The Economics of PEAT

Investigating the Adoption and Performance of Privacy-Enhanced Advertising Technologies

Garrett Johnson
Assistant Professor of Marketing
Questrom School of Business
@garjoh_canuck

Privacy Sandbox

Toulouse School of Economics

CONFERENCE SECRETARIAT

Florence CHAUVET

Christelle FAUCHIE

ORGANIZING COMMITTEE

Alexandre DE CORNIERE

Jacques CREMER

Paul SEABRIGHT

MERCI!

The Economics of PEAT

Investigating the Adoption and Performance of Privacy-Enhanced Advertising Technologies

Garrett Johnson
Assistant Professor of Marketing
Questrom School of Business
@garjoh_canuck

Privacy Sandbox

Toulouse School of Economics

Overview

- Privacy enhancing technologies & economics
- Privacy Sandbox primer
- Privacy Sandbox adoption
- Privacy Sandbox experiment results
  - Publisher revenue
  - Advertiser performance

Privacy enhancing technologies (PETs) & economics

Privacy & data economy: two sides of the same coin

Privacy & the data economy in tension

Promise of privacy-enhancing technologies (PETs)

PETs: Some examples

High-Level Use Case	PETs
Anonymizing Data for Analysis or Sharing	• Differential Privacy (DP) • K-Anonymity
Using Data Without Sharing It	• On-Device Computation • Trusted Execution Environments (TEE) • Federated Learning • Homomorphic Encryption
Combining Data Privately Between Parties	• Multi-Party Computation (MPC)

Why might firms use (costly) PETs?

Like air pollution scrubbers, PETs are costly but can enable new data sharing

- Regulation or self-regulation
- Improve quality: e.g., privacy, utility, personalization
- Restrict rival's capacity to process data
- Altruism/PR/"Privacy-washing"
- Gain capacity to process another's data

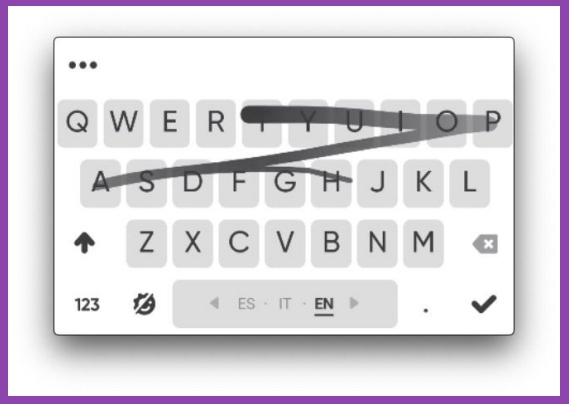

Ex.: Google uses **Federated Learning** to improve keystroke recognition by training models on users' devices without collecting personal data. Only model updates, not the raw typing data, are shared, enhancing privacy while refining the model across devices.

Why might economists care about PETs?

- It's coming for our data!
  - Debate over Differential Privacy in government statistics
  - Challenges for inference
  - NBER conferences on "Data Privacy Protection and the Conduct of Applied Research"
- Economists study tradeoffs: valuable to formalize tradeoffs & improve upon extant PETs
- Relevant subfields: empirical micro, econometrics, theory, industrial organization, innovation, digitization, public regulation
- We can't all study AI :P

Privacy Sandbox primer

Digital Ads: Privacy & the data economy in tension

Cookies enable cross-site ad targeting & measurement

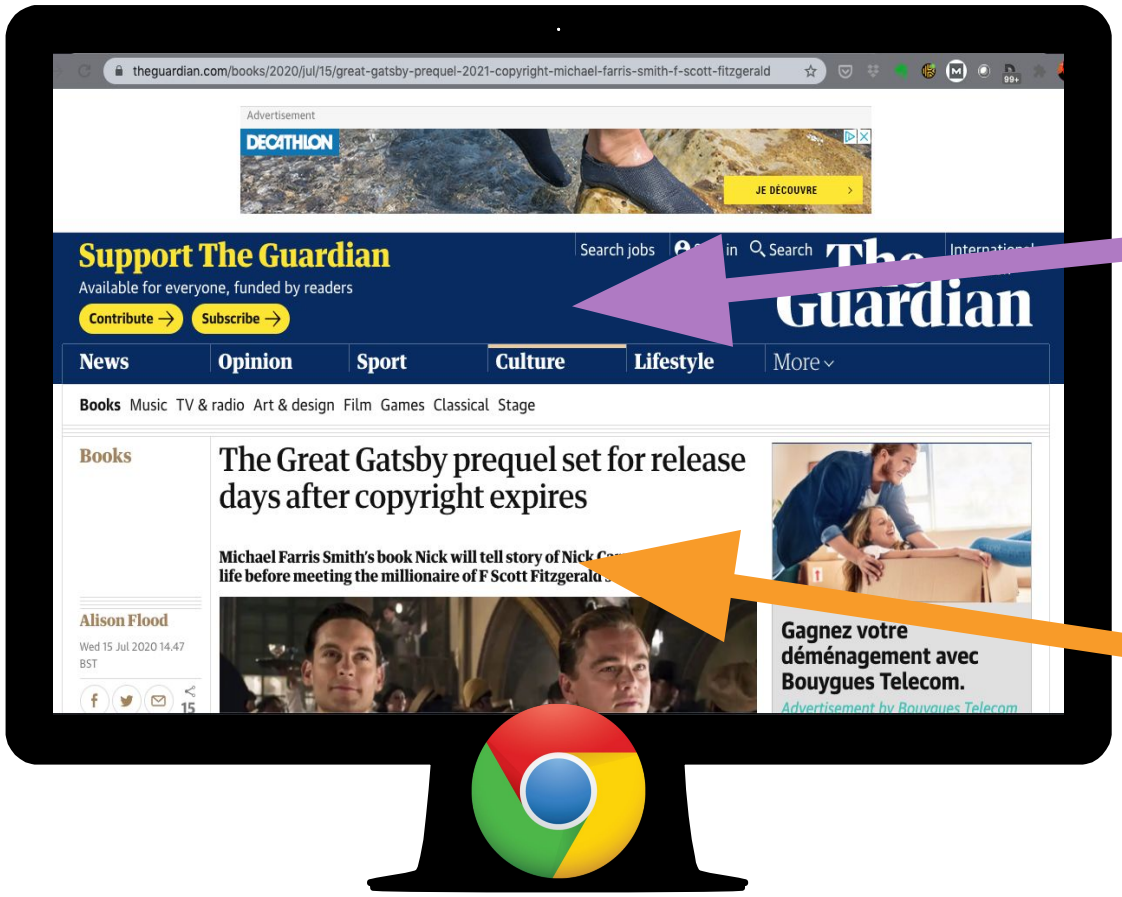

theguardian.com

1st party cookie
User ID = "ABC123"

doubleclick.net

3rd party cookie
User ID = "EFG456"

Ad identity increases ad revenue by 2X-3X*

(In the status quo)

*Source: Beales & Eisenach (2014); Johnson, Shriver, & Du (2020); Google (2019); UK CMA Report (2020); Alcobendas, Kobayashi, & Shum (2021); Sousa (2024)
NB: Laub, Miller, & Skiera (2022); Marotta, Abhishek, & Acquisti (2019) find smaller effects

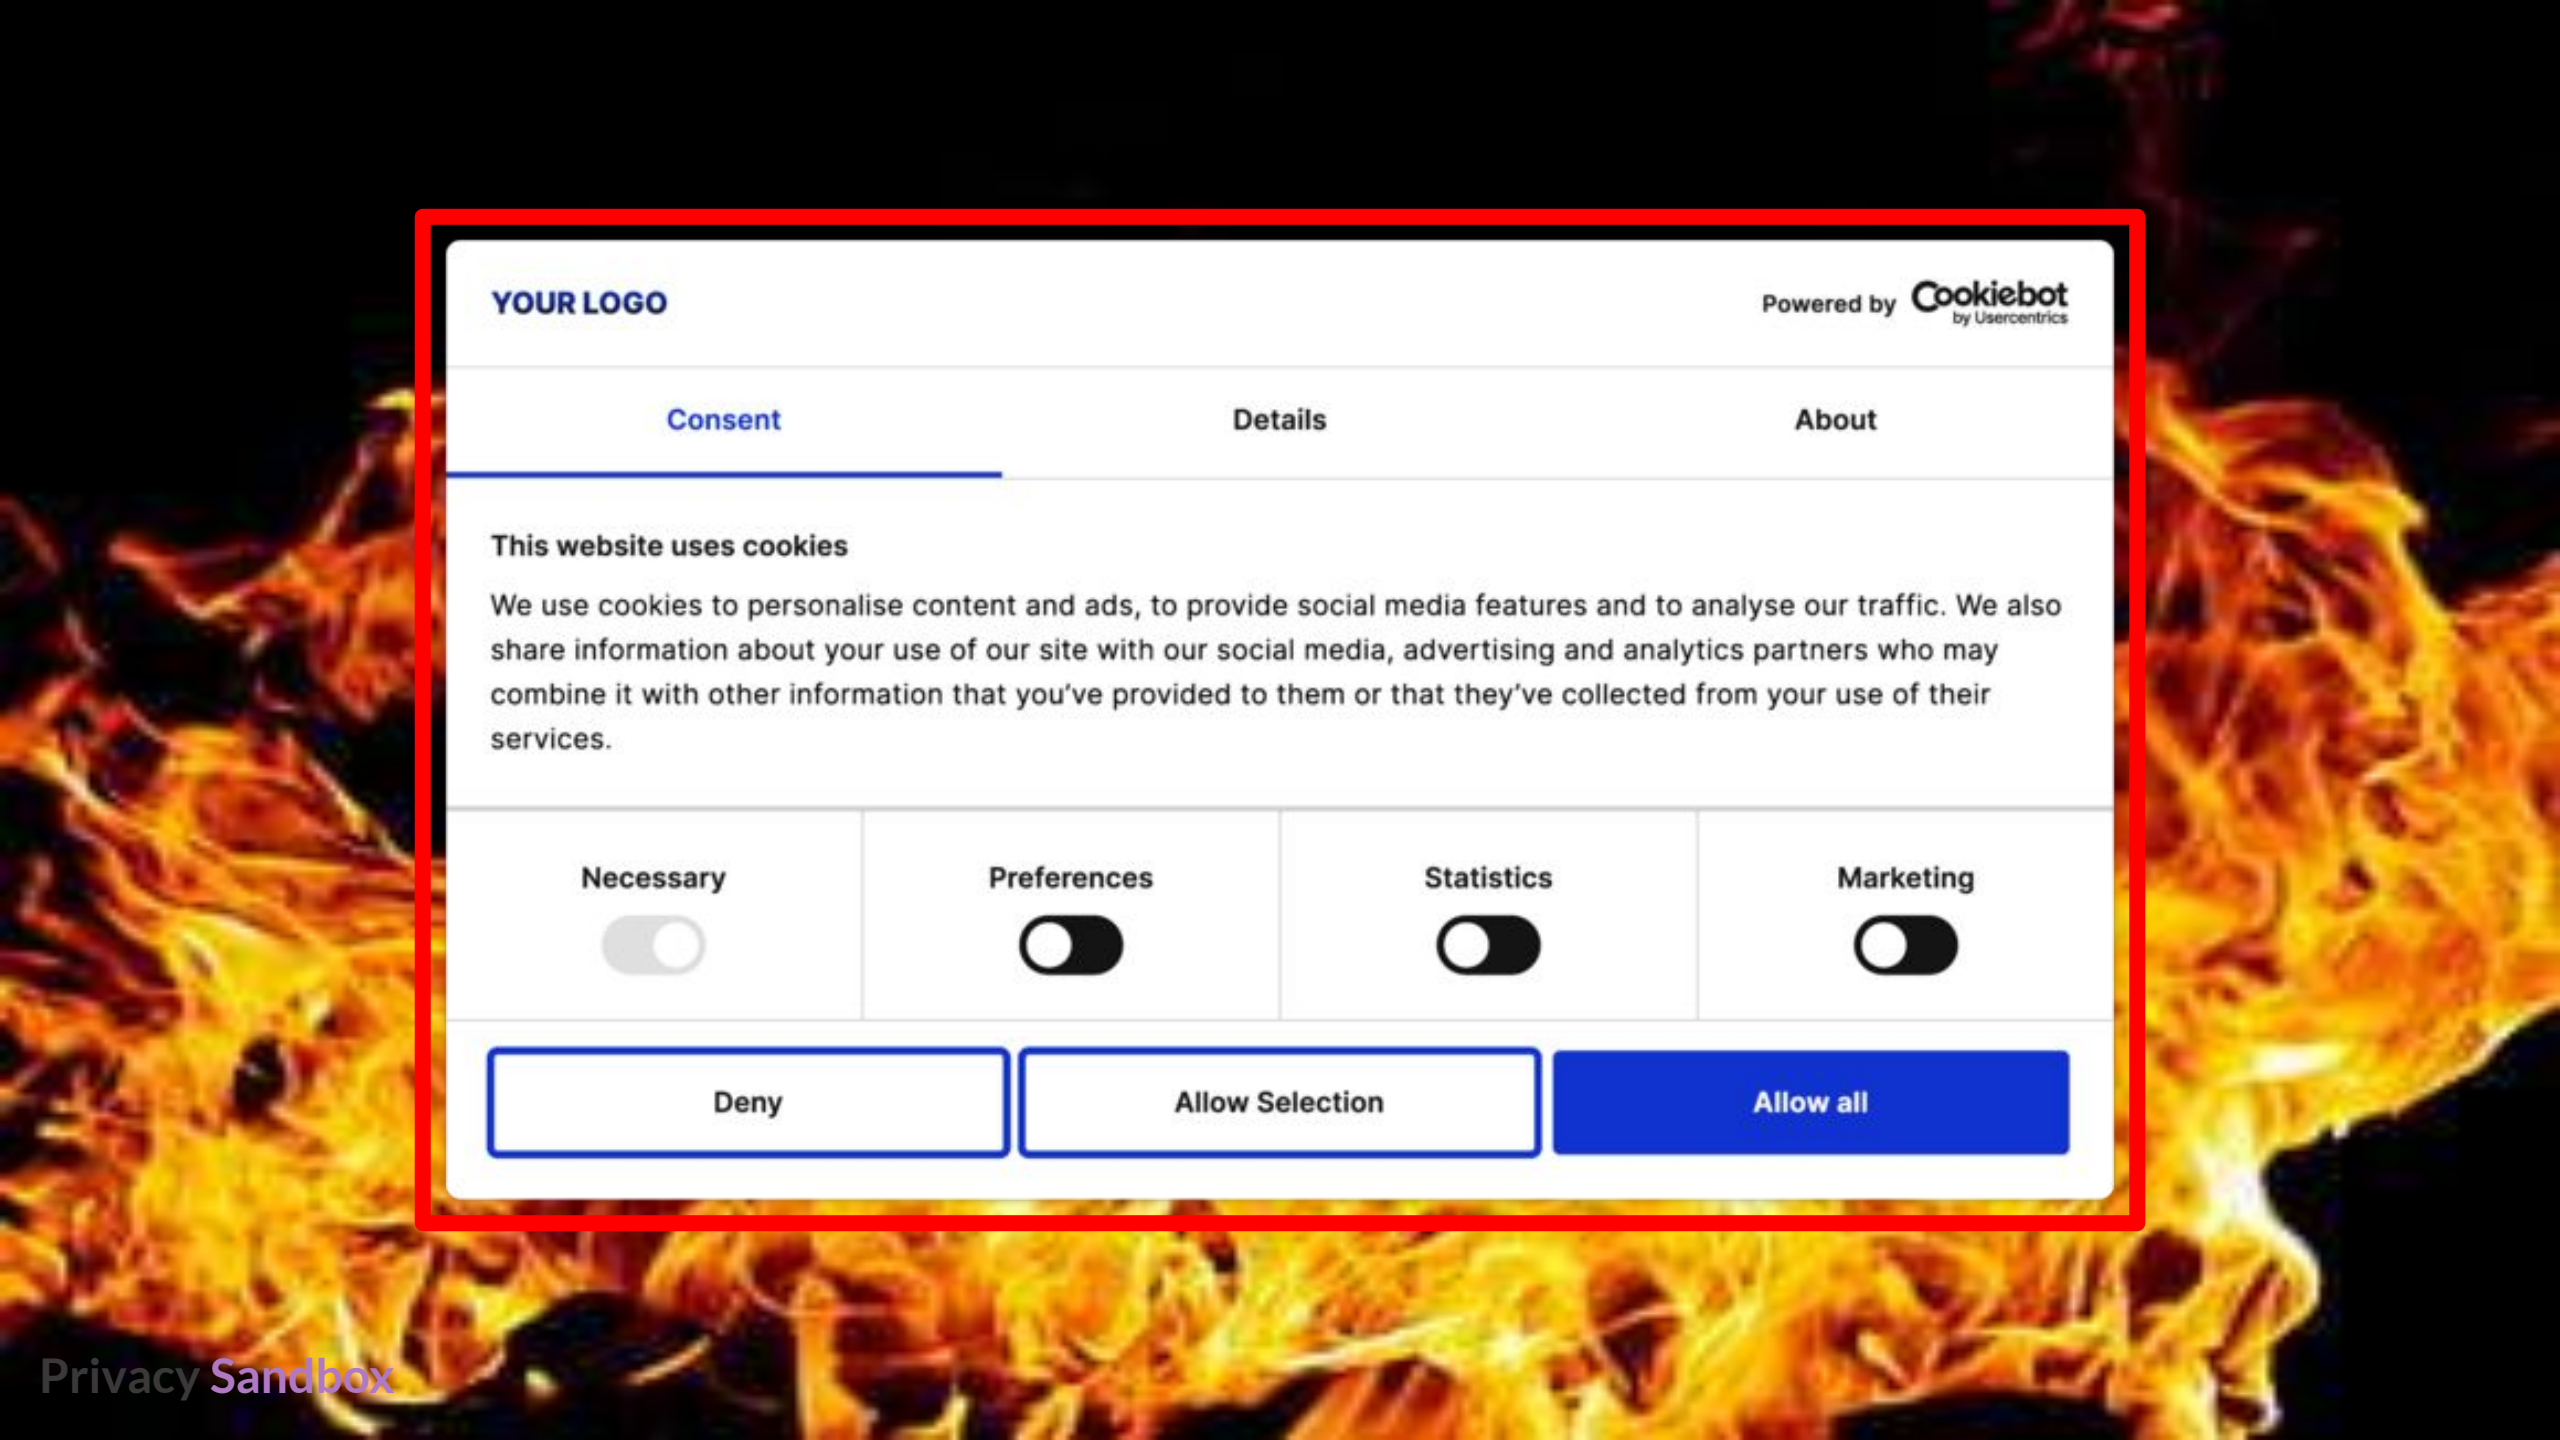
YOUR LOGO
Powered by Cookiebot by Usercentrics
Consent
Details
About
This website uses cookies
We use cookies to personalise content and ads, to provide social media features and to analyse our traffic. We also share information about your use of our site with our social media, advertising and analytics partners who may combine it with other information that you've provided to them or that they've collected from your use of their services.
Necessary
Preferences
Statistics
Marketing
Deny
Allow Selection
Allow all
Privacy Sandbox

Promise of privacy-enhancing advertising technologies (PEATs)

Privacy Sandbox & timeline

Google's "**Privacy Sandbox**" consists of multiple technologies that aim to preserve the benefits of cross-site/app identity in online advertising while offering superior privacy protection.

Aug. 2019

*Google **launches Privacy Sandbox** cookie-deprecation initiatives*

Sep. 2023

*Chrome made the Privacy Sandbox **APIs generally available***

Jan 2024

*Chrome deprecates cookies for **1% of user base** (30 million users globally)*

*Google/CMA **global experiment***

July 2024

*Google pivots from cookie deprecation to **user browser-level cookie consent** in 2025 (TBC)*

Overview: Privacy Sandbox technologies

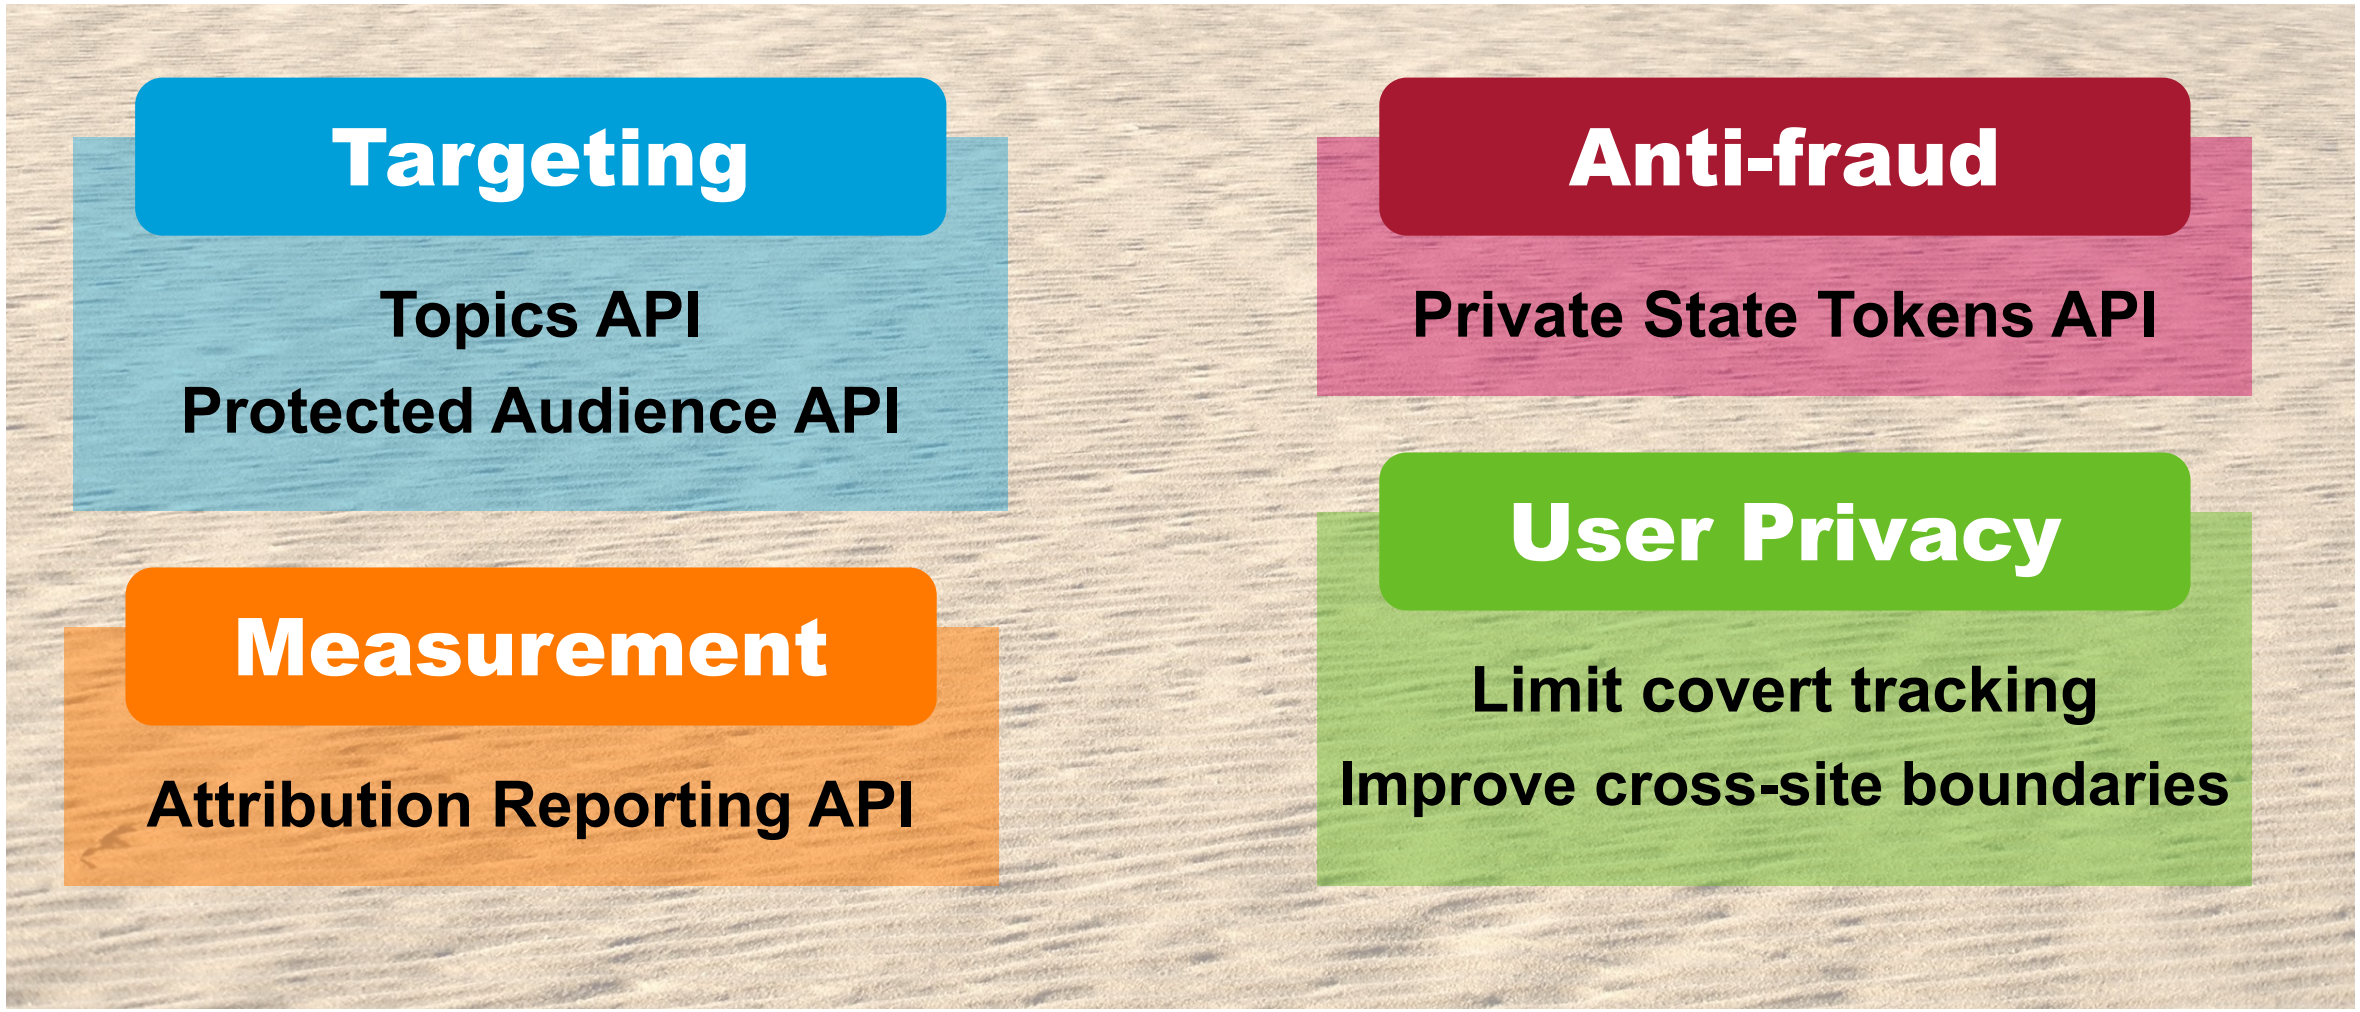

Privacy Sandbox adoption

"Unearthing Privacy-Enhancing Ad Technologies (PEAT): The Adoption of Google's Privacy Sandbox"

Supported by:

- **Sincera**
- **Digital Business Institute** (Questrom School of Business)
- **Rafik B. Hariri Institute** (Boston University)
- **Program on Economics & Privacy** (George Mason University)

Data: Tracking adoption

- **Data**: Track Privacy Sandbox adoption by 59.6K top, commercial websites & associated vendors
  - Focus on ad targeting: Protected Audience & Topics APIs
  - Source: Sincera is a start-up providing website metadata
- **Goal**: Descriptive research to inform
  - Academia: economics of innovation, web measurement in computer science
  - Public discussion: journalists, practitioners, & regulators
- Public dashboard at app.sincera.io/privacysandbox

Results: Protected Audience API Adoption

Protected Audience API (PA API): Key elements

- Site requests browser to add user to a **PA Interest Groups** in order to show related ads offsite
  - To protect user privacy, interest group membership never leaves the browser
- Site tells browser to host a **PA Auction** that selects a targeted ad based on the user's interest groups

PA Interest Groups over time

Vendor market shares (by site reach)

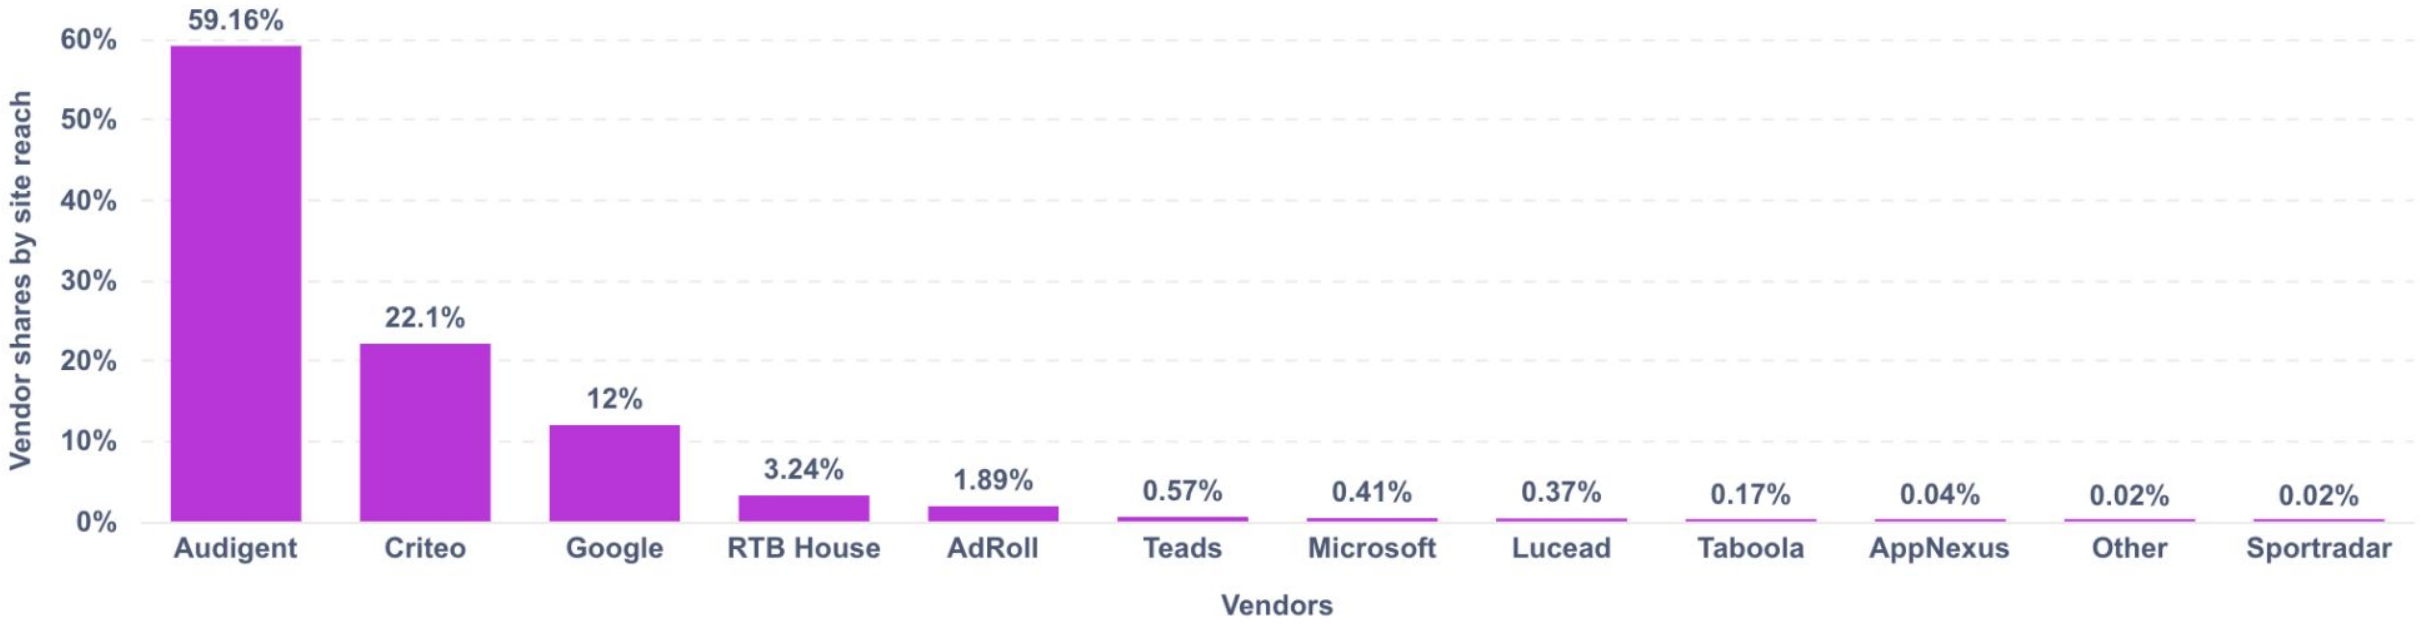

PA Auctions over time

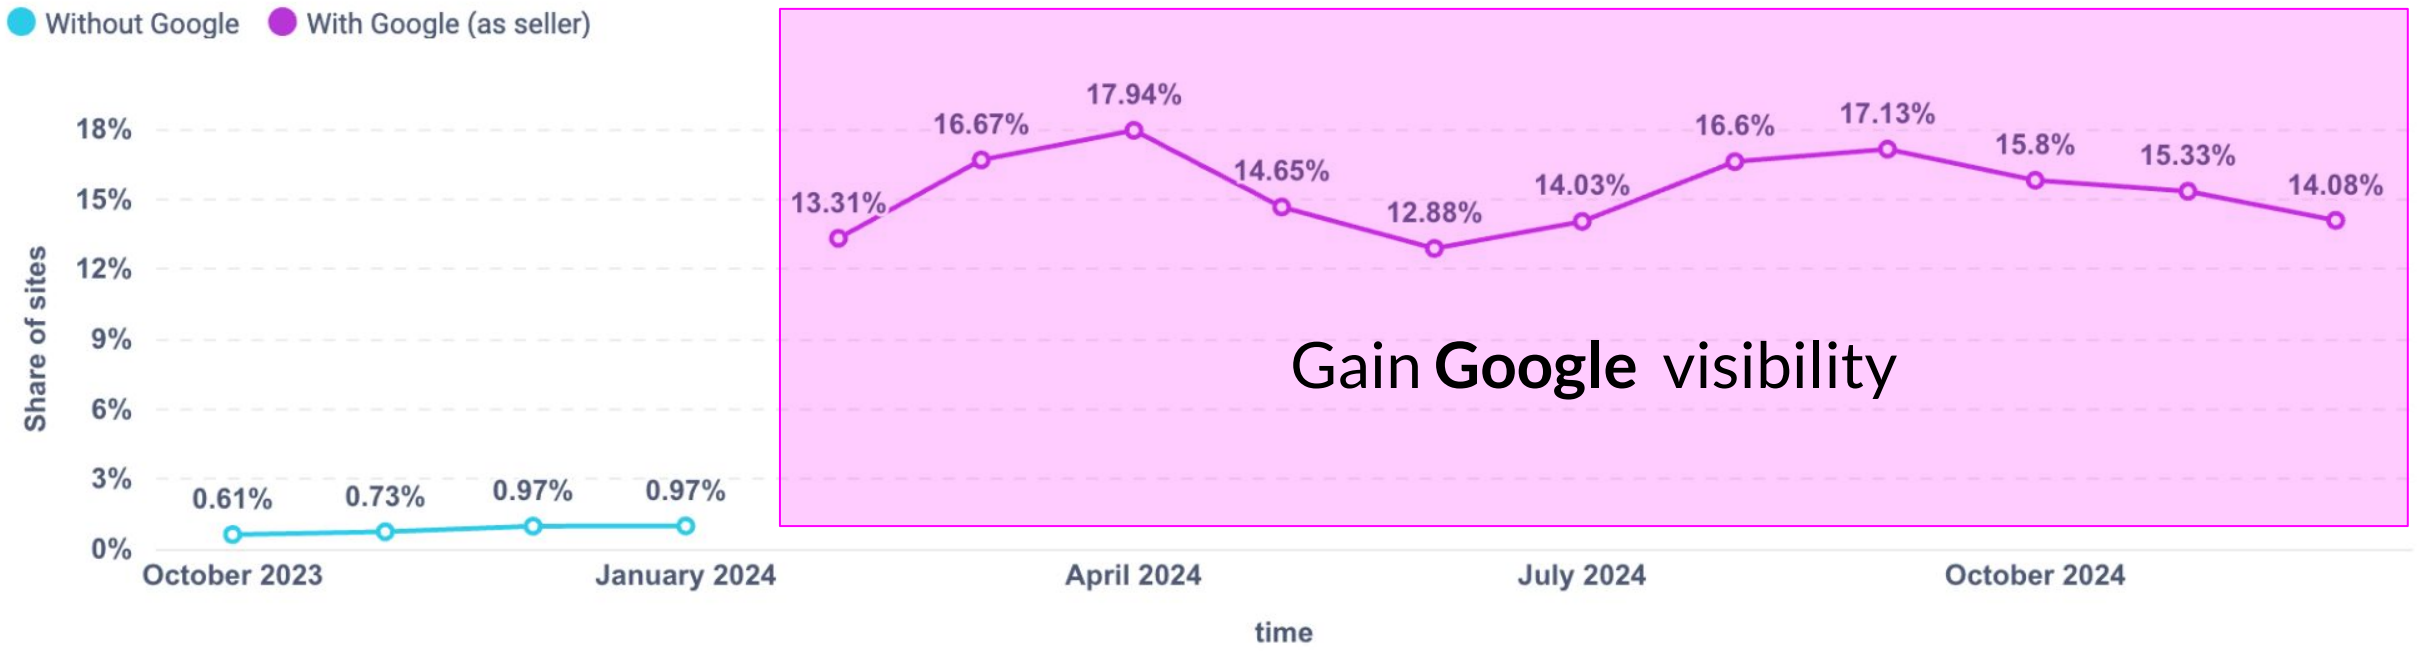

Vendor market more concentrated on auction side

Top-level seller market shares (by site reach)

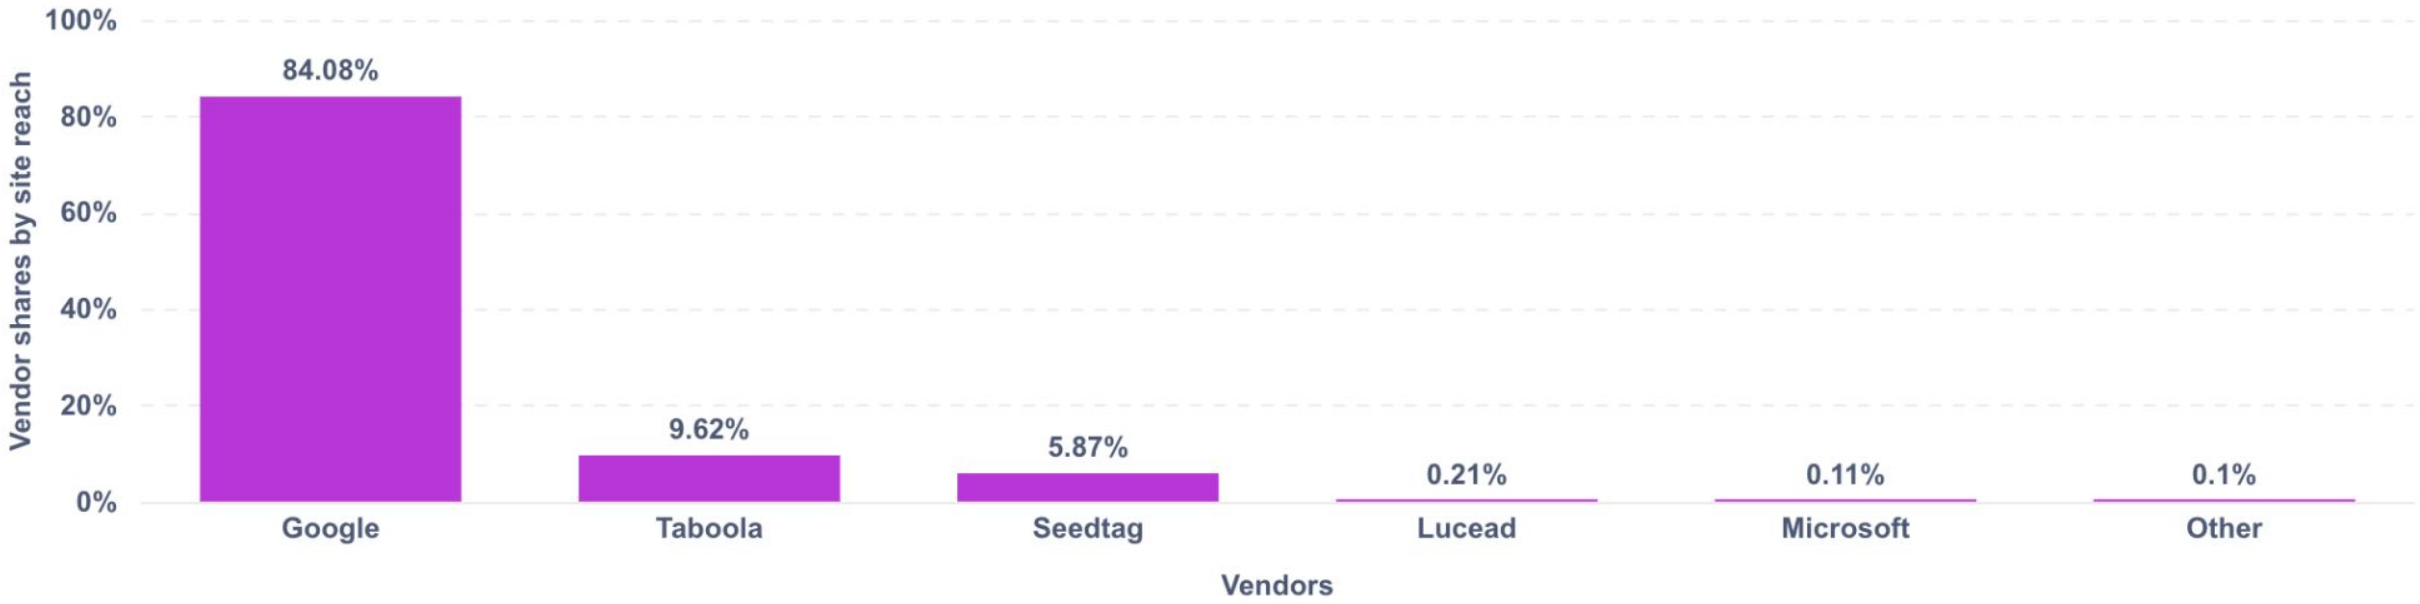

Results: Topics API Adoption

Topics API illustrated

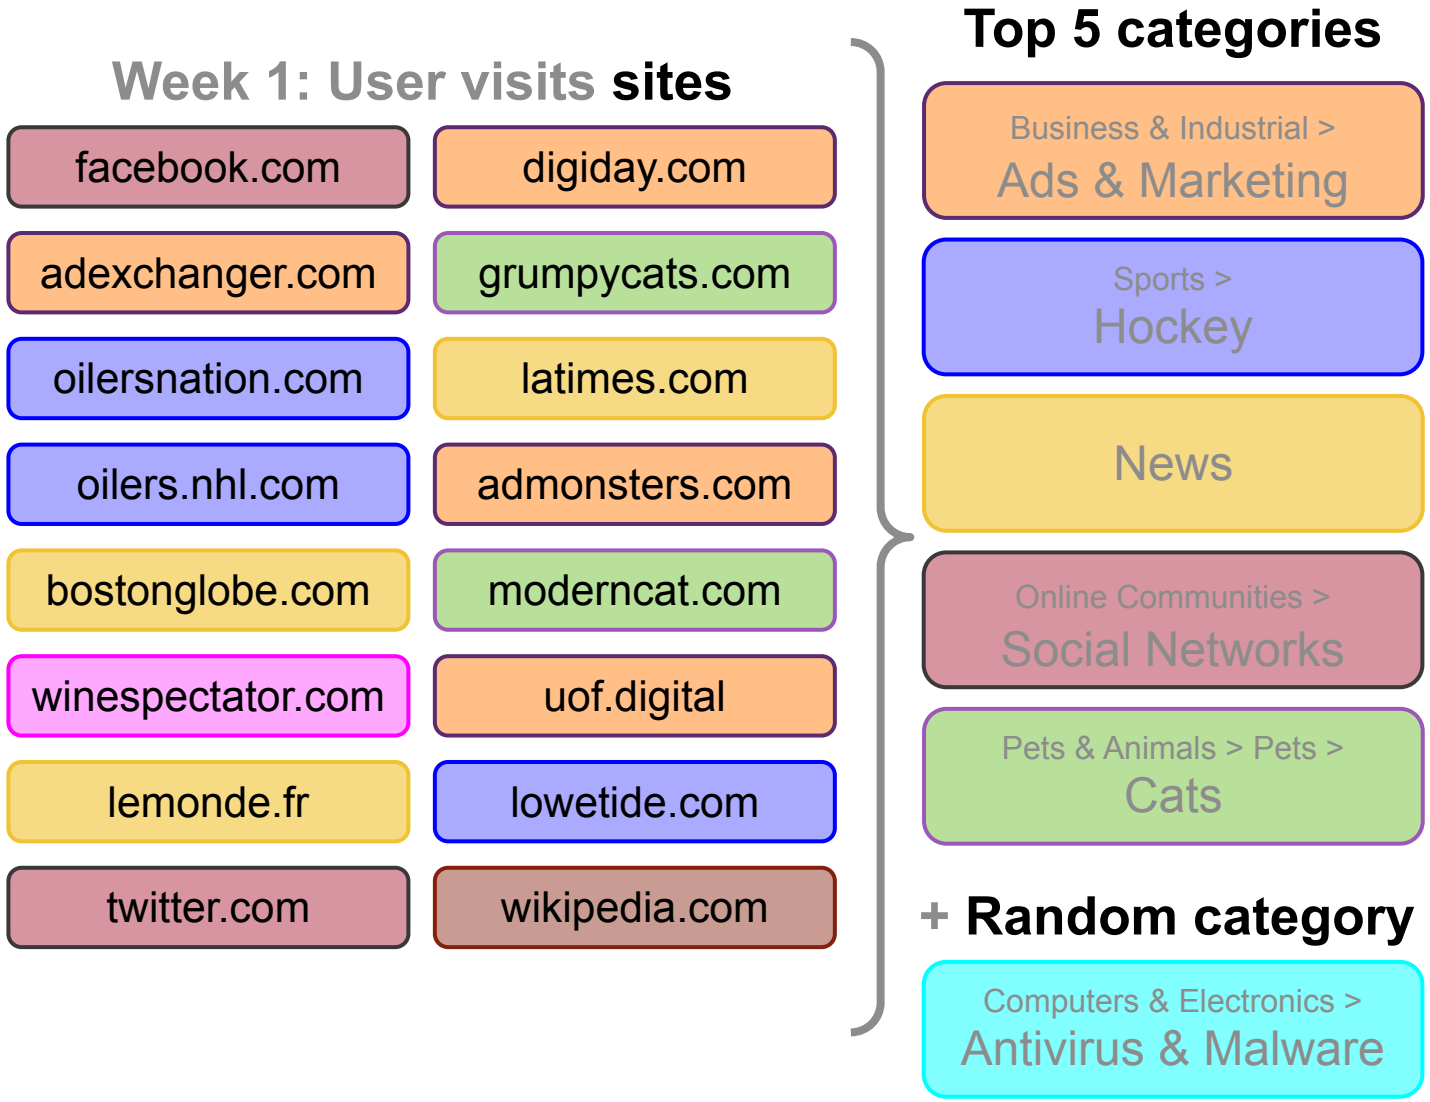

- 469 topic categories
  - Examples: Cats, Hockey
  - Avoids sensitive categories
- Each week, browser assigns user to 6 topics
  - Top 5 most visited topics (*after* sorting by commercial relevance)
  - 1 topic at random
- API returns up to 3 topics
  - 1 topic per user per site weekly for up to 3 weeks

Topics API calls dip after cookie deprecation pivot

Simplicity, incentives for wide coverage may favor Topics API use

Topics Calls: More vendors & less concentration

Vendor market shares (by site reach)

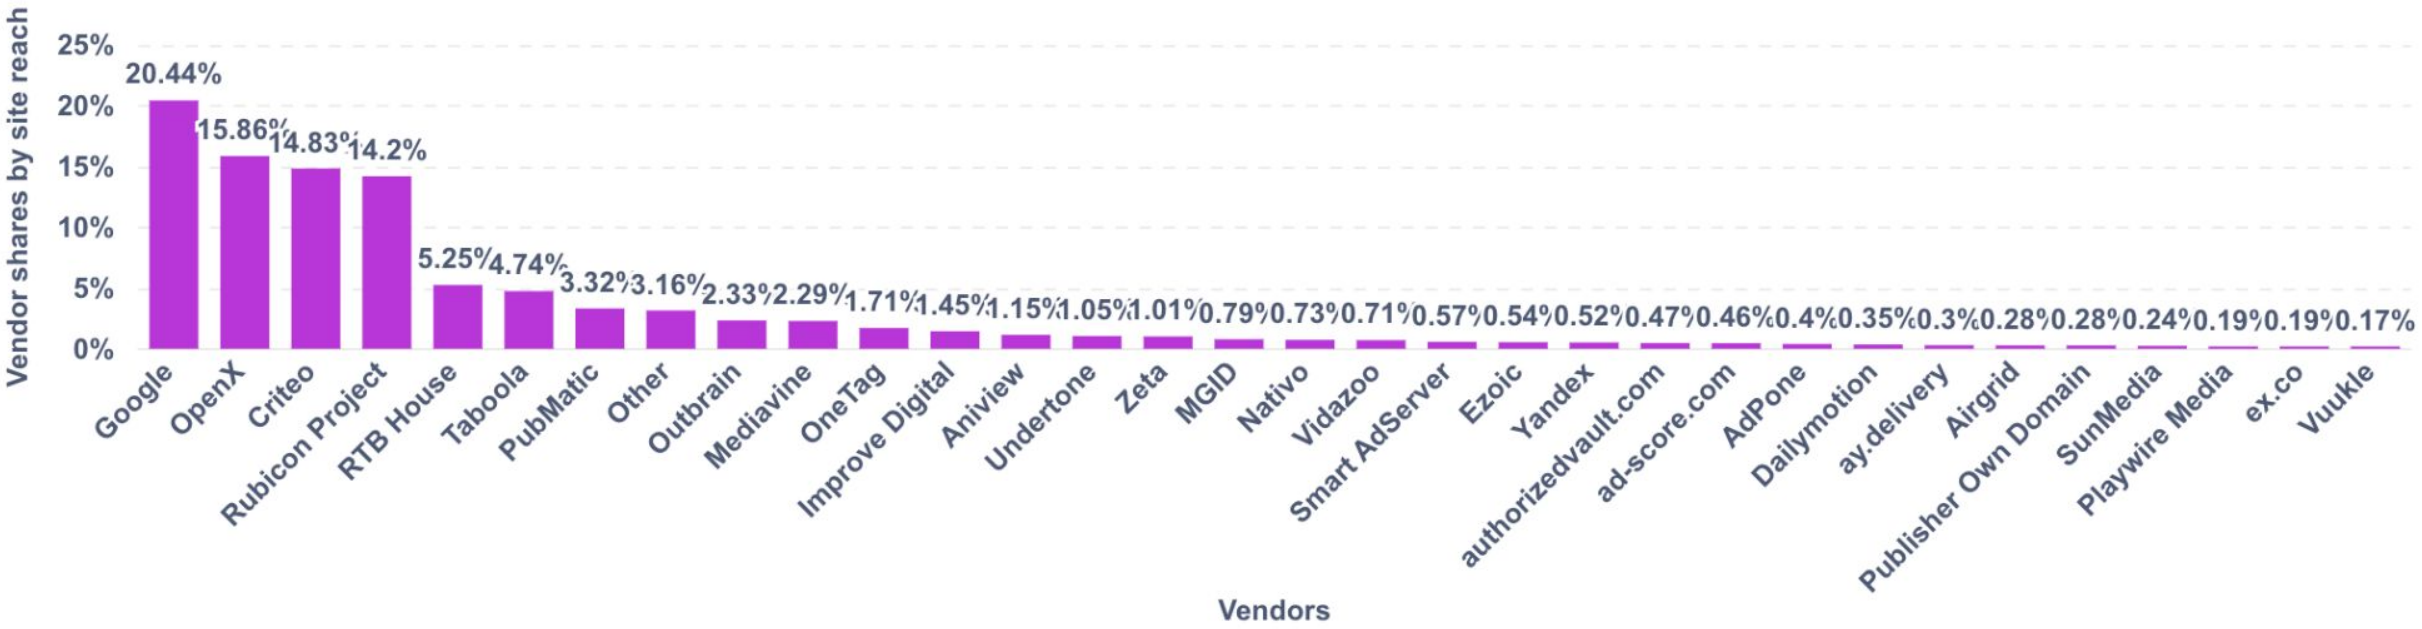

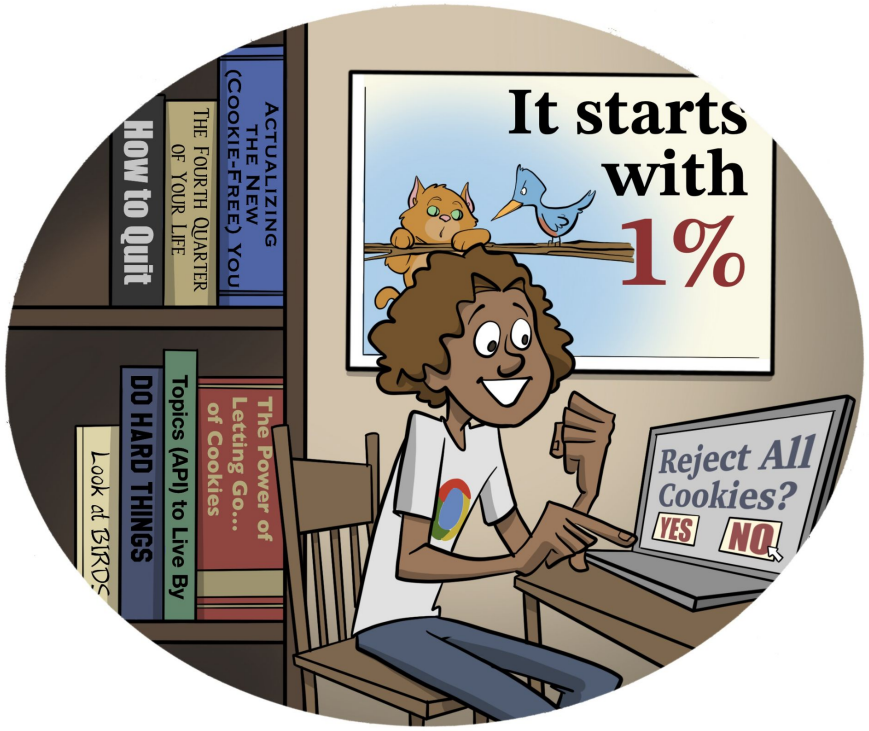

Privacy Sandbox experiment: Publisher results

"The Impact of Privacy-Enhanced Ad Technologies on Publisher Revenue: Evidence from an Industry-Wide Field Experiment" [work-in-progress]
Zhengrong Gu, Garrett Johnson, & Shunto Kobayashi

Supported by **The Center for Industry Self-Regulation**

Evaluating Privacy Sandbox in practice

- Ambitious, industry-wide experiment to help industry & regulators evaluate Privacy Sandbox's performance *ex ante*
  - *"Chicken & egg" problem*: Privacy Sandbox performance a function of industry adoption & investment in these tools
- We collaborate with **both** publisher- and advertiser-side partners
- Chrome users randomized into 3 main groups:

	3rd Party Cookies	Privacy Sandbox	Share
Status Quo	✔	✘	1%
Privacy Sandbox	✘	✔	0.75%
Cookieless	✘	✘	0.25%

Raptive: Publisher-side data source

- Raptive is the largest ad management company
  - Collectively, represent a top 10 digital property (Comscore)
  - 191M monthly site visitors, 5,200 US creators
  - Collectively, rank #1 in Food, Home, Lifestyle, and Family and #2 in Travel categories
- **Data**: Revenue per impression tied to experimental labels
- Obtain full, unselected view of publisher monetization
  - Vendors in the ad selling chain see selected subset of market
    - Selection varies by experimental group, violating experimental validity

Sandbox recovers 11% of lost revenue in August

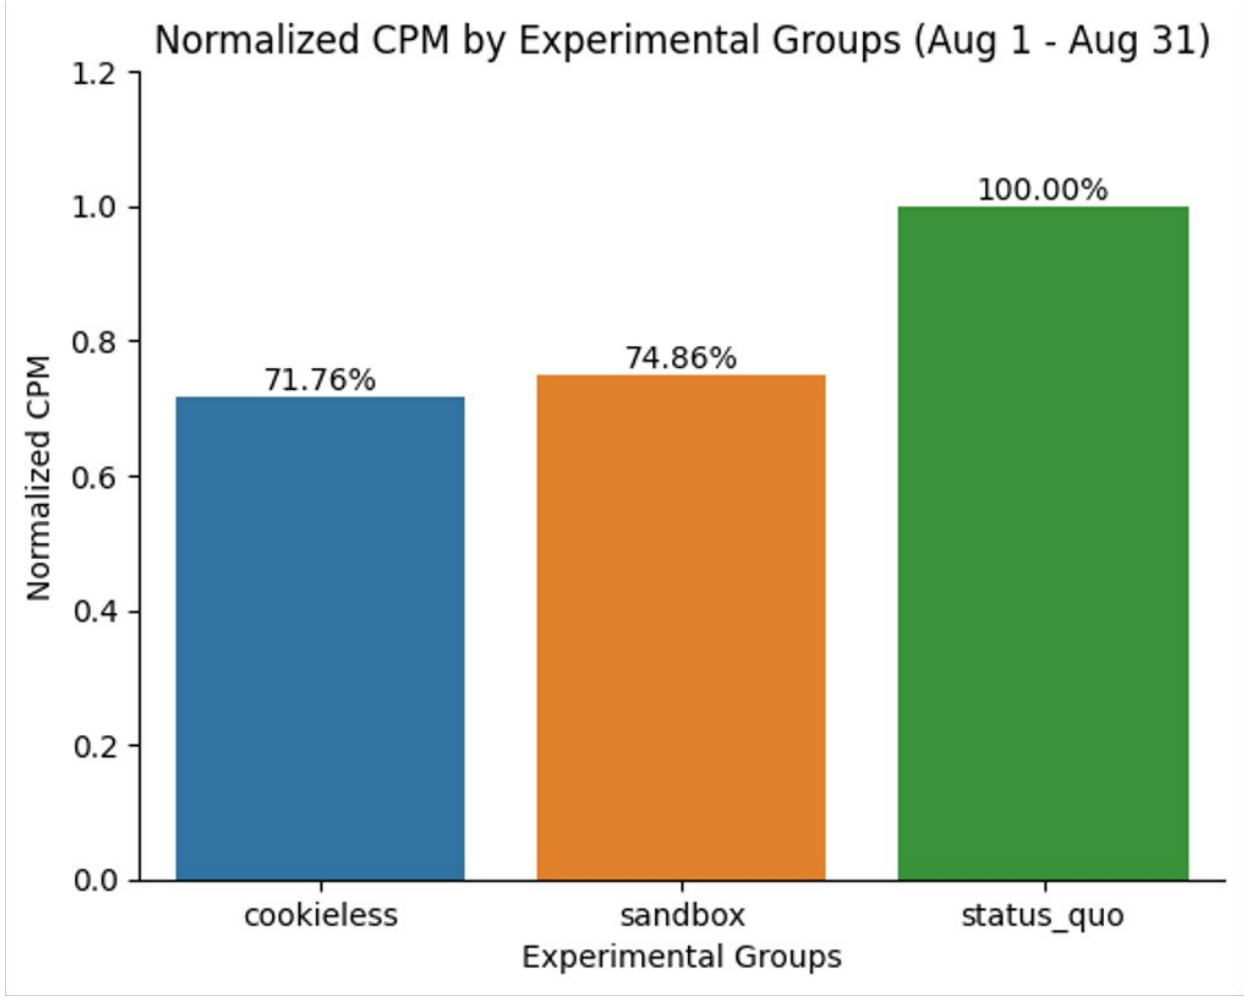

- Relative to status quo, average ad prices *fall*:
  - 28% without cookies
  - 25% with Privacy sandbox

- **"Recovery share"** is 11%: i.e., Sandbox recovers **11%** (3%/28%) of revenue lost from losing cookies

PA API increases latency in ad loads (June 7-15 data)

Note: Longer latency in Sandbox group regardless of whether ultimately sold via PA API

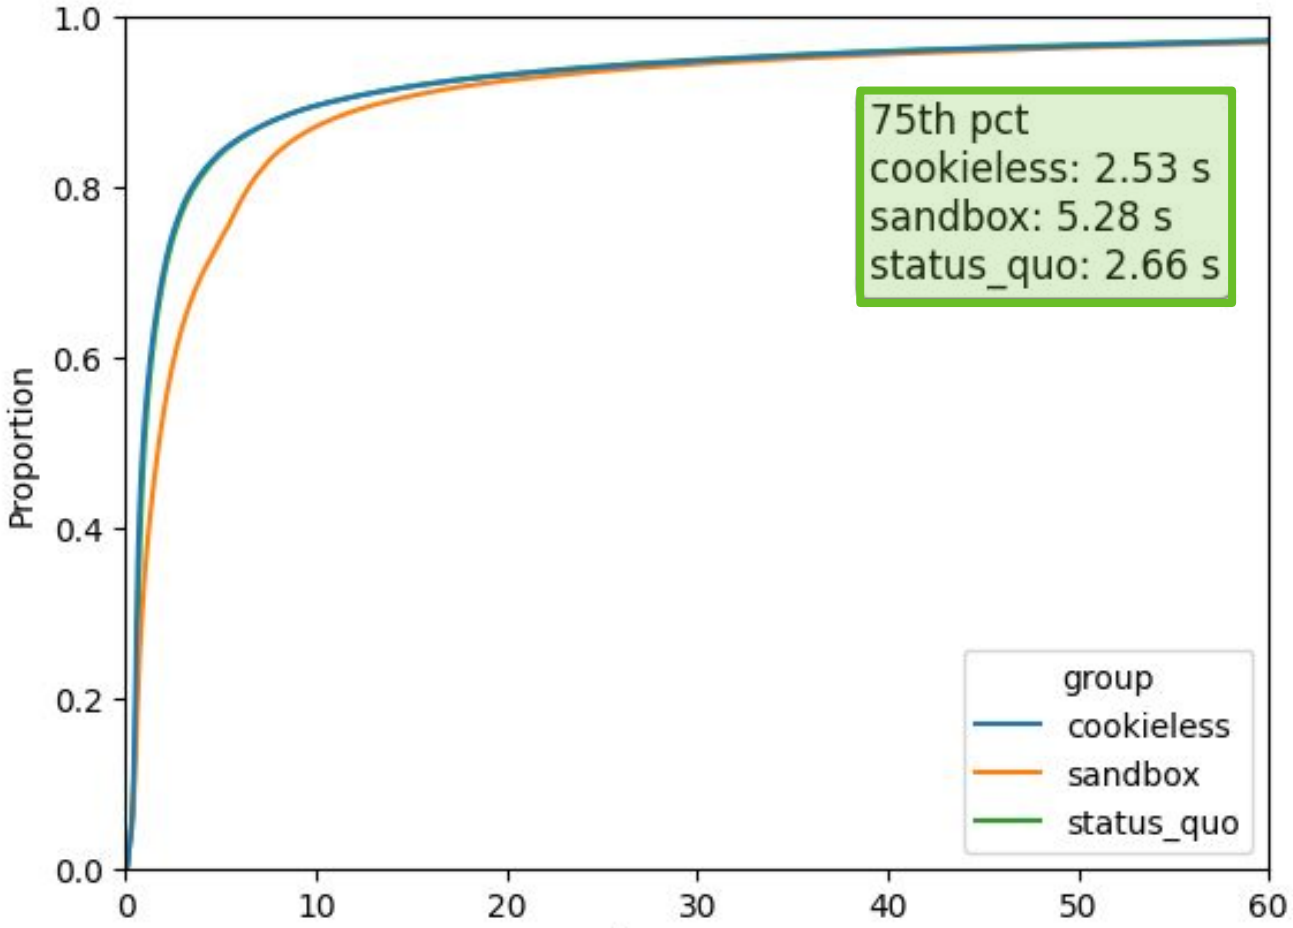

Latency (sec) by group

PA API seems to prevent some ads from loading

- If ad latency is too high, then ads won't load and publishers earn no revenue
- We estimate that Privacy Sandbox reduces ad impressions by 3.1%
  - How? Privacy Sandbox group has 3.1% fewer impressions than we expect given the impressions in other groups

- → Our earlier analysis overestimates Privacy Sandbox's benefit, because it omits that 3.1% of impressions yield 0 revenue
- → Revised recovery share falls to 2.8% (rather than 11%)
  - Revised revenue = 74.86*96.9% + 0*3.1%=72.54
  - Revised recovery share = (72.54-71.76)/(100-71.76)


PRIVACY
SANDBOX
RECOVERY SHARE
PUBLISHERS
RECOVERY
SHARE
FEWER IMPS
RENDERED
PUBLISHERS
imgflip.com

Privacy Sandbox experiment: Advertiser results

"Privacy-Enhanced versus Traditional Retargeting: Ad Effectiveness in an Industry-Wide Field Experiment"

Shunto Kobayashi, Garrett Johnson, & Zhengrong Gu

Supported by **NET Institute**

Advertiser performance in Privacy Sandbox

"Privacy-Enhanced versus Traditional Retargeting: Ad Effectiveness in an Industry-Wide Field Experiment"

- **Vendor partner:** Demand Side Platform (DSP) representing thousands of advertisers
- **Data**: Weekly total advertiser outcomes (clicks, click-through conversions, & conversions) by three experimental groups
  - "Google Analytics"-style data obtained by vendor's direct integration with advertiser's conversion tracking
  - >2K advertisers globally using retargeting campaigns
- **Advantages**
  - Observe all advertiser data rather than vendor's selected view
  - Avoids "apples-to-oranges" comparison of cookie- vs. sandbox-based measurement tools
  - Yields ad incrementality estimates (albeit imprecise)

Clicks fall -88.3% without cookies

"ID bridging" effectively preserves some cookie-based targeting in the short run

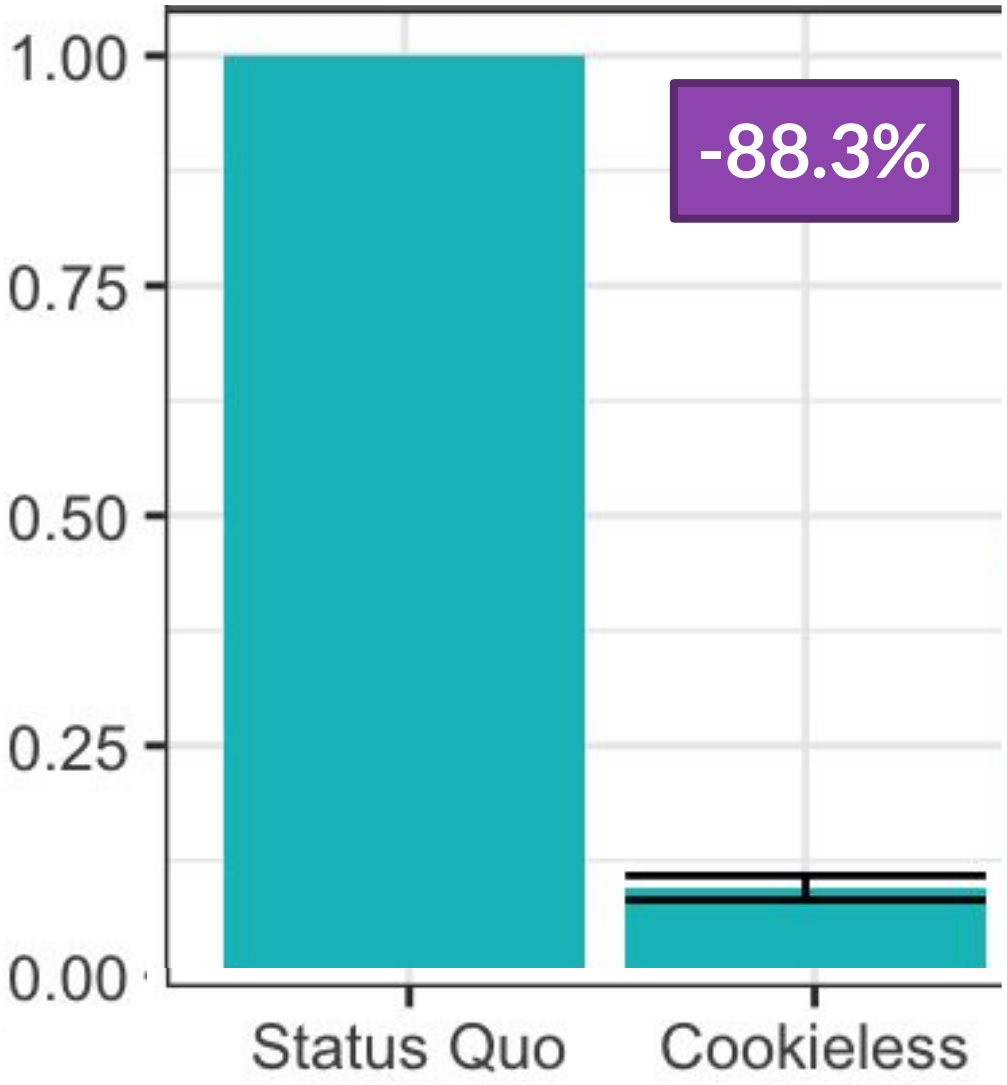

Privacy Sandbox

Note: Figures come from post-March data, but numbers come from full sample, so they slightly diverge—sorry!

Clicks: 46.3% recovery share with Privacy Sandbox

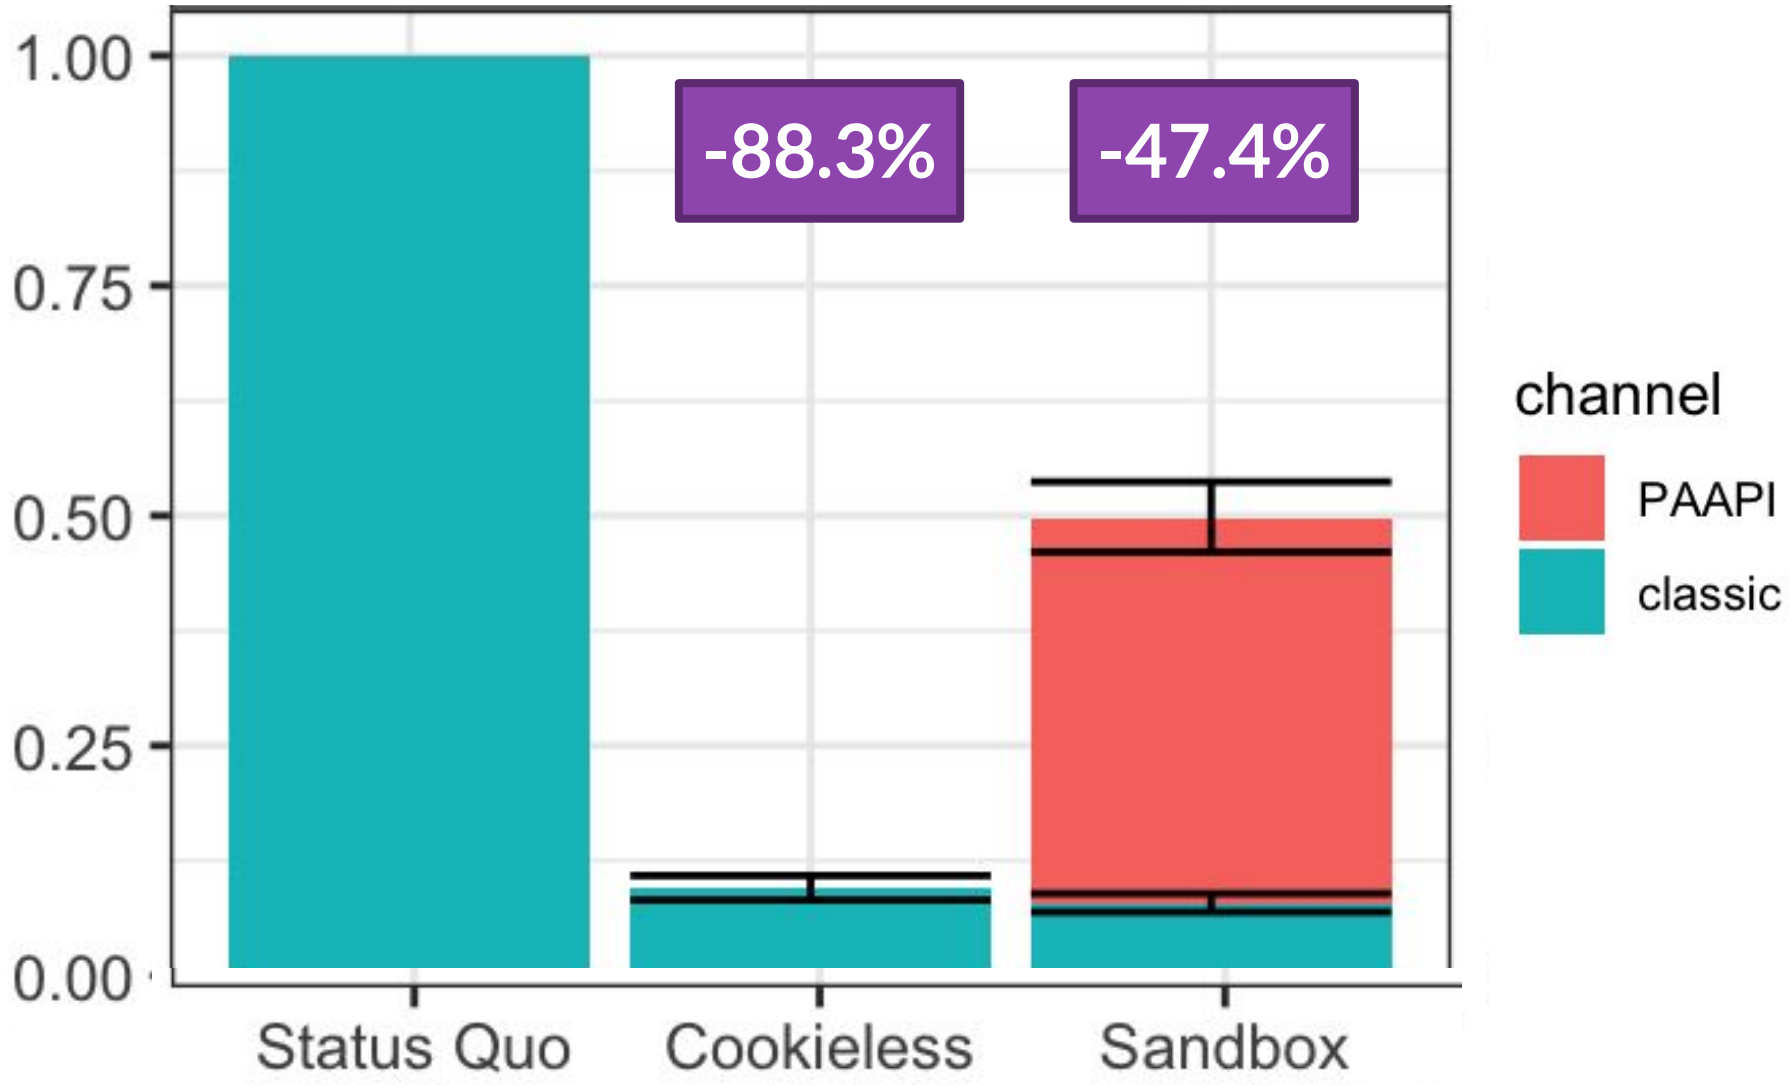

Note: Figures come from post-March data, but numbers come from full sample, so they slightly diverge—sorry!

Click-through conversions: 43.5% recovery share

Click-through conversion is a key onsite outcome (e.g., purchase) that originates with an ad click

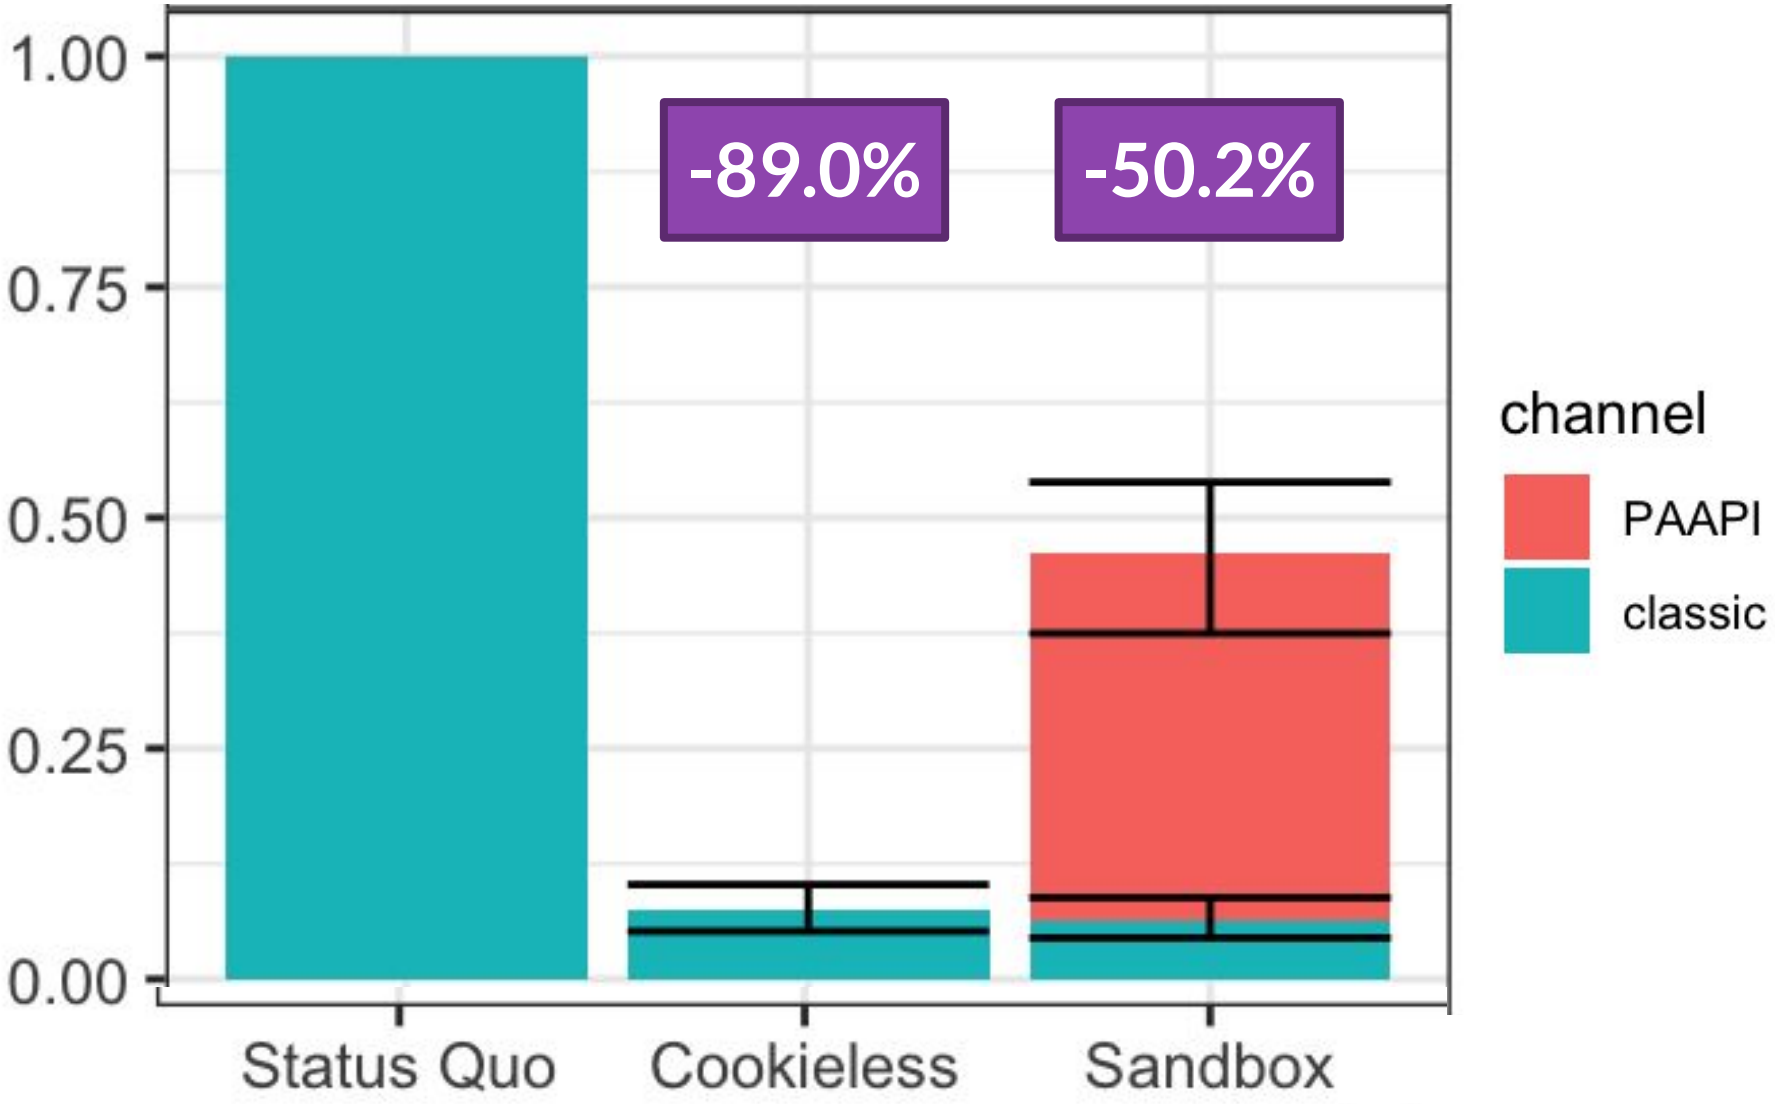

Privacy Sandbox

Note: Figures come from post-March data, but numbers come from full sample, so they slightly diverge—sorry!

Lower ad count & spend in Sandbox vs. Status Quo

Data from weeks 10+. Ad count and spend became accurately observed from week 10 (March 4th)

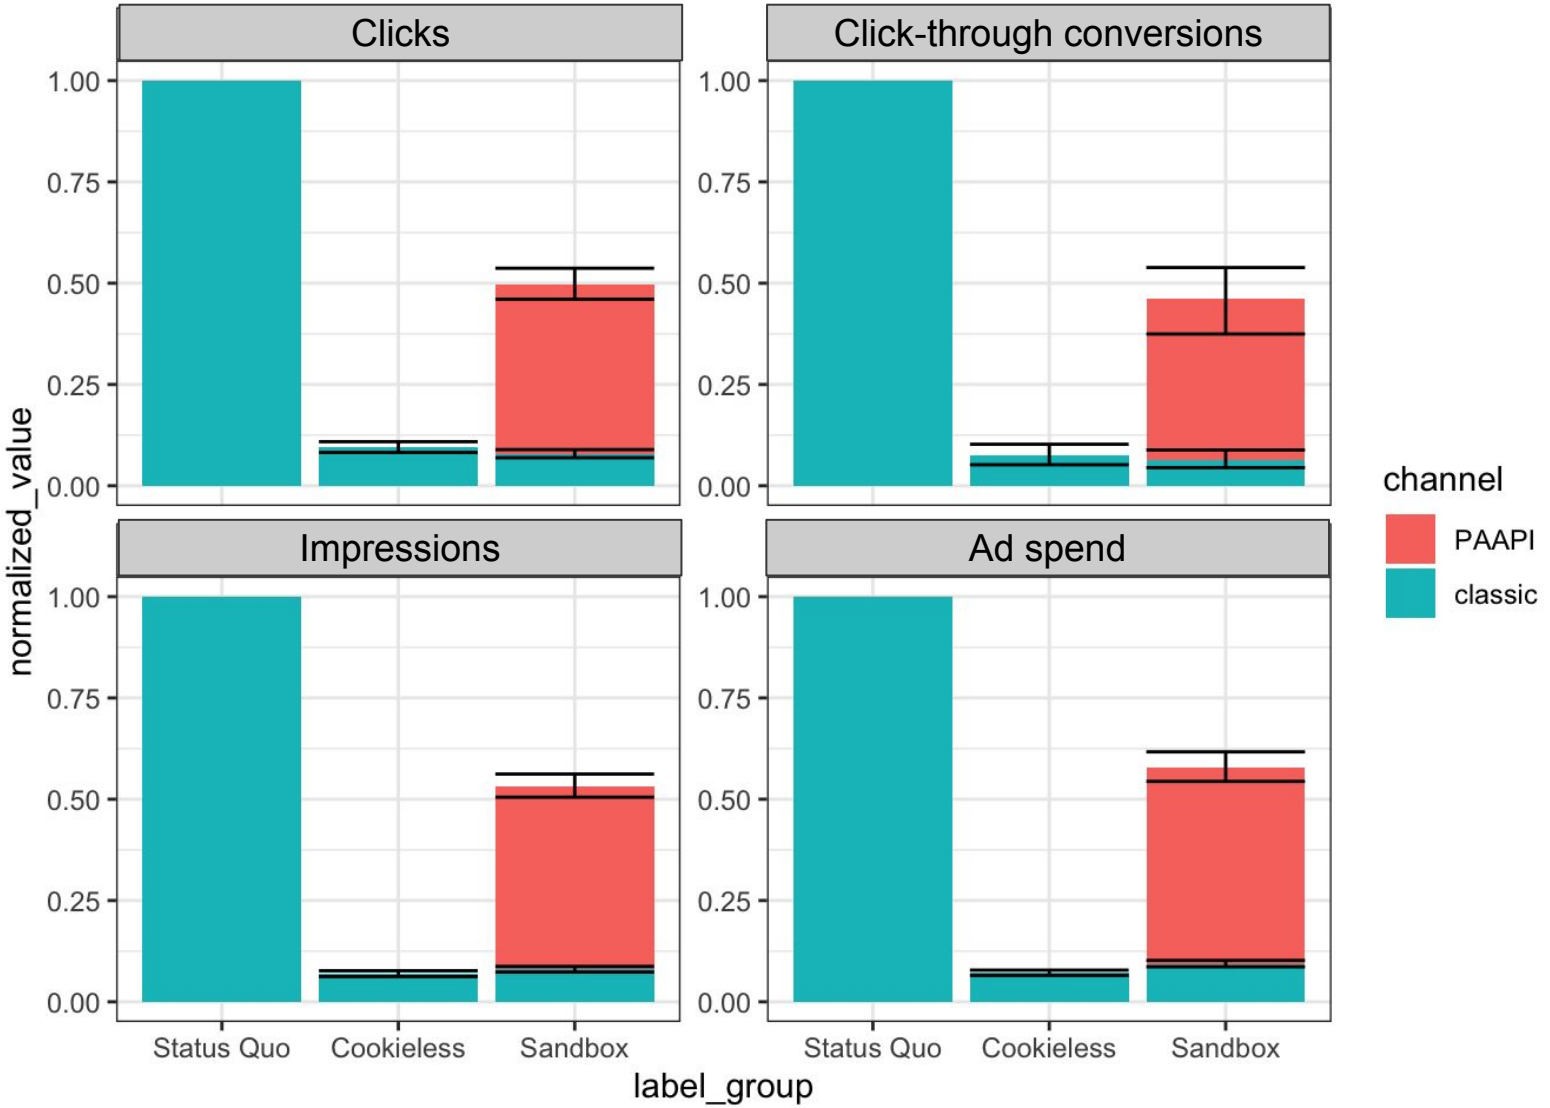

Why lower Sandbox spend?

- **Low supply of PA API inventory** (in H1 '24) at comparable prices seems to place ceiling on recovery share

Similar efficacy per dollar/impression

Ratio of sample average outcomes (weeks 10+)

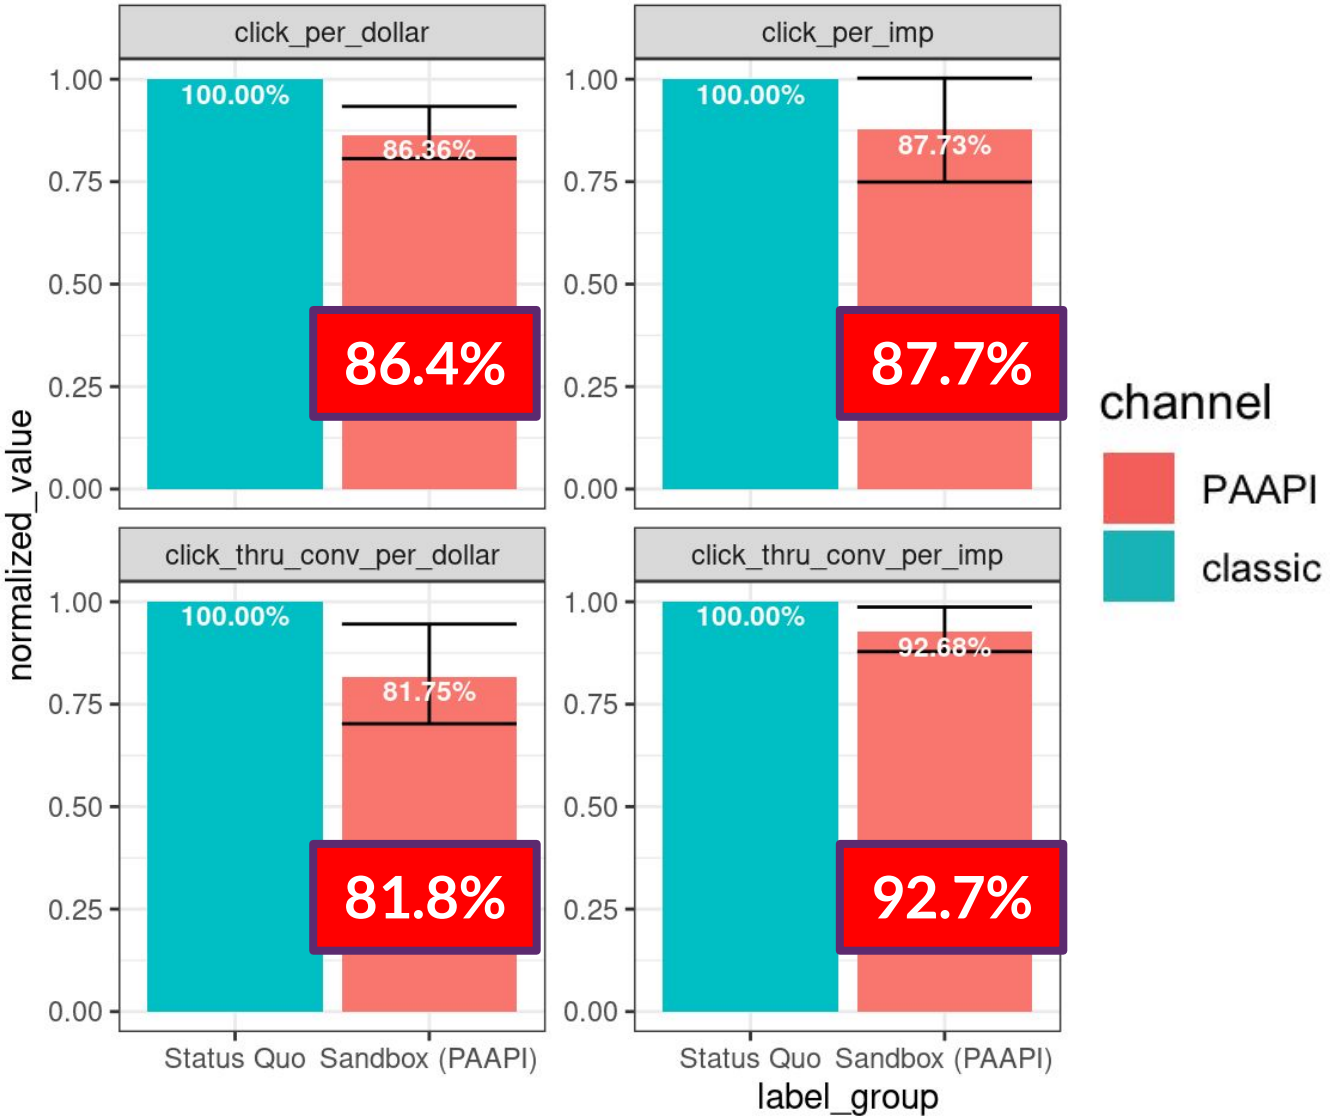

- Similar efficacy appears to be good news for advertisers
  - Further investment in Sandbox tools could help
  - By extension, more advertiser adoption would improve publisher revenue
- BUT, we do not observe the performance and adaptation of advertisers' campaigns that do not use retargeting

Retargeting effectiveness

- Retargeting effectiveness is **controversial**: retargeting targets favorably selected users and may provoke user privacy concern
  - Prior literature provides **case studies** (Sahni et al. '19; Johnson et al. '17)
- We estimate retargeting effectiveness **across 2K+ advertisers** by comparing Cookieless and Status Quo groups
- **Result**: Retargeting increases baseline conversions by >4.6%
  - **Note:** this is an *intent-to-treat* estimate on all site visitors rather than the subset that advertisers retarget

Conclusion

Conclusion

- Privacy-enhancing technologies are an exciting domain for economists to study
- Privacy Sandbox gained adoption among firms though this waned somewhat with Google's recent pivot on cookies
- Privacy Sandbox experiment show:
  - Advertisers can obtain similar ad performance per dollar/impression with Privacy Sandbox
  - Privacy Sandbox modestly improves publisher revenue, albeit somewhat undercut by significant latency issues
  - For both, soft market adoption restricts market performance
- Significant scope for increased industry adoption & adaptation to improve publisher- & advertiser-side outcomes

Thank you for your questions and comments!!

Topics opt-out rate falls over time

In July 2022, rate was much higher: bounded between 2.7% and 3.9%

Topics API illustrated

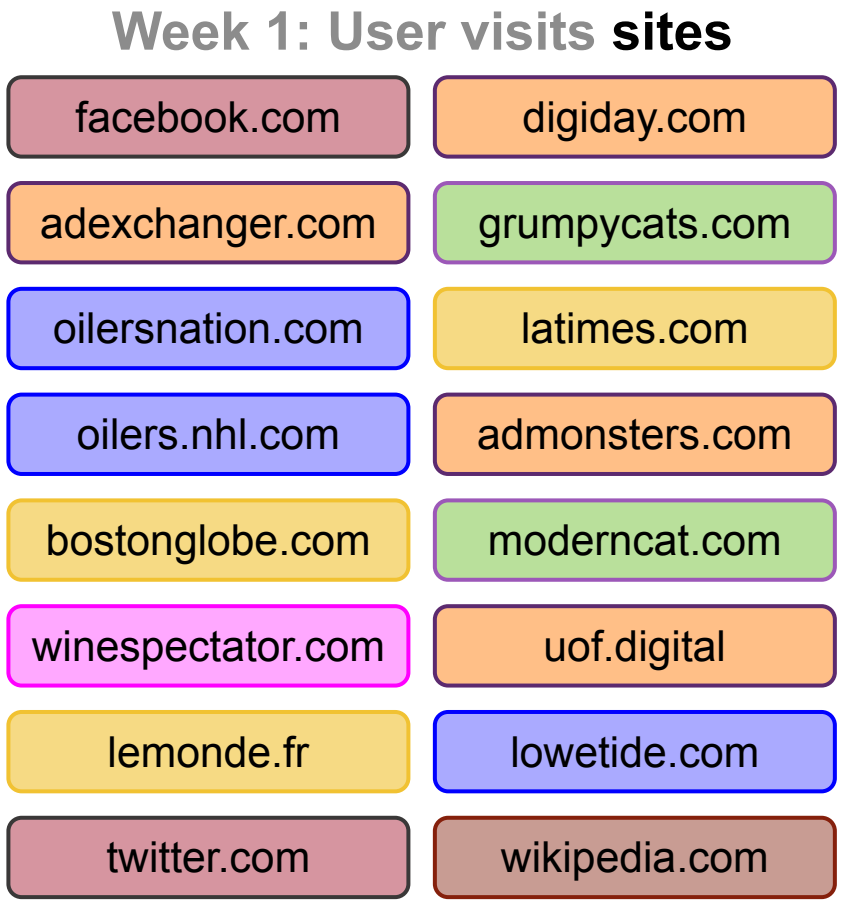

- Topics API enables interest-based advertising
- Browser classifies the topical category of domains that users visit
- Browser shares topic categories with ad tech vendors
  - Helps advertisers prospect & optimize ads
- Category generated from (sub)domain alone
  - e.g., tennis.site.com > Tennis
  - site.com/tennis > [topic of site.com]
- Privacy?
  - Users are "hidden in the crowd"
  - Users can opt out
  - Sites can opt out

Recovery share is low throughout experiment

Before August, this omits some PA API revenue (we think this has little impact on these results)

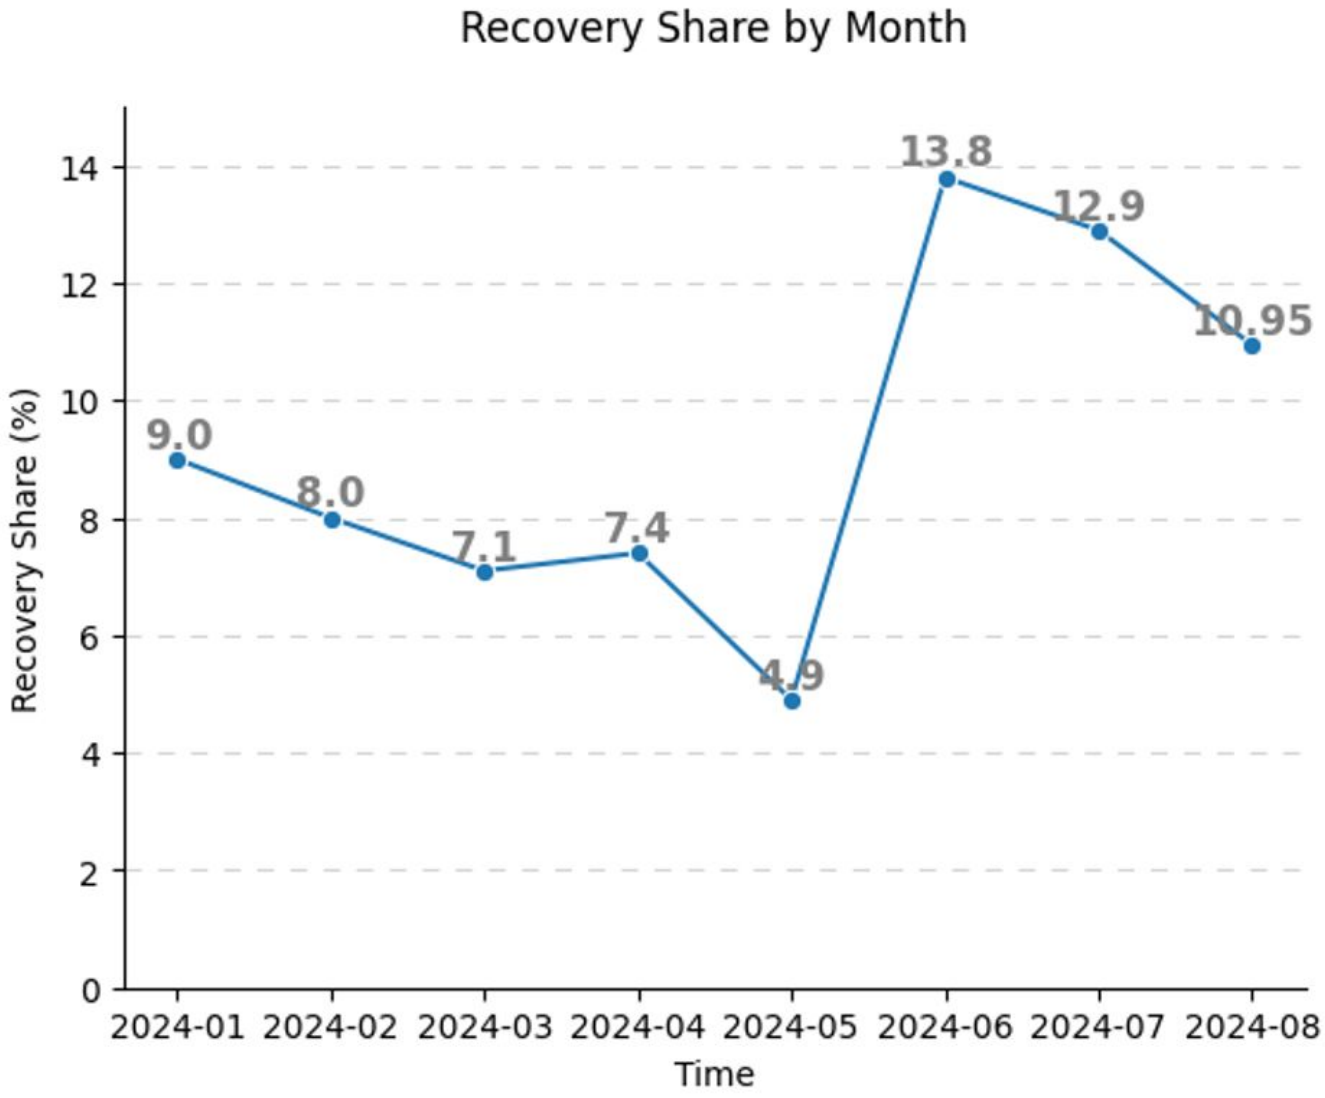

PA API increases latency in ad loads (June 7-15 data)

Note: Longer latency in Sandbox group regardless of whether ultimately sold via PA API

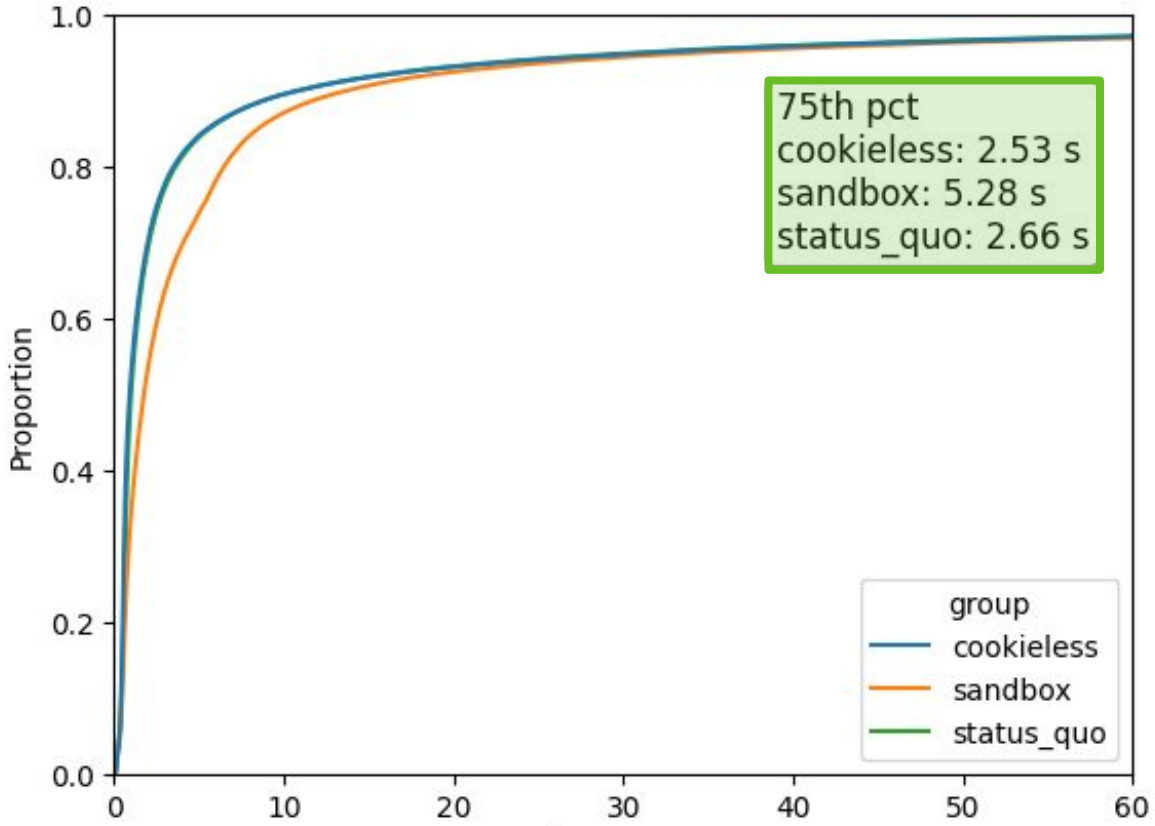

Latency (sec) by group

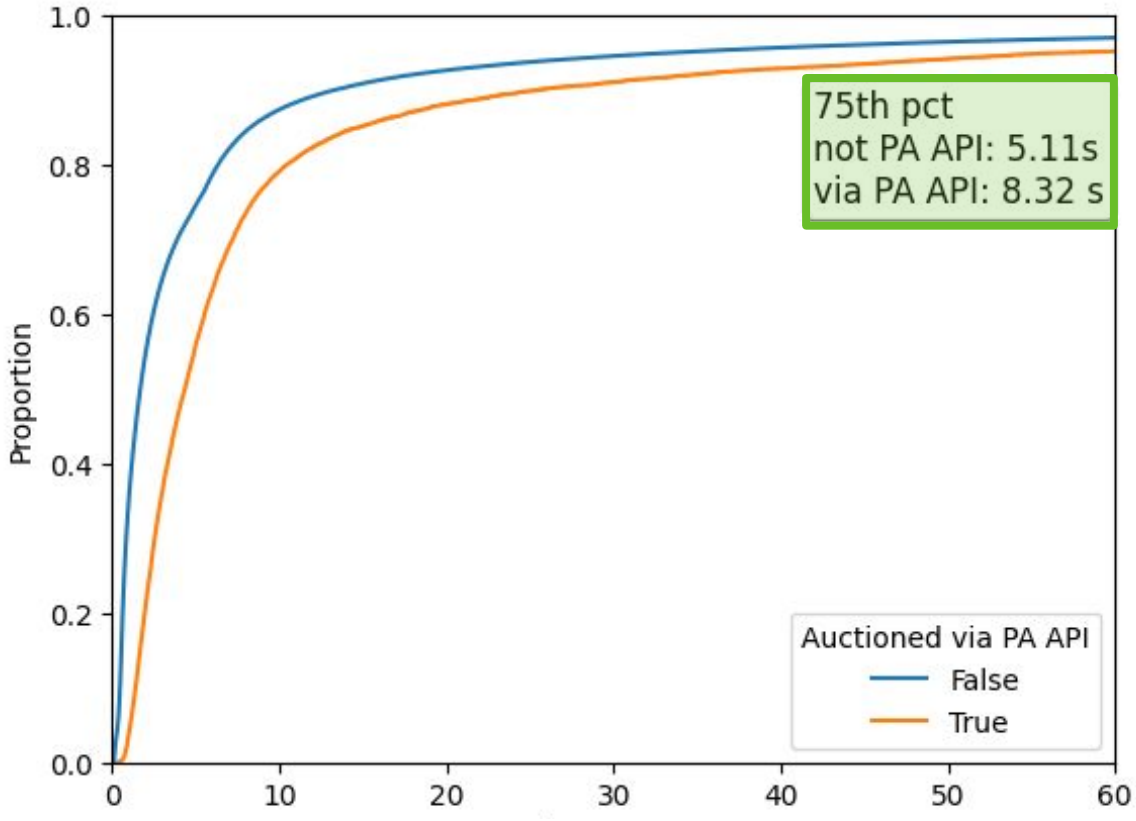

Latency (sec) by PA API (within Sandbox group)

~20-50% "recovery share" for advertisers

Full sample: weeks with active advertiser data sharing and PA API spend >0

	Cookieless vs Status quo	Sandbox vs Status quo	Recovery share
Ad clicks	-88.3%	-47.4%	46.3%
Click-through conversions	-89.0%	-50.2%	43.5%
Conversions	-4.38%	-3.60%	18.7% (noisy)

Cookieless reveals scale & decay rate of ID bridging

Sample fixed to initial 1.2K advertisers.

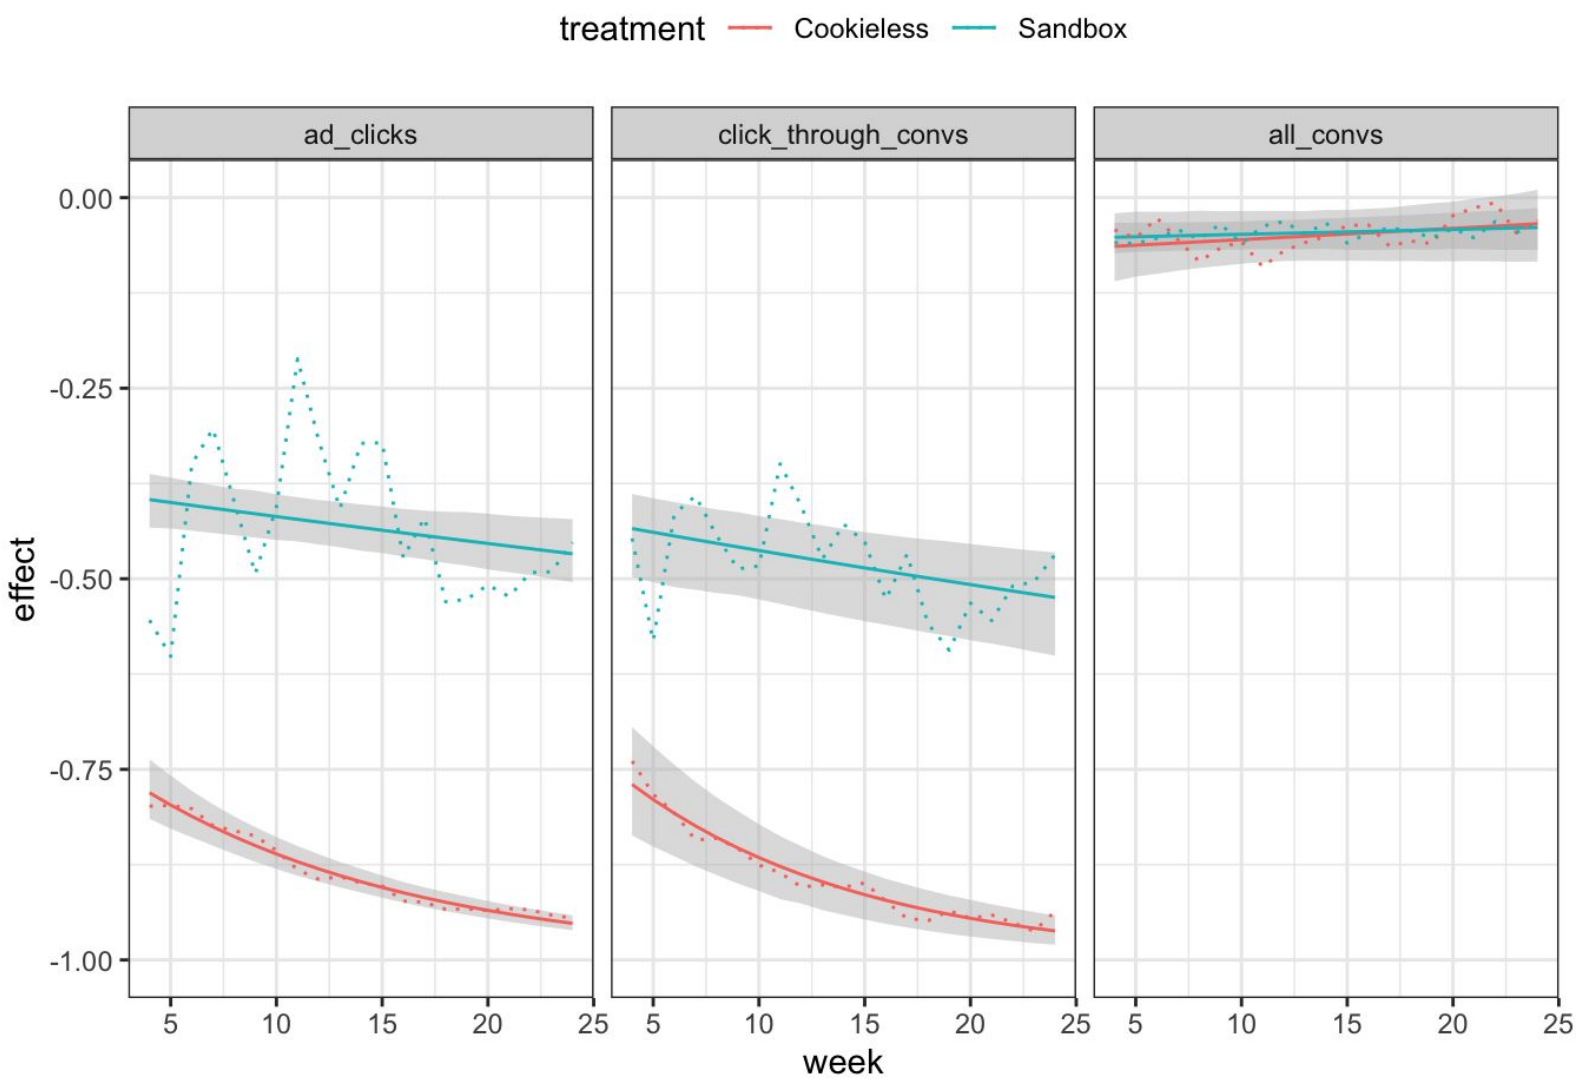

Dotted: Per-week estimate
Solid: Poisson QMLE

Bootstrapped with advertisers as clusters