Cart-ology: Intercepting Targeted Advertising via Ad Network Identity Entanglement

ChangSeok Oh, Chris Kanich[†], Damon McCoy[‡], and Paul Pearce
Georgia Institute of Technology, [†]University of Illinois Chicago, [‡]NYU

Problem

- Attacker can receive targeted ads intended for someone else
  - All that the attacker requires is the target's email address

- Attackers can execute this **Identity Entanglement** attack to compromise a victim's privacy and influence ads the victim sees

- Mitigations are possible, but require changes in ecosystem incentives

Targeted Advertising

Targeted ads are personalized advertisements that display merchandise with which the ad consumers interacted. The core enabler is cookie-based tracking.

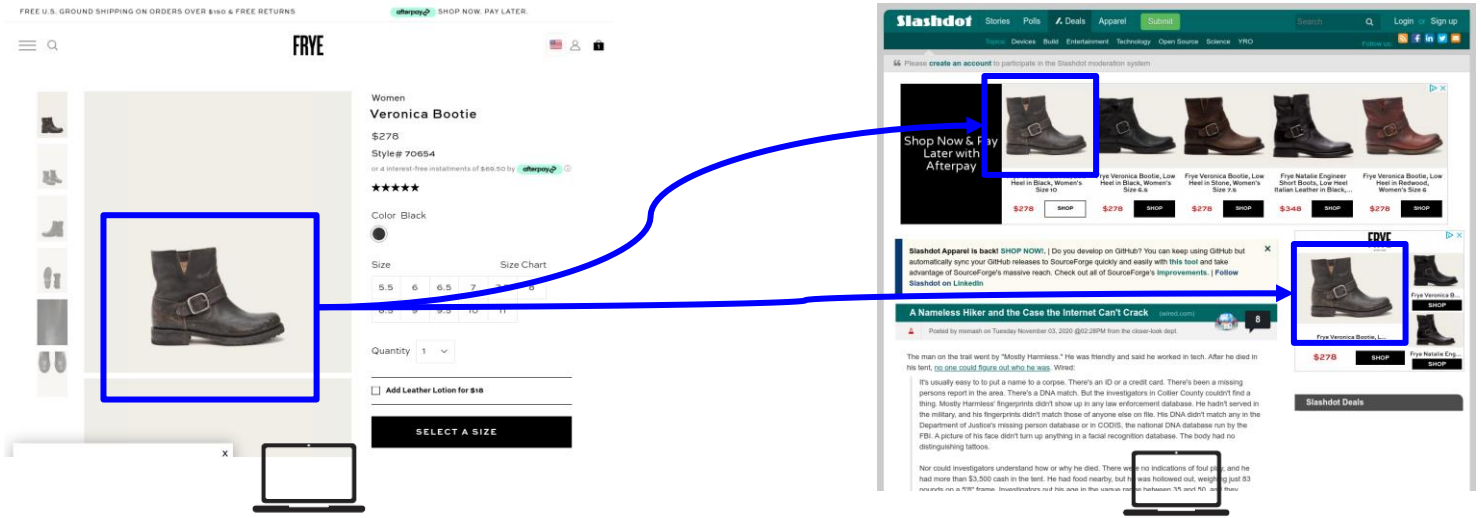

Cross-Device Tracking and Targeted Ads

Ad Networks identify users *across devices* in order to deliver consistent and personalized ads across contexts.

Targeted Ads on Cross-Devices

Third-party ad networks can use an email address to identify users with multiple devices.

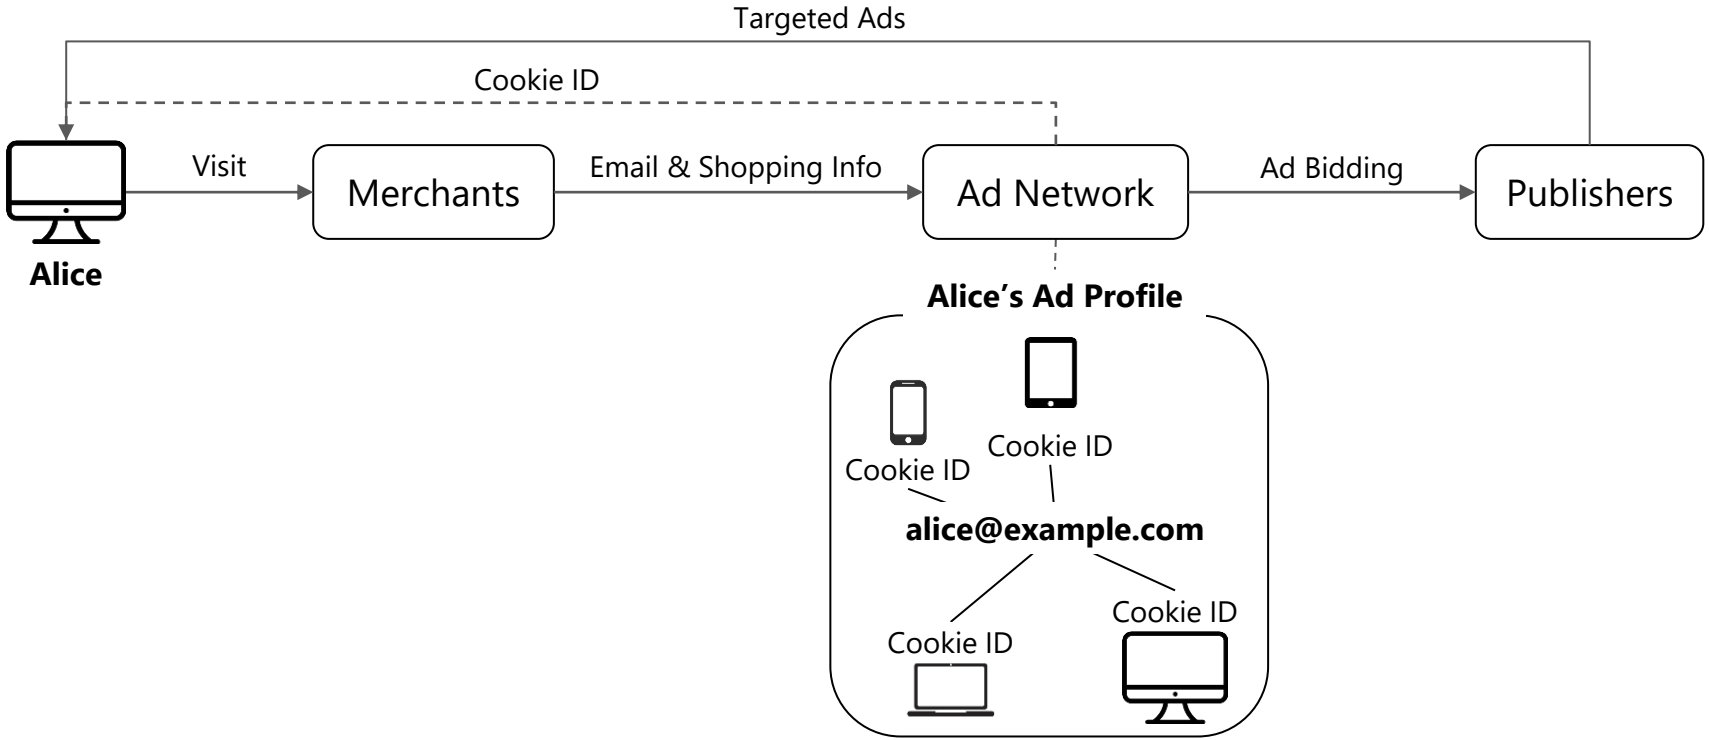

Identity Entanglement

The lack of integrity and authenticity in transferring the email address could confuse ad networks into delivering targeted advertisements to the wrong place.

- Ad companies cannot validate the information from their ad network.
- Many merchant websites do not check the ownership of the received email address.
- Attackers can exploit such tracking to entangle themselves with a victim on ad networks.

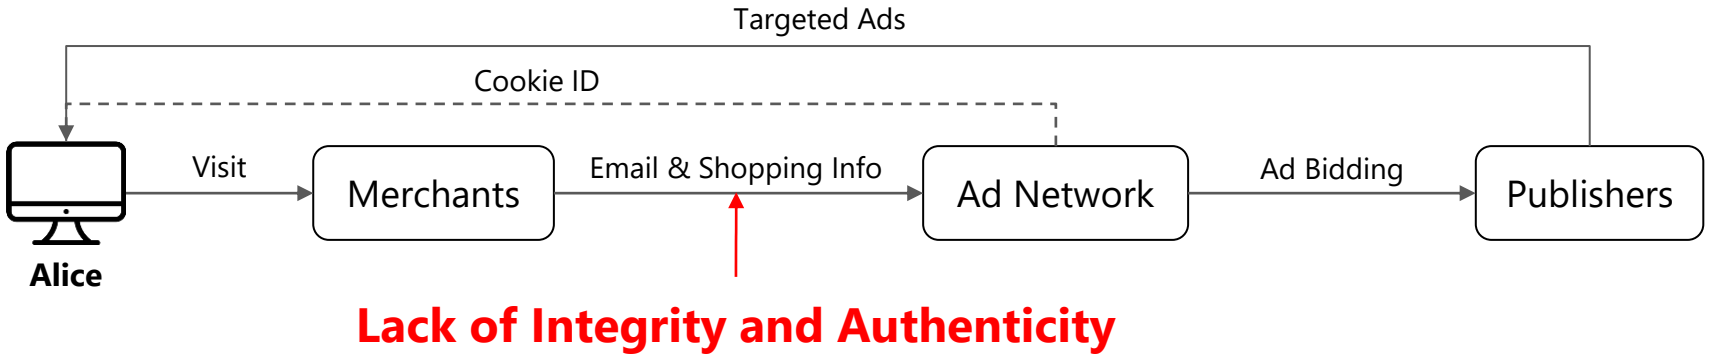

Identity Entanglement Attack

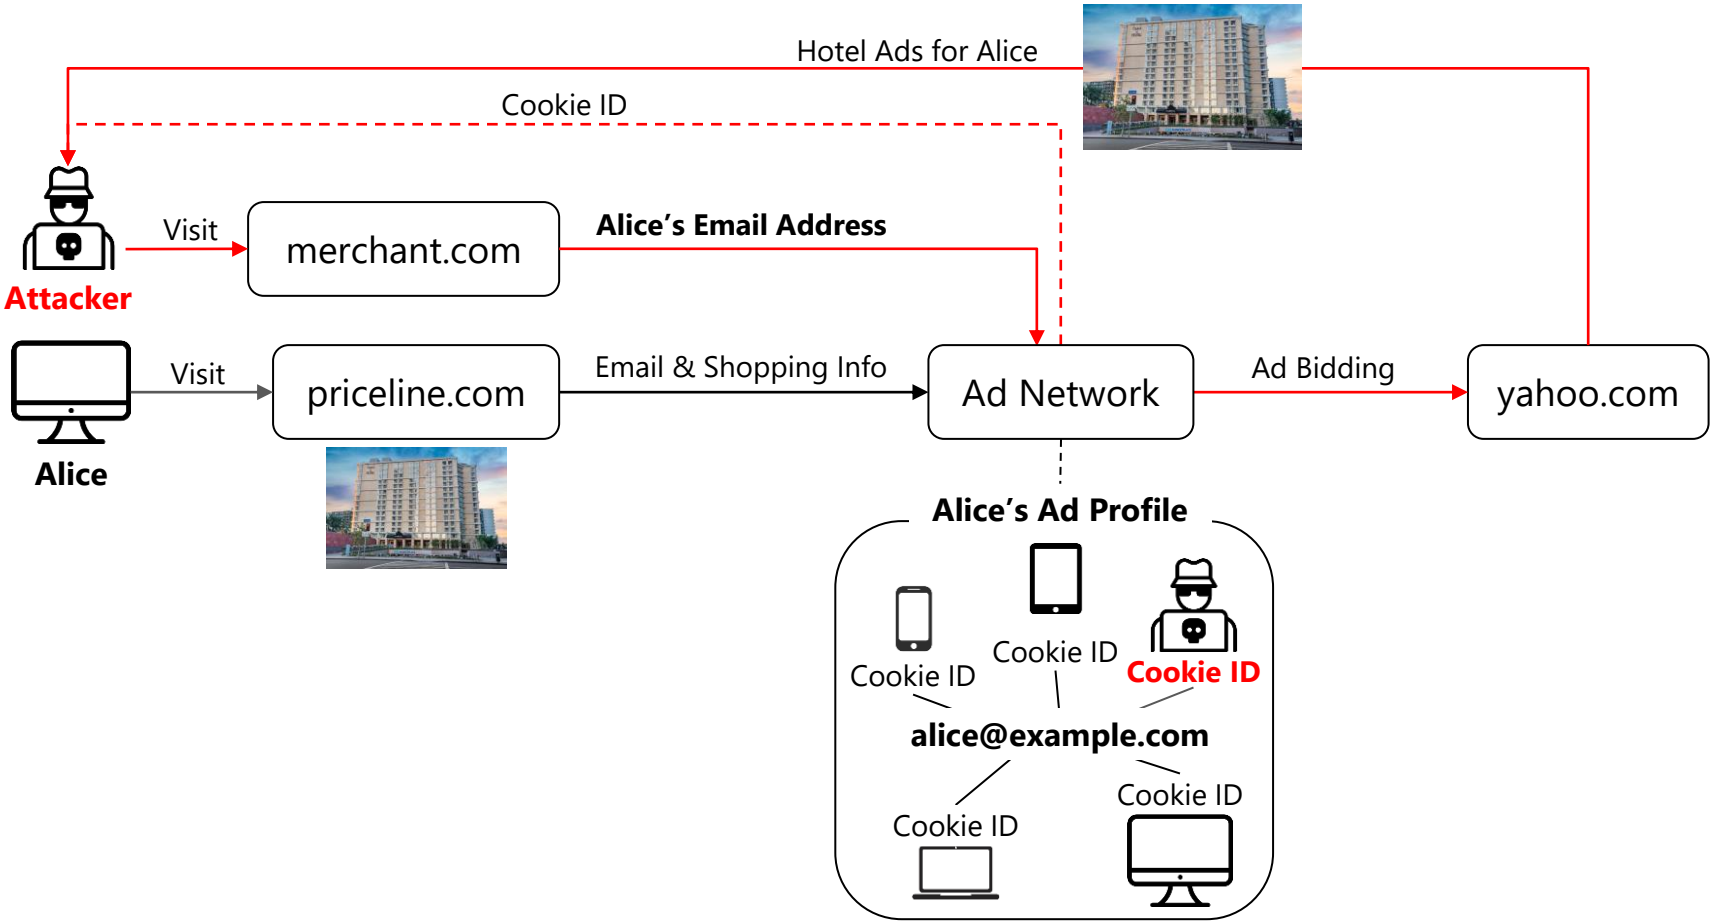

Privacy Implication of ID Entanglement Attacks

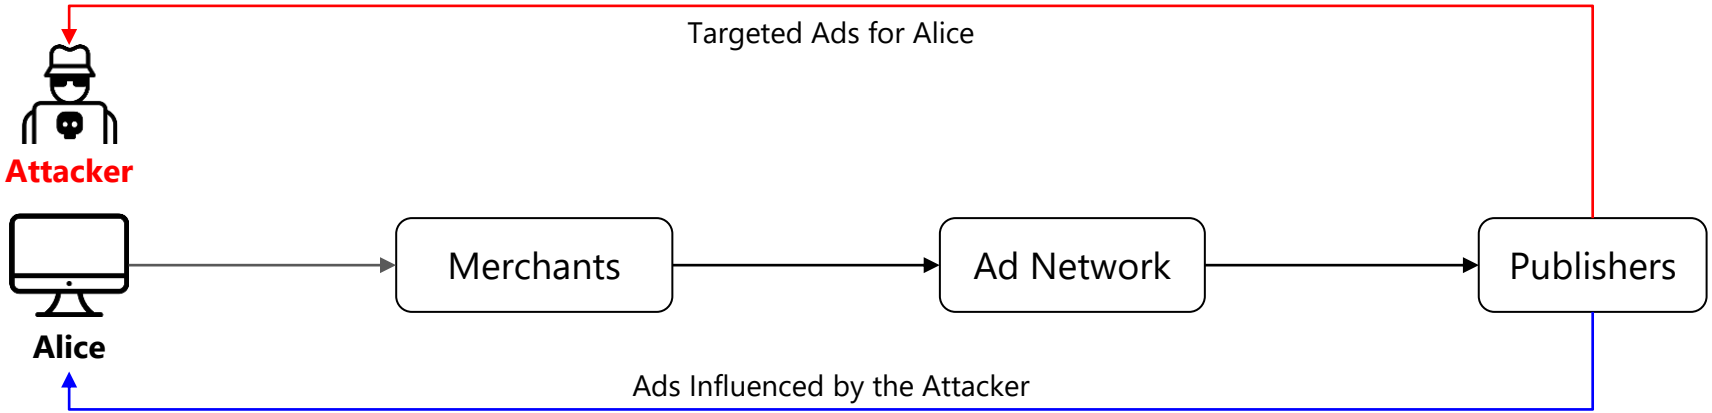

- Targeted ads contain privacy-sensitive information (e.g., where to live, where to visit, sexual orientation, etc.), so they could be weaponized as an attack channel.
- The attacker can violate the victim's privacy from the targeted Ads.
- The attack is bidirectional. The attacker can influence ads that the victim sees.

How can the attacker identify the victim’s ads?

The attack measures ads in entangled and baseline profiles while reloading the publisher's website, then finds outlier advertisers and items by the normalized difference.

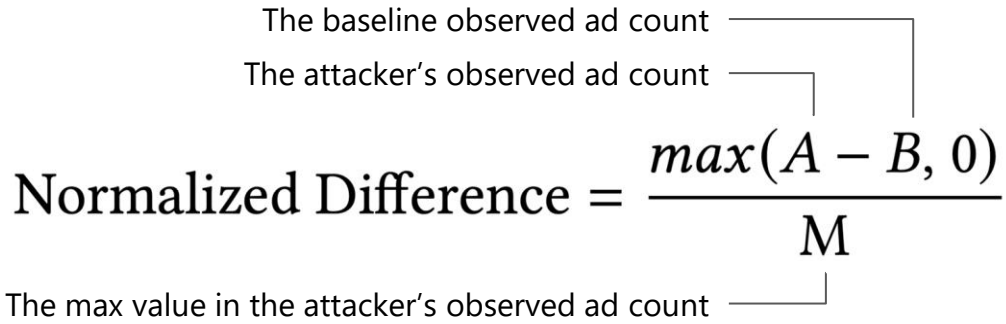

How can attackers identify what merchants a victim visits?

Using the normalized difference, the attacker can identify the merchant website the victim visited from general advertisers.

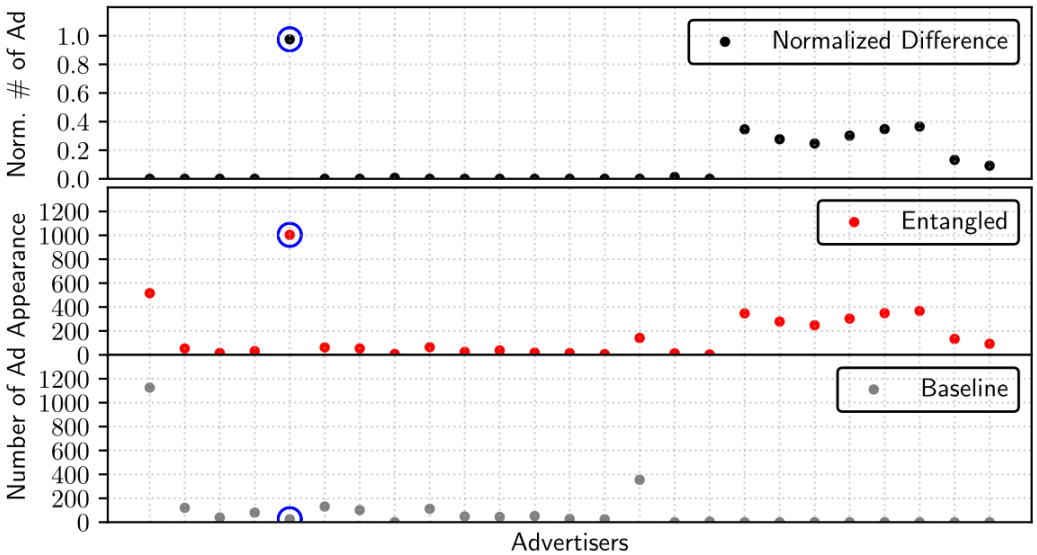

How can attackers identify the items a victim interacts with?

The specific product the victim engaged with shows up significantly more than other products.

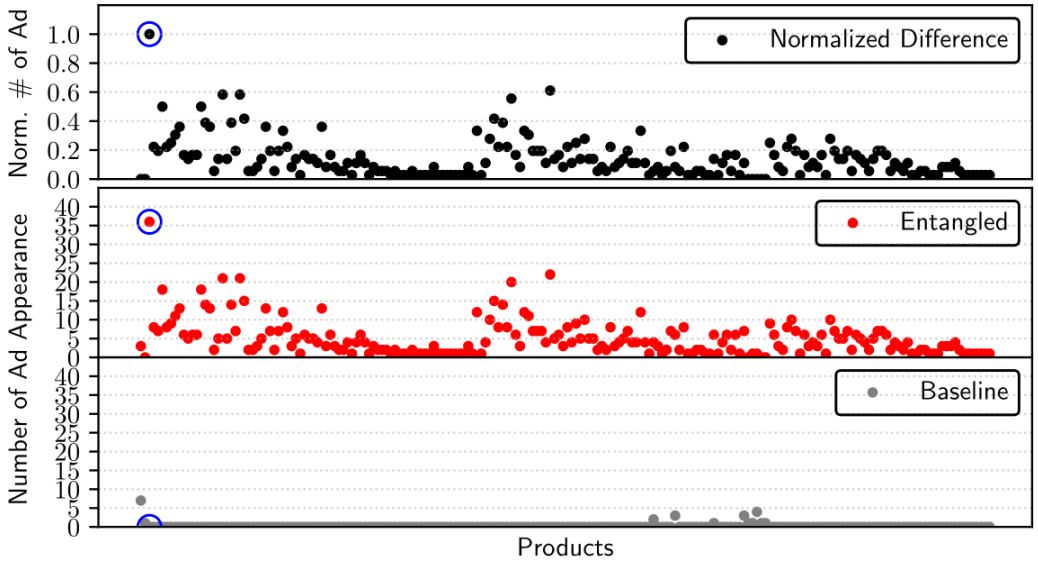

Attackers can influence what ads the victim sees.

The attack is bidirectional. Once attackers entangled their ID with the victim's ad profile, they can influence ads the victim sees.

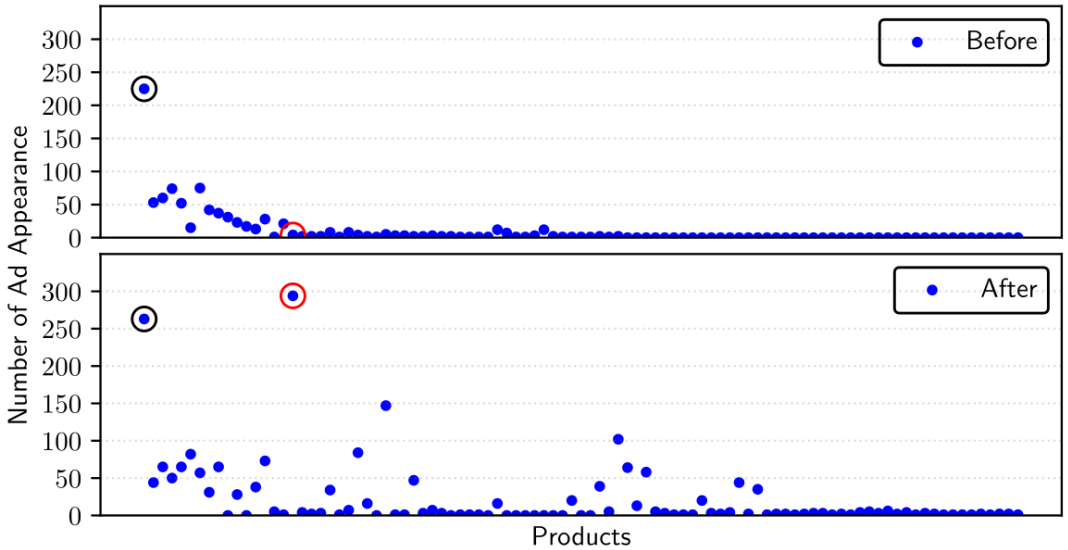

Takeaways

- We believe **any** leakage of user browsing behavior to an external attacker knowing only an email address should be **impossible.** We show it is possible and trivial.
- Targeted ads contain privacy-sensitive information and that, if delivered to the wrong person, represents a significant privacy problem.
- Attackers can execute this **Identity Entanglement** attack to compromise a victim's privacy and influence ads the victim sees.
- Mitigations are possible, but require changes in ecosystem incentives

Thank You