ITIF | INFORMATION TECHNOLOGY & INNOVATION FOUNDATION

Briefing By:

Daniel Castro
Vice President
Information Technology and Innovation Foundation

To the:

U.S. Senate AI Insight Forum

Hearing on:

"AI: Privacy & Liability"

November 8, 2023
Kennedy Caucus Room, Russell Senate Office Building
Washington, DC

INTRODUCTION AND SUMMARY

Leader Schumer, Senators Rounds, Heinrich, and Young, and distinguished members of the Senate, thank you for the opportunity to discuss privacy and liability issues related to artificial intelligence (AI). I am Daniel Castro, vice president of the Information Technology and Innovation Foundation (ITIF), a technology policy think tank, as well as director of ITIF's Center for Data Innovation.

AI AND PRIVACY

Regarding privacy, the development and deployment of AI involves the collection and use of vast amounts of data, including personally identifiable information (PII). Some of the potential privacy implications of AI include:

- Data Breaches: AI systems could expose PII to unauthorized individuals in a data breach. Users may share PII with AI chatbots, such as those providing legal, financial, or health services. For example, data breaches of transcripts of a conversation with an AI therapist or an AI girlfriend could reveal sensitive information.
- Data Leaks: AI may reveal PII included in training data. Affected individuals may have little control over the inclusion of this information. For example, their PII may be scraped from public websites and then used to train an AI model. Alternatively, an AI model may leak data if it trains on confidential user data, such as private contracts or medical records.
- Surveillance: AI makes it easier to analyze large volumes of data, including about individuals, which enables surveillance. Employers can use AI tools to closely monitor their employees, such as to track their keystrokes or how much time they spend at their desks. Similarly, governments can track individuals in public spaces, such as through facial recognition technology.
- Inferences: AI can infer information about people's identities, habits, beliefs, preferences, and medical conditions, including information that individuals may not know themselves, based on other data about those individuals. For example, AI systems can detect rare genetic conditions from an image of a child's face.[1]
- Deepfakes: AI makes it easier to create fake images, audio, and videos of individuals to harass them, harm their personal and professional reputations, and use fake media to conduct identity theft. In one recent case, students at a New Jersey high school allegedly used AI image generators to produce fake nude images of their female classmates.[2]

This list of privacy implications may appear alarming, but policymakers should remember that inflated fears about new technologies are quite common. Indeed, ITIF has documented the "technology panic cycle" (see Figure 1) often associated with novel technologies, especially the pattern of fear about potential privacy harms.[3] Concerns about a new technology grow rapidly at the beginning as critics make exaggerated and misleading claims about the likely risks, but eventually, public concerns decline over time when most of the predicted harms fail to materialize. In addition, concerns fade as courts clarify how existing laws and regulations apply, businesses address public feedback, social norms evolve, and consumers reconsider the tradeoffs once they understand better the benefits of the new technology. Concerns about AI, especially generative AI, are quickly reaching the height of hysteria. Policymakers should be careful not to overreact to

fears about how AI will impact privacy and rush to regulate the technology since history shows many of today's concerns will fade over time.[4]

Figure 1: Technology Panic Cycle

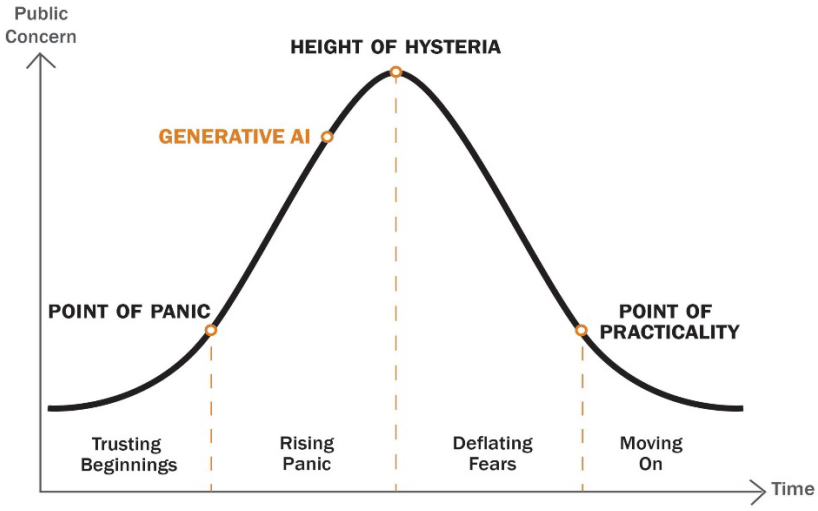

It is also important to note that AI does not present fundamentally new privacy concerns, although it may exacerbate some existing ones. Therefore, rather than pursuing AI-specific solutions, policymakers should seek to address broader privacy questions. For example, data breaches have been an unfortunate, yet regular, occurrence for the past two decades. Last year, there were nearly 1,800 data breaches in the United States impacting hundreds of millions of Americans.[5] Policymakers should address the larger problem of data breaches rather than focusing exclusively on data breaches involving AI systems.

Privacy concerns about surveillance and inference are similarly long-standing issues. In terms of surveillance, many employers regularly monitor their employees in the workplace, such as to track their attendance and Internet activity, and law enforcement uses technology like automatic license plate readers and surveillance cameras to maintain public safety. Many businesses use statistical techniques to infer information about individuals. For example, online advertising platforms may infer information about users, such as predicting their age or political leanings based on their online activity.[6] Privacy laws to address surveillance and inferred data can and should apply regardless of whether AI is involved.

Data leaks were an early concern about search engines, as attackers could use them to discover a trove of sensitive data, such as credit card information, Social Security numbers, and passwords, that were scattered across the Internet, often without the affected individuals' awareness.[7] However, individuals and organizations have since become more conscientious about what PII they publish on the Internet, and consumers can now use services like Google's "Results About You" tool to proactively monitor and remove PII from search engines.[8] Both with search engines and AI, the real privacy issue is that certain PII is publicly available on the Internet, not that there are tools that can find the information.

Finally, manipulated media, such as "Photoshopped" images, have a long history that predates recent advances in AI.[9] Indeed, many people do not realize that one of the most famous photographs of Abraham Lincoln is a composite image of Lincoln's head superimposed on an earlier photograph of John C. Calhoun

(see Figure 2).[10] The harm that can result from false and misleading media does not depend on whether those who created it used AI or other tools.

Figure 2: Photographs of President Abraham Lincoln (left) and Vice President John C. Calhoun (right).[11]

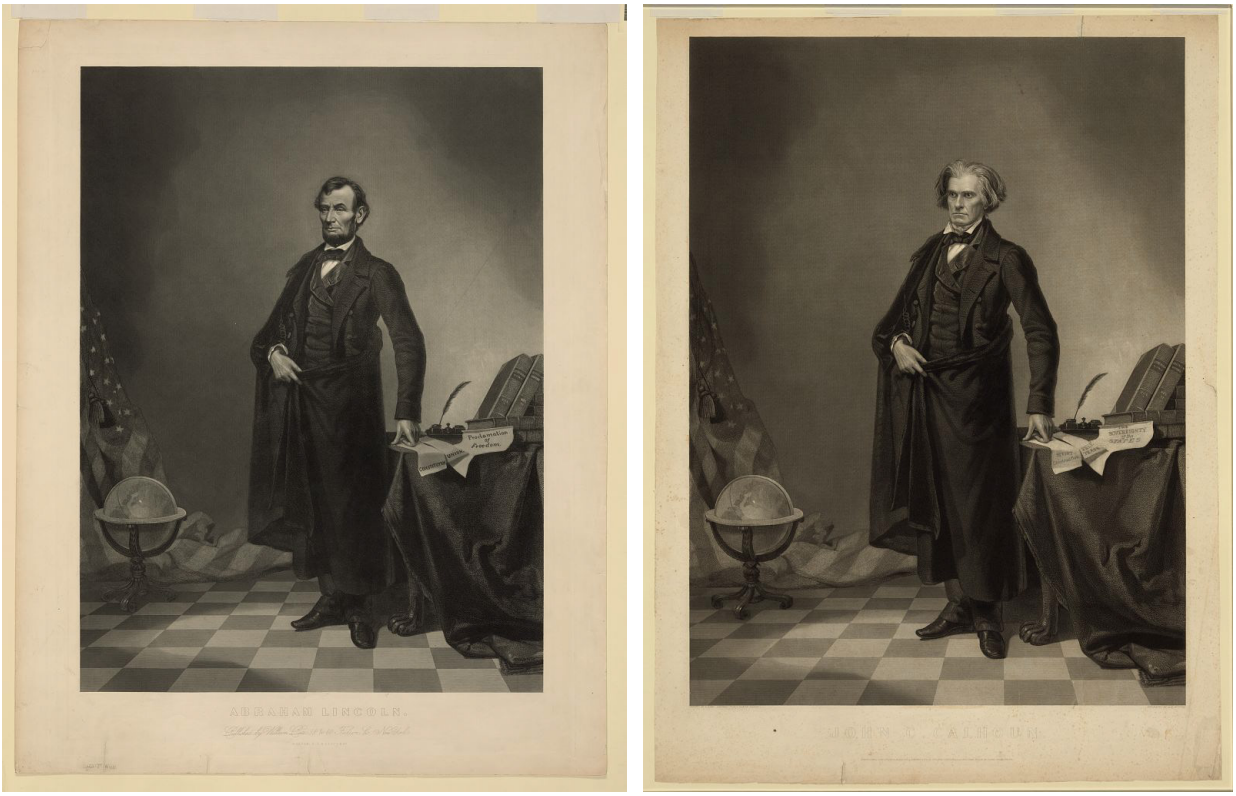

There are many commonsense solutions policymakers can take to address these concerns. First, and perhaps most importantly, Congress should not pass AI-specific privacy legislation, but instead pass comprehensive federal data privacy legislation that establishes basic consumer data rights (e.g., the right to opt out of data collection), preempts state laws, ensures reliable enforcement, streamlines regulation, and minimizes the impact on innovation.[12] A federal data privacy law would create legal safeguards to address many privacy concerns, such as setting rules on how organizations collect and share inferred PII and how employers collect and use PII about their employees. Importantly, such a law would apply irrespective of whether a data processor or data controller uses AI. A federal privacy law should also require data holders to disclose whether they make PII available to the governments of China, Russia, Iran, or North Korea which present unique threats given their histories of cyberattacks, including data theft, against the United States.

Second, Congress should pass targeted legislation to address specific privacy issues, including the following:

- Pass Federal Data Breach Notification: All 50 states, as well as the District of Columbia, Guam, Puerto Rico, and the Virgin Islands, have data breach laws, however, each jurisdiction has its own set of rules on how quickly to report a data breach or to whom a security incident should be reported.[13]

Congress should pass federal data breach legislation to create a single nationwide standard for reporting data breaches.

- Expand "Revenge Porn" Laws to Cover Manipulated Media: 48 states, Guam, and Puerto Rico have passed laws against nonconsensual distribution of intimate images.[14] In the absence of federal legislation, states should update these laws to ensure they cover fake media. Congress should pass legislation, such as the SHIELD Act, to address this problem at the federal level too.[15]
- Set Government Facial Recognition Performance Requirements: Congress should direct the General Services Administration to work with the National Institute for Standards and Technology to set performance standards—including for accuracy and error rates by age, race, and gender—for any facial recognition technology purchased using federal funds.[16] Congress should also direct the Department of Justice and the Department of Homeland Security to develop best practices on government use of facial recognition technology, including operational guidance and oversight protocols. The National Institute of Justice should also create best practices for state and local law enforcement to improve their use of facial recognition technology.[17]

Third, Congress should continue to invest in research for privacy- and security-enhancing technologies, as these will have important uses for AI. Additional research on topics such as secure multiparty computation, homomorphic encryption, differential privacy, federated learning, zero-trust architecture, and synthetic data can minimize or eliminate the need for AI-enabled services to process personal data while still maintaining the benefits of those services.[18] Many developers are already exploring solutions to address privacy concerns associated with large language models (LLMs). For example, some developers are exploring the use of "data privacy vaults" to isolate and protect sensitive data.[19] In this scenario, any PII would be replaced with deidentified data so that the LLM would not have access to any sensitive data, preventing data leaks during training and inference and ensuring only authorized users could access the PII.

AI AND LIABILITY

Legal liability has two main purposes: to compensate victims for any harm they have suffered and to internalize negative external effects to prevent or deter harm. By holding parties responsible for their actions, liability forces businesses to account for the costs of harm they cause. As a result, businesses have an incentive to develop safer products and unsafe products have a competitive disadvantage. In short, effective liability laws can spur consumer-friendly innovation.

Regarding liability for AI, policymakers should focus legal liability on operators (i.e., the parties responsible for deploying AI systems), rather than developers (i.e., the parties who make the underlying AI models), because operators make the most important decisions about how their applications impact others. For example, different applications, from customer service chatbots to healthcare applications, may use a particular LLM. The developer of the LLM is poorly situated to foresee and address all the potential risks of different applications, but the operator of the downstream application would have these insights.

Policymakers should promote a liability regime based on the idea of algorithmic accountability—the principle that operators should employ a variety of controls to ensure they can verify an algorithm works in accordance with their intentions and they can identify and rectify harmful outcomes.[20] In practice, algorithmic accountability requires regulators to consider factors such as whether an operator acted with negligence or

malicious intent and whether an operator took reasonable steps to minimize foreseeable risks from the use of its AI system. Operators who cause harm through negligence, such as by failing to consider the potential risks of their AI systems, should face more severe penalties than those acting in responsibly and in good faith.[21] An effective liability regime would incentivize AI operators to protect consumers from harm while also giving them the flexibility to manage their risk based on their knowledge of their products.

Policymakers should be careful to avoid imposing liability on online services that use AI when nearly identical services do not have liability as doing so would chill innovation. For example, Section 230 of the Communications Decency Act protects search engines from liability for defamatory statements that may be included in snippets of third-party content in search results. However, many legal observers do not believe that Section 230 applies to generative AI, which means that search engines using generative AI to answer user queries could face an influx of expensive lawsuits. To avoid this, Congress should amend Section 230 to shield online services and users from liability for content that is automatically generated using "information provided by another information content provider." This approach would recognize that using generative AI in this context is not unlike a search engine in the way it produces content: both require user input and generate results based on third-party information from outside sources. This approach would also ensure that online services can continue to experiment with emerging AI technologies and provide innovative and low-cost services for their users.[22]

CONCLUSION

There are important steps policymakers should take to address public concerns about the impact of AI on privacy and ensure liability fosters responsible innovation. However, policymakers should be careful not to overreact to concerns about AI and create laws and regulations with unintended consequences that harm innovation. Recent advances in AI have the potential to unlock substantial benefits in many areas of the economy and society, and the primary focus of policymakers should be on developing strategies to maximize those potential benefits rather than minimizing every possible harm.

REFERENCES

1. Tom Simonite, "This App Can Diagnose Rare Diseases From a Child's Face," *Wired*, March 8, 2022, https://www.wired.com/story/app-diagnose-rare-diseases-childs-face/.
2. Julie Jargon, "Fake Nudes of Real Students Cause an Uproar at a New Jersey High School," *The Wall Street Journal*, November 2, 2023, https://www.wsj.com/tech/fake-nudes-of-real-students-cause-an-uproar-at-a-new-jersey-high-school-df10f1bb.
3. Daniel Castro and Alan McQuinn, "The Privacy Panic Cycle: A Guide to Public Fears about New Technologies" (ITIF, September 2015), https://www2.itif.org/2015-privacy-panic.pdf.
4. Patrick Grady and Daniel Castro, "Tech Panics, Generative AI, and the Need for Regulatory Caution" (Center for Data Innovation, May 2023), https://www2.datainnovation.org/2023-ai-panic-cycle.pdf.
5. "2022 Annual Data Breach Report," Identity Theft Resource Center, January 25, 2023, https://www.idtheftcenter.org/post/2022-annual-data-breach-report-reveals-near-record-number-compromises/.
6. Nick Lally, "Examples of Data Points Used in Profiling," Privacy International, April 9, 2018, https://privacyinternational.org/sites/default/files/2018-04/data%20points%20used%20in%20tracking_0.pdf.
7. Johnny Long, "Google Hacking for Penetration Testers," Black Hat USA 2005, July 2005, https://www.blackhat.com/presentations/bh-europe-05/BH_EU_05-Long.pdf.
8. Reece Rogers, "How to Remove Your Personal Info From Google by Using Its 'Results About You' Tool," *Wired*, September 3, 2023, https://www.wired.com/story/results-about-you-remove-personal-info-from-google/.
9. Noah Shachtman, "Iran Missile Photo Faked (Updated)," *Wired*, July 10, 2008, https://www.wired.com/2008/07/iran-missile-ph/.
10. William Pate, engraver, *Abraham Lincoln*, Print. Boston, MA: L.A. Elliot & Co., c1865. Photograph. From Library of Congress: Popular Graphic Arts, https://www.loc.gov/pictures/item/2003654314/ (accessed November 5, 2023); Alexander Hay Ritchie, Engraver, *John C. Calhoun*, Print. New York, NY: A.H. Ritchie & Co, c.1852. Photograph. From Library of Congress, https://www.loc.gov/item/2003679757/ (accessed November 5, 2023).
11. Ibid.
12. Daniel Castro, Testimony to the House Administration Committee on "Big Data: Privacy Risks and Needed Reforms in the Public and Private Sectors," February 16, 2022, https://itif.org/publications/2022/02/16/testimony-house-administration-committee-big-data-privacy-risks-and-needed/.
13. "2022 Security Breach Legislation," National Conference of State Legislatures, September 29, 2022, https://www.ncsl.org/technology-and-communication/2022-security-breach-legislation.
14. "Nonconsensual Distribution of Intimate Images," Cyber Civil Rights Initiative, n.d., https://cybercivilrights.org/nonconsensual-distribution-of-intimate-images/ (accessed November 5, 2023).
15. SHIELD Act of 2023, S.412, 118th Congress (2023), Congress.gov, https://www.congress.gov/bill/118th-congress/senate-bill/412.
16. Daniel Castro, Hearing before the House Committee on Oversight and Reform, January 15, 2020, on "Facial Recognition Technology (Part III): Ensure Commercial Transparency & Accuracy," http://www2.itif.org/2020-commercial-use-facialrecognition.pdf.
17. Ibid.
18. National Science and Technology Council, "National Strategy to Advance Privacy-Preserving Data Sharing and Analytics" (Washington, D.C.: National Science and Technology Council, March 2023),

https://www.whitehouse.gov/wp-content/uploads/2023/03/National-Strategy-to-Advance-Privacy-Preserving-Data-Sharing-and-Analytics.pdf.

19. Joseph Williams, Lisa Nee, "Privacy Engineering," *Computer*, Vol 55, No 10, October 2022, https://ieeexplore.ieee.org/document/9903879.
20. Joshua New and Daniel Castro, "How Policymakers Can Foster Algorithmic Accountability" (Center for Data Innovation, May 2018), http://www2.datainnovation.org/2018-algorithmic-accountability.pdf.
21. "AI Risk Management Framework (AI RMF 1.0)," National Institute of Standards and Technology, January 2023, https://www.nist.gov/itl/ai-risk-management-framework.
22. Ashley Johnson, "Generative AI Is the Next Challenge for Section 230," Information Technology and Innovation Foundation, April 12, 2023, https://itif.org/publications/2023/04/12/generative-ai-is-the-next-challenge-for-section-230/.