about how their technological devices work and what they do, there are more hidden factors at play.

The implication of not explicitly presenting the entire truth to customers in this way is the potential for large-scale and long-term privacy violations. These types of violations are technically legal due to the lack of federal privacy laws in the US (Klosowski, 2021). Amazon recognizes the absence of congressional action and states that they support local privacy laws. However, Amazon and other big tech companies are specially positioned within this issue because they have the power to lobby against bills they do not want, and they have indeed used this power. According to the journalism site Reuters, Amazon has lobbied to undermine and kill several proposed privacy protection bills across the US over the years, and has declined to comment on why they opposed the bills (Dastin et al., 2021).

Conclusion

Transparency is endorsed a multitude of times by Amazon on their website (Amazon, n.d.). By focusing on that ideal, they aim to create an image of a trustworthy company. They then aim to support that image with evidence by providing simplified descriptions of their services and products and allowing people to peer inside the technological black box of a smart speaker device. However, the level of transparency here is highly selective and limited to only the information about which Amazon desires to be transparent, an advantage that is created by the power dynamic between the public as consumers and big technology companies as producers. Amazon's selective transparency is a tactic to gain consumer trust, and it has arguably been quite effective on consumers. People feel like Amazon's current level of transparency warrants their trust, but they don't understand all the ways in which the technology is embedded in society, i.e., its social construction as described in the SCOT framework. This combines with the ease of forming trust in anthropomorphized products to foster trust and to allow Amazon to discreetly use consumer data from smart speaker devices to target advertisements.

It is my position that Amazon needs to be fully transparent rather than selectively transparent in regards to all aspects of their smart speaker technology; not only how it works but also how the collected consumer data is truly used. The US has a need for more strict and more comprehensive data privacy and transparency laws that apply to companies like Amazon, and in order to facilitate this, we need stronger regulations and limitations on corporate lobbying. This would decrease the ability for technology companies to mislead consumers. Until then, it is my hope that this analysis helps to provide an avenue through which we as consumers can think critically about our use of household digital assistants, what it means to place our trust in them and their producers, and how we can take part in more informed interactions with technology companies.

References

Amazon. (n.d.). Alexa, Echo Devices, and Your Privacy. *Amazon.com.* https://www.amazon.com/gp/help/customer/display.html?nodeId=GVP69FUJ48X9DK8V

Amazon. (n.d.). How Alexa works: Wake word. *Amazon.com*. https://www.amazon.com/b/?node=23608571011&tag=googhydr-20&hvadid=661712030093&hvpos=&hvnetw=g&hvrand=18392446379828946748&hvpone=&hvptwo=&hvqmt=e&hvdev=c&hvdvcmdl=&hvlocint=&hvlocphy=9008337&hvtargid=kwd-2086027308707&ref=pd_sl_4gozlak3z2_e&gclid=Cj0KCQjwwYSwBhDcARIsAOyL0fjm5Ko64iLf-Y_osCX1LU1MYFaQxLOb3yf2Kt8qEYZdcxOnYoGnfhMaAgOnEALw_wcB

Amazon. (n.d.). Interest-Based Ads. *Amazon.com*. https://www.amazon.com/gp/help/customer/display.html?nodeId=GLVB9XDF9M8MU7UZ

Amazon. (n.d.). Amazon.com Privacy Notice. *Amazon.com*. https://www.amazon.com/gp/help/customer/display.html?nodeId=GX7NJQ4ZB8MHFRNJ

Ananny, M. & Crawford, K. (2018). Seeing without knowing: Limitations of the transparency ideal and its application to algorithmic accountability. *New Media & Society*, *20*(3), 973-989. https://doi.org/10.1177/1461444816676645

Combs, M., Hazelwood, C., & Joyce, R. (2022). Are you listening? – an observational wake word privacy study. *Organizational Cybersecurity Journal: Practice, Process and People, 2*(2), 113-123. https://doi.org/10.1108/OCJ-12-2021-0036

Dastin, J., Kirkham, C., & Kalra, A. (2021). Special Report: The Amazon lobbyists who kill U.S. consumer privacy protections. *Reuters*.

https://www.reuters.com/legal/litigation/amazon-lobbyists-who-kill-us-consumer-privacy-protections-2021-11-19/

Feltman, C. (2009). The thin book of trust. Thin Book Pub. Co.

Fetterolf, E. & Hertog, E. (2023). It’s not her fault: Trust through anthropomorphism among young adult Amazon Alexa users. *Convergence, 0*(0), 1-18. https://doi.org/10.1177/13548565231200337

Foehr, J. & Germelmann, C. (2020). Alexa, Can I Trust You? Exploring Consumer Paths to Trust in Smart Voice-Interaction Technologies. *The Association for Consumer Research, 5*(2), 129-237. https://doi.org/10.1086/707731

Grattan, Debbie. (2016). 6 Reasons People Trust a Female Voice Over Male Voices. *Debbie Grattan Voiceover Talent*. https://www.debbiegrattan.com/blog/why-trust-female-voice-over-male-voice/

Gray, S. (2016). Always On: Privacy Implications of Microphone-Enabled Devices. *Future of Privacy Forum*. https://fpf.org/wp-content/uploads/2016/04/FPF_Always_On_WP.pdf

Heath, J. (2023). Study Shows Alexa Invades Privacy, Collects User Data for Ad-Targeting. *UC Davis College of Engineering*. https://engineering.ucdavis.edu/news/study-shows-alexa-invades-privacy-collects-user-data-ad-targeting

Iqbal, U., Bahrami, P., Trimananda, R., Cui, H., Gamero-Garrido, A., Dubois, D., Choffnes, D., Markopoulou, A., Roesner, F., & Shafiq, Z. (2023). Tracking, Profiling, and Ad Targeting in the Alexa Echo Smart Speaker Ecosystem. *Proceedings of the 2023 ACM Internet Measurement Conference (IMC ’23), October 24–26*. 1-15. https://doi.org/10.1145/3618257.3624803

Klosowski, T. (2021). The State Of Consumer Data Privacy Laws In The Us (And Why It Matters). *The New York Times*. https://www.nytimes.com/wirecutter/blog/state-of-privacy-laws-in-us/

Klein, H. & Kleinman, D. (2002). The Social Construction of Technology: Structural Considerations. *Science, Technology, & Human Values, 27*(1), 28-52. https://www.jstor.org/stable/690274

National Public Media. (2022). The Smart Audio Report. *National Public Media*. https://www.nationalpublicmedia.com/insights/reports/smart-audio-report/

Neville, S. (2020). Eavesmining: A Critical Audit of the Amazon Echo and Alexa Conditions of Use. *Surveillance & Society, 18*(3), 343-356. https://doi.org/10.24908/ss.v18i3.13426

Nguyen, T. (2021). Transparency is Surveillance. *Philosophy and Phenomenological Research, 105*(2), 331-361. https://doi.org/10.1111/phpr.12823

Stern, J. (2017). Alexa, Siri, Cortana: The problem with all-female digital assistants. *The Wall Street Journal*. https://www.wsj.com/articles/alexa-siri-cortana-the-problem-with-all-female-digital-assistants-1487709068

STS Infrastructures. (n.d.). SCOT. https://stsinfrastructures.org/content/scot#

Vaccaro, A. & Madsen, P. (2009). Corporate dynamic transparency: the new ICT-driven ethics? *Ethics and Information Technology, 11*(2), 113-122. https://doi.org/10.1007/s10676-009-9190-1

Wojton, H., Porter, D., Lane, S., Bieber, C., & Madhavan, P. (2020). Initial Validation Of The Trust Of Automated Systems Test (TOAST). *The Journal Of Social Psychology, 160*(6), 735-750. https://doi.org/10.1080/00224545.2020.1749020