Before the
National Telecommunications and Information Administration
Washington DC

In the Matter of	)	
	)	
Developing the Administration’s Approach	)	Docket No. 180821780–8780–01
	)	
to Consumer Privacy	)	

Comments by Jerry Y. Wei, Graduate Researcher, Princeton University

Submitted November 9, 2018

I am Jerry Y. Wei, Graduate Researcher in computer science at Princeton University's Center for Information Technology Policy.

I am responding to the National Telecommunications and Information Administration's Request for Comments on Developing the Administration's Approach to Consumer Privacy. I urge the Administration to encourage the research and use of privacy preserving technologies that would sustain both private companies' business models as well as individual consumer privacy.

Specifically, I argue that:

Innovation and consumer privacy protection need not be mutually exclusive, nor opposing forces in the Administration's approach. The correct use of different implementations in products, otherwise known as privacy substitutes, can be used in in conjunction with policy and educational actions to maintain consumer privacy while not stymying features nor functionalities the technology industry's services.

In almost all scenarios, there exist more than one design that can serve both the needs of product functionality while protecting user privacy. Alternative technical implementations to realize a product, known as privacy preserving technologies or privacy substitutes[1], should be viewed as a spectrum of possible choices, each with their own advantages and tradeoffs with respect to functionality and privacy. In many applications, a more privacy-preserving substitute can often provide the same functionality as a privacy-intrusive alternative.

We begin with a non-technical example from two Stanford University researchers to illustrate a real-world analogy to privacy preserving technology:

[1] Mayer and Narayanan, Privacy Substitutes (2013) (https://www.stanfordlawreview.org/online/privacy-and-big-data-privacy-substitutes)

> *"Suppose a coffee shop offers a buy-ten-get-one-free promotion. One common approach would be for the shop to provide a swipe card that keeps track of a consumer's purchases, and dispenses rewards as earned. An alternative approach would be to issue a punch card that records the consumer's progress towards free coffee. The shop still operates its incentive program, but note that it no longer holds a record of precisely what was bought when; the punch card keeps track of the consumer's behavior, and it only tells the shop what it needs to know."* (Mayer and Narayanan, Privacy Substitutes)

This latter implementation roughly illustrates how privacy preserving technologies work in web advertising: instead of sending data on user behaviour to advertisers, privacy preserving technologies store a user's online habits within the web browser itself, as well as selectively parceling out information derived from those habits.

To explore this concept further, consider the application where a first-party website has embedded third-party elements that display advertisements to viewers. A privacy-intrusive implementation might send the user's browsing history as well as their exact location to a third-party advertisement publisher in order to select which advertisement they display. A more privacy-preserving technology may simply send an approximate location or region including the user's location, as well as a heuristic for the user's browsing data, such as summarizing browsing data using keywords instead of disclosing all past web history. An even more privacy-preserving substitute may be to simply store all user-related data locally on the device itself and send nothing to the third-party website; rather, the third-party publisher sends a few options for marketing, and the local application decides which advertisement is most appropriate for the user.

The notion of privacy preserving technologies extends beyond online advertising. Consider an application that receives certificates that require authentication before trusting a service, such as logging onto a social media account. Current implementations of centralized identity management systems require applications to send certificates to a central system that validates a certificate's authenticity. The necessary communication between the user's application and the central system leads to the system knowing about user's activities. A privacy preserving substitute to this would be to have embedded features in the certificate, so that an application can validate its authenticity without the need to communicate with a central server. Using an analogy, one can imagine validating a one-dollar bill. The centralized system method would require users to bring the dollar note to a bank to be authenticated. As an alternative, the privacy preserving technology would simply ask users to look for anti-counterfeiting measures on the currency to validate its authenticity.

The use of privacy preserving technology extends beyond academic research and is widely used in industry as well. Apple has been vocal about its adoption of differential privacy, a privacy preserving technology that proves that personal data cannot be "de-anonymized" to learn about specific individuals, in its various products and services. Other major players in the technology space, such as Google[2] and Uber[3], have also been quick to investigate the benefits of differential privacy.

A simplified explanation of how differential privacy techniques works and how it could be applied to anonymize user data can be viewed through an example where a coffee shop wishes to know the most popular items on their menu. Instead of recording data on each

[2] Jordan Novet, Following Apple, Google is exploring differential privacy in Gboard for Android (2017) (https://venturebeat.com/2017/04/06/following-apple-google-tests-differential-privacy-in-gboard-for-android/)

[3] Katie Tezapsidis, Uber Releases Open Source Project for Differential Privacy (2017) (https://medium.com/uber-security-privacy/differential-privacy-open-source-7892c82c42b6)

customer and their specific orders, the shop can just tally up the number of times a particular item has been ordered and analyze this “aggregated” data at the end of the data. Some differential privacy techniques further protect user privacy by adding random noise to the data collected – say, adding and subtracting a few tallies for some items on the menu – without greatly affecting the validity of the end results.

Analogous to the example above, Apple uses differential privacy to “protect the privacy of user activity in a given time period, while still gaining insight that improves the intelligence and usability” of certain features, such as QuickType suggestions, Emoji suggestions, and sources behind battery usage[4]. To do this, Apple first transforms the data into a “differentially private” form, before it transmits data to be analyzed[5].

Other companies have dedicated resources to further researching and developing platforms to enable differential privacy. Google AI, the search company’s research division, presented Randomized Aggregatable Privacy-Preserving Ordinal Response, or RAPPOR in 2014. The goal is to provide a technology to “crowdsource statistics from end-user client software, anonymously, with strong privacy guarantees.” Its creators have described RAPPOR’s functionalities as “allow[ing] the forest of client data to be studied, without permitting he possibility of looking at individual trees.” [6]

While there are factors that discourage the adoption of privacy preserving technologies, such as implementation costs and the lack of understanding on the side of technology organizations, there are many actions that the Administration can pursue to encourage the use of

[4] Apple, Differential Privacy (2017) (https://www.apple.com/privacy/docs/Differential_Privacy_Overview.pdf)

[5] Andrew Greenberg, Apple’s Differential Privacy is About Collecting Your Data – But Not Your Data (2016) (https://www.wired.com/2016/06/apples-differential-privacy-collecting-data/)

[6] Erlingsson, Pihur, and Korolova, RAPPOR: Randomized Aggregatable Privacy-Preserving Ordinal Response (2014) (https://storage.googleapis.com/pub-tools-public-publication-data/pdf/42852.pdf)

such privacy preserving technologies. These actions, like the Administration's overall approach to consumer privacy, must also consider the context of each situation. The Administration must first understand the benefits and risks associated with possible privacy preserving technologies in a consumer privacy problem. Then it can begin incentivizing deployment of more privacy-preserving technologies that do not compromise on functionality, nor inhibit a sector's ability to innovate. These incentives can come in the form of legislation, media scrutiny, or offering safe harbour for companies that implement sufficiently secure and privacy-preserving technologies. Likewise, the Administration should also engage with industry to discuss privacy-preserving alternatives, and fund initiatives that serve to research or educate the use of privacy preserving technologies.

Conclusion

In my comment, I discussed how the use of privacy preserving technologies can offer a solution to protecting individual privacy without hindering innovation. The use of privacy preserving technologies to empower both users and companies alike to protect and limit the flow of consumer data. The Administration should urge the use of privacy preserving technologies in its proposal, while also encouraging further research through the allocation of resources on the development and effects of these privacy preserving technologies. While further steps are necessary to supplement user protection, I believe that these are important considerations for the Administration in its approach to consumer privacy.