Online Advertising, Data Sharing, and Consumer Control

Thomas Jungbauer, Cornell University
(with Justin P. Johnson and Marcel Preuss)

ASSA 2024
San Antonio, TX

Changing Advertising Landscape

The market for online advertising is enormous (e.g., Facebook's 2022 revenue was $115 Billion) and rapidly changing, for example:

- Deprecation of third-party ad technologies (e.g. cookies, ATT)
- Category-based replacements (e.g. Google Topics)
- Consumer control over data (e.g. GDPR)
- Generative AI
- Campaign Management

Our goal: analyze some of these changes (the first three) in the context of *advertiser property rights over data*

The Consumer Journey and Advertising

A typical third-party cookie-enabled re-targeting timeline:

1. A consumer goes to the Nike website to look at Air Jordans
2. Nike *voluntarily* places a cookie in the consumer's browser *on behalf of an "ad exchange"*
3. The consumer visits other (third-party) websites
4. The consumer may be re-targeted by Nike on these websites

But What About Adidas?

- Perhaps a consumer who looked at Nike might prefer Adidas

(a) Nike Air Jordan 1 Mid

(b) Adidas Harden Volume 7

But What About Adidas?

- But unless the consumer visited Adidas they are unlikely to see an Adidas ad

(a) Nike Air Jordan 1 Mid

(b) Honor Among Thieves

But What About Nike?

- Similarly, consumers who visited only Adidas are unlikely to see a Nike ad

(a) Honor Among Thieves

(b) Adidas Harden Volume 7

Re-targeting and Cross-targeting

- *Cross-targeting* refers to a situation in which Adidas can target consumers who *only* visited Nike
  - (or vice versa)
- In principle, an ad exchange could enable cross-targeting
  - After all, it places cookies in the browsers of consumers who visit either Nike or Adidas
- In practice, this does not seem to be the case

Re-targeting and Cross-targeting

- *Cross-targeting* refers to a situation in which Adidas can target consumers who *only* visited Nike
  - (or vice versa)
- In principle, an ad exchange could enable cross-targeting
  - After all, it places cookies in the browsers of consumers who visit either Nike or Adidas
- In practice, this does not seem to be the case
- *This paper asks, "Why not?"*

Why No Cross-targeting?

- **Our answer:** Because Nike/Adidas *voluntarily choose to share data with an ad exchange*, and because they may not wish to share this data with rivals, ad exchanges don't offer this option

Why No Cross-targeting?

- **Our answer:** Because Nike/Adidas *voluntarily choose to share data with an ad exchange*, and because they may not wish to share this data with rivals, ad exchanges don't offer this option
- That is, advertisers have property rights over data regarding *intent to purchase*
- These property rights lead to inefficiencies in equilibrium
- **We also ask** whether certain industry changes (GDPR, 3rd-party cookie replacements) might improve efficiency

Related Literature

Privacy and data control

- Taylor (2004), Villas-Boas (2004), Montes, Sand-Zantman, and Valletti (2019), Choi, Jeon, and Kim (2019), Markovich and Yehezkel (2021), Dosis and Sand-Zantman (2022), Miklós-Thal, Goldfarb, Haviv, and Tucker (2023)

Online Advertising / Targeting

- Iyer, Soberman, and Villas-Boas (2005), Johnson (2013), Bergemann and Bonatti (2011), Athey, Calvano, and Gans (2018), Goldfarb and Tucker (2011), Villas-Boas and Yao (2021)

Data Markets/ Intermediaries

- De Corniere and De Nijs (2016), Bergemann and Bonatti (2015), Choi, Mela, Balseiro, and Leary (2020), Bergemann, Bonatti, and Gan (2022), Ichihashi (2022)

Online Advertising w/ focus on institutional details

- Sayedi (2018), Kraemer, Schnurr, and Wohlfarth (2020), and D'Annunzio and Russo (2020)

Ingredients to Our Model

- Advertisers have property rights over consumer *intent to purchase* data
- No individual advertiser can identify all consumers with intent
  - So: possible social gains to sharing data
- To re-target consumers, advertisers must share data (e.g. place cookies) with an ad exchange
  - Data sharing enables tracking and re-targeting

Ingredients to Our Model

- Advertisers have property rights over consumer *intent to purchase* data
- No individual advertiser can identify all consumers with intent
  - So: possible social gains to sharing data
- To re-target consumers, advertisers must share data (e.g. place cookies) with an ad exchange
  - Data sharing enables tracking and re-targeting
- An ad exchange decides whether to *further share* this data or not
  - That is, whether to allow *cross-targeting or not*
  - So that Nike can target Adidas' "exclusive" website visitors and vice versa

Base Model: Advertisers

- Two advertisers, A_1 and A_2
- A_1 and A_2 are competing for a unit mass of consumers
- A mass $\alpha < 1/2$ consumers have recently and exclusively visited A_i's "website." The remaining $1 - 2\alpha > 0$ consumers have recently visited both websites
  - **So:** For each advertiser, $1 - \alpha$ consumers have visited it, where α of these are exclusive visits

Base Model: Advertisers

- Two advertisers, A_1 and A_2
- A_1 and A_2 are competing for a unit mass of consumers
- A mass $\alpha < 1/2$ consumers have recently and exclusively visited A_i's "website." The remaining $1 - 2\alpha > 0$ consumers have recently visited both websites
  - **So:** For each advertiser, $1 - \alpha$ consumers have visited it, where α of these are exclusive visits
- A visit means consumers have *intent to purchase*
- To consummate a purchase, these consumers must be *re-targeted*, i.e. they must receive an ad (from that specific seller)

Base Model: Ad Exchanges

- There are two advertising exchanges, ADX_1 and ADX_2
- Each ADX_i permits *re-targeting*: advertiser A_i can reach all $1-\alpha$ consumers who have visited A_i
  - Only possible if A_i agrees to share its data with ADX_i

Base Model: Ad Exchanges

- There are two advertising exchanges, ADX_1 and ADX_2
- Each ADX_i permits *re-targeting*: advertiser A_i can reach all $1-\alpha$ consumers who have visited A_i
  - Only possible if A_i agrees to share its data with ADX_i
- (Only) ADX_1 may choose to allow *cross-targeting*: A_i can also reach the $1-\alpha$ consumers who visited its rival A_{-i}
  - Only possible if A_{-i} agrees to share its data with ADX_1

Base Model: Ad Exchanges

- There are two advertising exchanges, ADX_1 and ADX_2
- Each ADX_i permits *re-targeting*: advertiser A_i can reach all $1-\alpha$ consumers who have visited A_i
  - Only possible if A_i agrees to share its data with ADX_i
- (Only) ADX_1 may choose to allow *cross-targeting*: A_i can also reach the $1-\alpha$ consumers who visited its rival A_{-i}
  - Only possible if A_{-i} agrees to share its data with ADX_1
- (Only) ADX_1 lets advertisers reach *look-alike* consumers
  - There are η_1 look-alikes if one advertiser shares data with ADX_1, and $\eta_2 > \eta_1$ if both do

Base Model: Ad Exchanges

- There are two advertising exchanges, ADX_1 and ADX_2
- Each ADX_i permits *re-targeting*: advertiser A_i can reach all $1-\alpha$ consumers who have visited A_i
  - Only possible if A_i agrees to share its data with ADX_i
- (Only) ADX_1 may choose to allow *cross-targeting*: A_i can also reach the $1-\alpha$ consumers who visited its rival A_{-i}
  - Only possible if A_{-i} agrees to share its data with ADX_1
- (Only) ADX_1 lets advertisers reach *look-alike* consumers
  - There are η_1 look-alikes if one advertiser shares data with ADX_1, and $\eta_2 > \eta_1$ if both do
- Exchanges offer *frequency capping*

Base Model: Payoffs and Timing

- Each ADX_i charges an exogenous fee $w > 0$ to place an ad
- Each ADX_i receives an exogenous share of w per ad placed
  - (Unmodeled publishers receive the remaining share)
- A consumer who receives only one ad is worth M to the advertiser, whereas a consumer who receives two ads is worth $C < M$
- Product prices are fixed and $2C > M$, so that joint advertiser revenue is higher when consumers see two ads

Timing

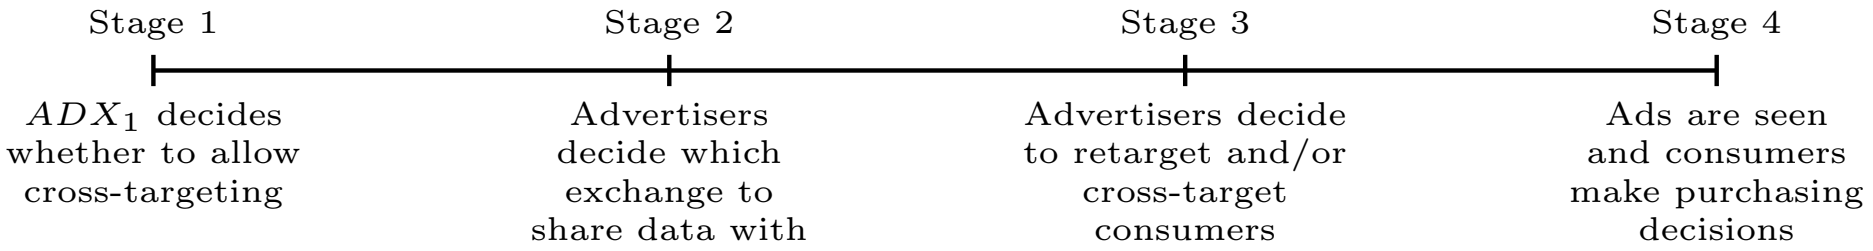

Importantly, the options available at Stage 3 depend on earlier Stage outcomes

1. A_i can only re-target on ADX_i if it shared data with ADX_i
2. A_i can cross-target A_{-i} only if
   - ADX_1 allows cross-targeting, AND
   - if A_{-i} shares data with ADX_1

Advertiser Perspecive: Risks of Data Sharing and Cross-targeting

- Suppose ADX_1 allows cross-targeting and both advertisers share data with it
- Cross-targeting pros and cons for A_i:
  - A_i can reach its rival's α exclusive consumers
  - A_i's α exclusives can be reached by its rival
- And sharing data with ADX_1 always increases the look-alike audience size
  - (Even if cross-targeting not permitted)

Data Sharing Preferences

- Advertisers may jointly prefer an equilibrium in which cross-targeting takes place (in addition to re-targeting):

$$\alpha(M - C) < \alpha(C - w) \iff M < 2C - w,$$

  i.e., the lost monopoly power over "exclusive" customers is made up by the (net) value of reaching the rival's exclusive customers
- (Assuming both advertisers use ADX_1)

Data Sharing Preferences

- Advertisers may jointly prefer an equilibrium in which cross-targeting takes place (in addition to re-targeting):

$$\alpha(M - C) < \alpha(C - w) \iff M < 2C - w,$$

  i.e., the lost monopoly power over "exclusive" customers is made up by the (net) value of reaching the rival's exclusive customers
- (Assuming both advertisers use ADX_1)
- Because $2C > M$, this holds when w is "small"
- However, when w is too small, things get complicated...

Stay with ADX_1 or Leave?

The "Leave" defection is profitable if and only if

$$\underbrace{\alpha(M-C)}_{\substack{\text{Defend Your}\\ \text{Exclusives}}} \geq \underbrace{\alpha(C-w)}_{\substack{\text{Loss From Not}\\ \text{Targeting Rival's}\\ \text{Exclusives}}} + \underbrace{(\eta_2-\eta_1)(C-w)}_{\substack{\text{Smaller Look-alike}\\ \text{Audience}}}.$$

- Recall that if $\alpha(M-C) < \alpha(C-w)$ then advertisers *jointly prefer cross-targeting* (assuming both use ADX_1)

Stay with ADX_1 or Leave?

The "Leave" defection is profitable if and only if

$$\underbrace{\alpha(M-C)}_{\substack{\text{Defend Your}\\ \text{Exclusives}}} \geq \underbrace{\alpha(C-w)}_{\substack{\text{Loss From Not}\\ \text{Targeting Rival's}\\ \text{Exclusives}}} + \underbrace{(\eta_2-\eta_1)(C-w)}_{\substack{\text{Smaller Look-alike}\\ \text{Audience}}}.$$

- Recall that if $\alpha(M-C) < \alpha(C-w)$ then advertisers *jointly prefer cross-targeting* (assuming both use ADX_1)
- Moreover, if A_i finds it profitable to "Leave" then its rival is unharmed by A_i leaving
- So the option of "leaving" is not problematic from advertisers' joint perspective

The "Leave and Snipe" Deviation

But advertisers have an additional defection from a data-sharing equilibrium

- A_i could try and get the best of both worlds
- Share data (only) with ADX_2 but continue to cross-target A_{-i}

$$\alpha M + (1-\alpha)C - 2(1-\alpha)w + \eta_1(C-w).$$

- Leave & Snipe is profitable if and only if

$$\underbrace{\alpha(M-C)}_{\substack{\text{Defend Your} \\ \text{Exclusives}}} \geq \underbrace{(1-2\alpha)w}_{\substack{\text{Wasted Impressions} \\ \text{from Using} \\ \text{Two Exchanges}}} + \underbrace{(\eta_2-\eta_1)(C-w)}_{\substack{\text{Smaller Look-alike} \\ \text{Audience}}}.$$

The "Leave and Snipe" Deviation

But advertisers have an additional defection from a data-sharing equilibrium

- A_i could try and get the best of both worlds
- Share data (only) with ADX_2 but continue to cross-target A_{-i}

$$\alpha M + (1-\alpha)C - 2(1-\alpha)w + \eta_1(C-w).$$

- Leave & Snipe is profitable if and only if

$$\underbrace{\alpha(M-C)}_{\substack{\text{Defend Your} \\ \text{Exclusives}}} \geq \underbrace{(1-2\alpha)w}_{\substack{\text{Wasted Impressions} \\ \text{from Using} \\ \text{Two Exchanges}}} + \underbrace{(\eta_2-\eta_1)(C-w)}_{\substack{\text{Smaller Look-alike} \\ \text{Audience}}}.$$

- Moreover, A_i's rival is *always harmed* when A_i "Leaves & Snipes"

Equilibrium of Base Model

Proposition

In equilibrium, both advertisers share their data with ADX_1, and retarget their exclusive website visitors and the look-alike audience of size η_2 using ADX_1. In addition, ADX_1 allows cross-targeting if and only if neither "Leave" nor "Leave & Snipe" are profitable, in which case both advertisers also cross-target on ADX_1.

- If ADX_1 allowing cross-targeting causes both advertisers to share data with it, then ADX_1 indeed allows cross-targeting
- Also ensures the size of the look-alike audience is maximized
- *Consumers are harmed if ADX_1 does not allow cross-targeting in equilibrium—but regulating this outcome need not help*

A Prisoner's Dilemma

- Recall that advertisers jointly prefer a full cross-targeting outcome if and only if

$$\alpha(M - C) < \alpha(C - w) \iff M < 2C - w$$

- If this holds, the "Leave" defection is not profitable
- However, "Leave & Snipe" may be profitable

$$\underbrace{\alpha(M - C)}_{\substack{\text{Defend Your} \\ \text{Exclusives}}} \geq \underbrace{(1 - 2\alpha)w}_{\substack{\text{Wasted Impressions} \\ \text{from Using} \\ \text{Two Exchanges}}} + \underbrace{(\eta_2 - \eta_1)(C - w)}_{\substack{\text{Smaller Look-alike} \\ \text{Audience}}}.$$

for example, if w and $\eta_2 - \eta_1$ are small.

A Prisoner's Dilemma

- A prisoner's dilemma caused by strong advertiser property rights over consumer intent data
- Note that the 3rd-party cookie system indeed grants advertisers strong property rights
- Hence, these advertiser "data rights" are potentially problematic
- Could eliminating these rights benefit consumers? Or even advertisers?

Extension: Category-Based Advertising

- Google says it is phasing out 3rd-party cookies in Chrome
- Says it will replace it with Topics
  - Chrome will track which websites a consumer visits and use on-browser ML techniques to assign the consumer to categories
  - Advertisers can bid to reach consumers in different categories
- Advertisers complain Topics will be less accurate
- Note that Topics *weakens advertiser property rights* compared to the cookie system

Category-Based Advertising ADX_1

We model a Topics-inspired category system as follows:

- ADX_1 offers category-based advertising
- All advertiser data is shared by default
- But category system is inaccurate and results in some wasted impressions, indexed by a parameter τ
- An advertiser must send $\tau > 1$ ads to reach a unit mass of consumers, so $\tau - 1$ represents wasted ads

Category-Based Advertising ADX_2

- ADX_2 uses an alternative ad technology that only allows re-targeting (no data sharing between advertisers)
- These systems may also be imperfect. For example, may require user opt-in
- Suppose tech allows $\sigma < 1$ of website visitors to be re-targeted

Extension: Category-Based Advertising

Proposition

Suppose that the Topics system is not too inaccurate, that is if $\tau > 1$ is sufficiently close to 1. Then in the equilibrium of the Topics model, both advertisers use ADX_1 to re-target and cross-target.

- So Topics makes it is easier to support cross-targeting
- An advertiser could use ADX_2. . . but, unlike in base model, using ADX_2 does not prevent a rival from cross-targeting
- But not all consumers are better off: because $\tau > 1$, some consumers are inaccurately targeted

Topics: Reduced Accuracy and Advertiser Expropriation

Proposition

In the region of Topics accuracy where both advertisers strictly prefer to bid on ADX_1: (i) ADX_1's profits strictly increase as the Topics technology becomes less accurate, that is as τ increases. Moreover, (ii) this effect increases as its rival ADX_2 is less efficient (σ is smaller).

- ADX_1 may have incentives to reduce accuracy—which may increase its ad sales
- So, depriving advertisers of data property rights has mixed effects
- **Note**: Similar effects may exist on any platform that has "full data control," e.g. ads on Facebook or iOS, where the platform "sees everything"

Other Extensions Considered

- Heterogenous consumers can opt out of tracking
- Consumers who visit both websites are more likely to make a purchase
- Cross-targeting is more expensive (due to increased competition for eyeballs)
- Some consumers buy without being re-targeted with an ad
- Ad prices are endogenous and vary by ADX

Conclusion

- We consider targeted advertising when ad exchanges and advertisers play important roles in collecting/using data
- We identify a new insight: strong advertiser property rights over data leads to insufficient sharing of information
- Weakening advertiser property rights may benefit consumers—a potentially positive role of central platform control by firms like Google, Apple, or Meta
- Giving consumers partial data property rights makes many consumers better off by changing ADX/advertiser actions

Online Advertising, Data Sharing, and Consumer Control

Thomas Jungbauer, Cornell University
(with Justin P. Johnson and Marcel Preuss)

ASSA 2024
San Antonio, TX