Lauren Holt

AI Pavillion Seminar

Professor Evans

10 December 2018

How Much Do Smart Speakers Really Hear and Who is Listening?

As we witness technology's rapid advancements, it seems that soon nearly everything with a plug or battery will be able to respond to a voice command. Always-on microphones are implemented in many popular tech products today, allowing humans to interact with virtual voice assistants, such as Amazon's Alexa and Google's Assistant. When a user verbalizes a keyword or phrase, known as the wake word, the device is activated and "begins" listening for further instruction (Kruzel). It doesn't take long to realize that the microphone must have been listening in the background, waiting to pick up on the vocalized "Alexa" or "Hey, Google." Questions are raised around whether the background conversation is being recorded and saved somewhere in the cloud, and who actually has access to the recordings' data. Although voice assistant devices are typically viewed as convenient aids for consumers, it must be questioned if these devices are being used as the trojan horse for data collection and government surveillance. This piece will discuss the privacy concerns of always-on microphones, beginning with official privacy statements from Amazon and Google and then moving into the government's potential access to personal data recorded from always-on microphone devices. It will close by raising concern of false positive wake words and discussing their legal relevance.

Amazon and Google's Privacy Statements

Dominating 70% of the smart speaker market, Amazon has high standards to live up to in the eye of the consumer (Koetsier). Having an effective privacy policy and terms of conditions is a crucial method to ensure trust among the company and the consumer. With a quick Internet search, Amazon's policies are clearly posted, but upon closer examination, statements dealing with data sharing become rather opaque. Google's privacy policy and terms of conditions are laid out similarly. The transparent aspects of both company's policies will first be discussed.

Recording and Storage of the Data

As previously mentioned, smart speakers are designed to activate in the presence of a wake word. The FAQ sections of both Amazon and Google's websites outline that although the smart devices are said to be "always listening," they are *technically* only listening when they are prompted, as the background recordings are on an automatically deleting loop, resetting every few seconds until they identify the proper acoustics of the wake word ("Data Security & Privacy on Google Home"). The companies assure users that no personal audio leaves the device until proper initiation.

Once the user begins to interact with a smart speaker device, the data is uploaded and stored in "the cloud," a server residing in the company's data center. Google and Amazon both ensure that the data is automatically encrypted to enhance its security in the server. According to the Google Home Help Center, the conversation data remains saved until one manually deletes it, which can be done by searching through in the settings tab of one's account. Amazon's data storing functions similarly, where recordings can be both viewed and manually deleted through account settings.

Enhancing User Personalization

As it is now understood that the conversation data is indefinitely stored in the cloud, it is still not completely clear how the data is used and who it is shared with. This is where company privacy policies become more vague. Amazon states that voice recordings are simply used "to answer your questions, fulfill your requests, and improve your experience and services"("Alexa Device FAQs"). The generation of unique voice profiles, which allow Alexa to identify who is talking to her, is an example of how the analysis of stored user data contributes to personalization. Google has a very clearly laid out privacy policy, affirming the user that collected data is used to make "services faster, smarter, and more useful" ("Data Security & Privacy on Google Home"). Beyond direct smart speaker personalization, the data contributes to maximizing a user's Google experience as a whole, explaining why an advertisement related to an exchange previously made with one's Google Home might appear in their web browser. In regards to marketing, conversation with a smart speaker can be thought of like an Internet search, where websites collect data to build a user profile and improve target advertisements.

Third Party Data Access

In addition to using the data within the company to improve the product and user experience, a certain amount is shared with external third parties. Amazon's privacy policy explains that third parties "have access to personal information needed to perform their functions, but may not use it for other purposes." For example, if one requested an Uber through their smart speaker, their location and other data relevant to the service would be sent to the party. The definition of "relevant data" varies among services and is therefore difficult to quantify. Google makes a clear statement that they do not sell personally identifiable user data to

external companies, though just as Amazon does, they do share information with associated third parties when necessary.

Government Data Access

When it comes to legal matters, Google and Amazon's data sharing policies become more ambiguous and situation based. Both companies receive regular requests from governments and courts to disclose private data, though they generally try to avoid doing so unless the data release is crucial to the case. Google's privacy policy expresses that personal data can be shared outside of the company when there is "reasonable belief" that it is necessary to fulfill "any applicable law, regulation, legal process, or enforceable governmental request." They explain how their legal team handles the decision process addressing government requests for data, making sure that the requests satisfy both legal requirements and company policies. As of January 2018, they fulfilled an average of 67% of legal requests for data that they received ("Google Transparency Report").

Amazon is known to be more strict with releasing consumer data, prioritizing the privacy of their users. An Amazon Spokesman stated that they "will not release customer information without a valid and binding legal demand properly served" (Flynn). This is exemplified in recent news cases where Amazon has denied government requests to access to a murder suspect's Amazon Echo recordings. It is likely that Amazon does not want to set a precedent where suspicion of criminal activity commensurates with the government's right to private data from their devices. It is difficult to determine whether the presence of a smart speaker device in a suspect's home is enough to claim probable cause for further investigation of corresponding audio recordings. If Amazon Echo data was easily accessible and admissible in courts, Amazon

would be put at a marketing disadvantage, as their devices would be viewed as witnesses and the company would be viewed as snitches (Ferguson). It should also be noted that it is unlikely that the conversations containing relevant legal evidence are among the user and smart speaker, and thus would not have been recorded and saved by the device in the first place.

Government Policy of User Data Acquisition

Though Amazon and Google both have policies in place to handle attempted government intrusions of user data, there is only so much control they have as a third party in legal situations. Government policies that regulate user data sharing will now be discussed. The historical context of the legislation surrounding personal data will first be reviewed and then current principles and policies will be examined.

The Extension of the Fourth Amendment through Katz v. US (1967)

The most notable ruling surrounding the government's access to personal data was a result of the 1967 Supreme Court case Katz v. the United States. In regards to the Fourth Constitutional Amendment, protection from unreasonable search and seizure, the case discussed the definition of "reasonable search." Federal agents attached a listening device to a public phone booth to eavesdrop on Katz, based solely off of their suspicion that he was engaging in illegal activity. Katz was convicted based on the recordings, but appealed the case, claiming that it was a violation of the Fourth Amendment and his right to a reasonable expectation of privacy. It was then ruled that the unwarranted wiretapping of a phone was an equivalent invasion of privacy as physical intrusion, and would therefore require a warrant ("Katz v. United States"). This case set a precedent for the consideration of reasonable privacy that would later become relevant to digitized data.

The Third Party Doctrine

Though a reasonable expectation of personal privacy is a strong value of the US Constitution, the right was limited within the decade following the Katz case. Both the US v. Miller (1976) and Smith v. Maryland (1979) cases established that a user should have no expectation of privacy for information voluntarily shared with a third party (Kerr). Under the Third Party Doctrine, Fourth Amendment rights are only applicable to the third party that holds the user information, such as Amazon and Google, rather than to the user themselves. Though the Third Party Doctrine obliterated a user's Fourth Amendment rights to their data, the privacy rights did not disappear entirely. Instead they became a dominion of the third party that holds the data. If the government wants to acquire personal information about a user, they must deal with the third party on a legal level (Kerr).

The Electronic Communications Privacy Act (ECPA)

With the responsibility to protect an abundance of user data, many companies have established their own privacy policies, such as those previously mentioned of Amazon and Google. Though Amazon and Google have policies in place to handle legal requests for data and to prevent unnecessary government data seizing, the presence of a subpoena, court order, or search warrant can still force a company to abandon their policies and forfeit data, as stated by the Electronic Communications Privacy Act (ECPA) of 1986 ("Google Transparency Report"). The ECPA was enacted to extend the legal protection that prevented the wiretapping phones to also protect digitally transmitted computer data. Though in concordance with the Third Party Doctrine, it is not the user that is protected from unreasonable search and seizure, but the company to which they granted their data.

As time advances, the Fourth Amendment and protection of privacy is becoming more and more limited in the digital world. Users are forced to forfeit personal data to third parties by simply interacting with common technologies like cellphones and search engines, even when carrying out the most mundane tasks. Most consumers are not even aware of the degree of data that they are sharing. Due to the fact that third parties now hold an unprecedented amount of user data, there is an ongoing push to overrule the Third Party Doctrine and to regain privacy rights for the consumer (Solove). With the increased prevalence of smart devices in homes, particularly smart speakers, the legalities surrounding user privacy must be revisited and updated. There needs to be more regulation of what defines "voluntary data" if users are at risk of losing all privacy rights when forfeiting it to a third party.

False Positive Wake Words

What is perhaps most concerning regarding whether smart speaker data is voluntarily shared or not stems from the possibility of false positive wake words. As the name implies, a false positive wake word is when a device misinterprets a sound for a wake word, and therefore begins streaming an involuntary recording. A paper from the USENIX Security Symposium delved into research on the accuracy of Alexa's speech interpretations and found that only 68.9% of the words in their tested data set were correctly identified by the device (Kumar, Deepak, et al. 36). Though the words with the highest misinterpretation rates were generally homophones, words that sound alike but have different spellings and meanings (such as "cell" and "sell"), the research implies that there is some degree of uncertainty in the accuracy of a smart speaker's identification of a wake word. Additionally, the phonetic structure of how a user pronounces a

word can also lead to confusion of the devices audio transcriptions, for example the word “wet” is commonly misinterpreted as “what.”

Earlier this year in Portland, Oregon, a woman’s Amazon Echo recorded and sent her private conversation to one of her husband’s colleagues. When the case was brought to Amazon for investigation they ruled that,

> The Echo woke up due to a word in background conversation sounding like ‘Alexa,’ Then, the subsequent conversation was heard as a ‘send message’ request. At which point, Alexa said out loud ‘To whom?’ At which point, the background conversation was interpreted as a name in the customer’s contact list. Alexa then asked out loud, ‘[contact name], right?’ Alexa then interpreted background conversation as ‘right.’(Chokshi)

The sequence of events might seem to be a rare occurrence, and the Echo device did seek clarification of the misinterpreted commands multiple times, but it was not the first report of its kind. As the use of smart speaker devices continues to multiply, such events will become more common. This raises concerns about the security of the smart speaker and the data it captures.

Additionally, there have been user reports where their Amazon Echo is falsely activated by a Lexus commercial playing on the television in the background. Other users have said that words like “Alaska” cause the their Amazon Echo to light up. Berkeley researchers even found that they could embed wake words and commands into music, where the sounds went unnoticed by the human listener, but were picked up by both the Amazon Echo and Google Home (Smith). These findings open new doors of concern where malicious audio files, whether music on the radio or advertisements on the television, could be used to make unauthorized purchases and unlock doors through smart speakers.

Wayne Kurtzman, Research Director at International Data Corporation (IDC), speaks on the wake words of the Google Home, “Hey, Google” and “OK, Google," stating that, “using the two word wake system causes significantly fewer false 'listenings' than Alexa,” (Koetsier). Though this two word activation method *reduces* accidental device arousal, researchers have found phonetically similar phrases to activate the device, such as “cocaine noodles,” which is frequently misinterpreted as “Hey, Google” (Smith). Despite the seemingly higher level of accuracy, Google Home users have still reported times where they were in conversation and realized that their device was wrongfully recording.

Implications of False Positive Wake Words

If a smart speaker’s data was seized for legal use but was not actually a product of a user’s intentional interaction with the device, it would be a violation of the Third Party Doctrine. It could be argued that the user still has privacy rights over the data, as they did not voluntarily initiate the streaming of the recordings. Not only would this create issues within Google and Amazon about acquiring unauthorized data, but it would create major privacy invasion disputes on a legal level, should the data be seized for government use. As controversies rise over the ownership of digital data, there are not sufficient legal procedures in place to rightfully handle them. Legislation must be updated in correspondence to new technologies, keeping the right to reasonable privacy central.

Paul Rosenzweig, a Law Professor at George Washington University Law School warns users that Google Home and other other digital home assistant devices should be thought of as "home wiretapping devices, designed to collect every bit of information about you” (The Federalist Society). The occurrence of inadvertent recordings of private conversations through

false positive wake words supports his claim and raises concern around what personal data consumers are blindly handing off to third parties like Amazon and Google, and potentially even the government. The presence of home smart speakers weaken consumers' Fourth Amendment rights and puts virtually all of their personal information at risk to government scrutiny.

Conclusion

Following extensive research of smart speakers and their implications, one can surmise that although recording private conversations is against both the policies of Amazon and Google, the aforementioned cases prove that technological flaws cause it to remain an issue. Though companies typically oppose legal interactions, the government still holds authority over the user data and can access it through legal measures. Juniper Research predicts that by 2021 more than fifty percent of American households will have at least one smart speaker (Smith). With a society that is so drastically shifting towards automation, consumers are blindly forfeiting Constitutional rights to powerful companies and legal bodies, leading to the demise of privacy for the individual.

There are solutions to the digital data dispute, but they require the voice of consumers to fight for their privacy and security rights within the law. This must begin with public awareness of the matter. Users must be more conscious of the capabilities of the devices that they install in their private homes and must decide if what seems to be a simple convenience is worth the privacy risk. Since avoiding the use of smart technology is not feasible for society as a whole, laws and regulations are the only things that will shape what the data can be used for; whether it is primarily for maximizing the consumer experience or potentially intruding on consumer privacy.

Works Cited

"Alexa and Alexa Device FAQs." *Amazon*, Amazon, www.amazon.com/gp/help/customer/display.html?nodeId=201602230.

"Amazon Privacy Notice." *Amazon*, Amazon, www.amazon.com/gp/help/customer/display.html?nodeId=201909010&pop-up=1.

Chokshi, Niraj. "Is Alexa Listening? Amazon Echo Sent Out Recording of Couple's Conversation." *The New York Times,* The New York Times, 25 May 2018, www.nytimes.com/2018/05/25/business/amazon-alexa-conversation-shared-echo.html.

"Data Security & Privacy on Google Home." *Google*, Google, support.google.com/googlehome/answer/7072285?hl=en.

"Google Privacy & Terms." Google, Google, policies.google.com/privacy#footnote-legal.

"Google Transparency Report." *Google Transparency Report*, Google, transparencyreport.google.com/user-data/overview.

"Katz v. United States." Oyez, 8 Dec. 2018, www.oyez.org/cases/1967/35.

Kerr, Orin. "Third Party Rights and the Carpenter Cell-Site Case." *The Washington Post*, WP Company, 15 June 2017, www.washingtonpost.com/news/volokh-conspiracy/wp/2017/06/15/third-party-rights-and-the-carpenter-cell-site-case/?noredirect=on&utm_term=.d5f155e2b180.

Koetsier, John. "Amazon Echo, Google Home Installed Base Hits 50 Million; Apple Has 6% Market Share, Report Says." *Forbes*, Forbes Magazine, 2 Aug. 2018, www.forbes.com/sites/johnkoetsier/2018/08/02/amazon-echo-google-home-installed-base-hits-50-million-apple-has-6-market-share-report-says/#382ec920769c.

Kruzel, John. “Is Your Amazon Alexa Spying on You?” *Politifact*, 31 May 2018, www.politifact.com/truth-o-meter/statements/2018/may/31/ro-khanna/your-amazon-alexa-spying-you/.

Kumar, Deepak, et al. *Skill Squatting Attacks on Amazon Alexa*. 2018, p. 36, *Skill Squatting Attacks on Amazon Alexa.*

Smith, Craig S. “Alexa and Siri Can Hear This Hidden Command. You Can't.” *The New York Times,* The New York Times, 10 May 2018, www.nytimes.com/2018/05/10/technology/alexa-siri-hidden-command-audio-attacks.html?module=inline.

Solove, Daniel. “10 Reasons Why the Fourth Amendment Third Party Doctrine Should Be Overruled in Carpenter v. US.” *TeachPrivacy*, 14 Dec. 2017, teachprivacy.com/carpenter-v-us-10-reasons-fourth-amendment-third-party-doctrine-overruled/.

The Federalist Society. YouTube, YouTube, 13 Nov. 2017, www.youtube.com/watch?v=HV9q3PXqGPY.