

What the Machines Can See

The state of mobile app growth in 2026. The engines learn from what they can see early and claim. Durable value depends on what the reports leave out: what would have happened anyway, what the store keeps, and what a customer is worth after a year.

By **Evgeny Popov** · Published by No Fluff Advisory · US-first, with global context

Research cutoff: 27 September 2026 · Information after this date is not used

621 distinct source addresses consulted	272 findings recorded; 272 re-checked against their sources	33 headline claims in the ledger	37 / 64 companies / products in the vendor universe	11 + 5 execution products and benchmark channels scored twice	119 logged corrections
--	---	---	--	---	----------------------------------

ABSTRACT

What this paper found

App-growth engines are built to learn from what they can see early and claim. In 2025 the money went where value per user was higher, even where the engines see less. Durable growth depends on what the reports do not show: what would have happened anyway, what the store keeps, and what a customer is worth after a year.

This paper asks one question: **what produces extra durable app value, for which buyers, under which limits, and at what cost?** It answers from the US market first, with global evidence where US data runs out. It keeps apart things the industry often blends: installs and customers, credit and cause, forecast and realized value, and a product's existence and its effect.

The research ran as ten evidence streams, three gap-fill passes and a review of the author's own work. Five fact-check passes then re-opened the sources behind every recorded finding, two separate passes scored the vendor panel, a hostile reviewer attacked the argument, and an auditor rechecked the numbers. ChatGPT (OpenAI) then reviewed three builds, and each of its factual points was checked against the sources. It consulted 621 distinct source addresses and recorded 272 findings. The 33 headline claims each carry a record that links prose, evidence and calculation. Chapter 12 describes the process as it actually ran, including what it could not reach.

Six findings drive the rest.

01

The engines learn from what they see; the money does not simply follow.

Value bidding needs dense, early signal: Google's app campaigns need at least 10 conversions a day to bid for a return target. Yet in 2025 acquisition spend grew 35% on iOS, where signal is thinnest, and fell 1% on Android. It grew 18% in non-game apps and 3% in games. Value per user, category mix and a surge of China-based shopping budgets explain where money went at least as well as signal does CL-016.

02

Remarketing grew faster than acquisition, and we found no independent test of it in apps.

It rose 37% to \$31 billion in 2025, against 13% for acquisition, and made up 29% of app marketing spend in AppsFlyer's panel estimate. Randomized studies of web retargeting find real but modest lift. We did not locate a public independent randomized app study that validates its attributed returns [CL-005](#).

03

The specialist with the highest adjusted EBITDA margin in our table owns an auction its model learns from, but ownership alone does not explain the lead.

AppLovin reported \$5.5 billion of net revenue in 2025 at an 82% adjusted EBITDA margin, and its model uses outcomes from its own mediation auction [CL-007](#). Liftoff, with no such auction, reports a 55% margin on the same net basis, and Unity's rebound followed a new model more than a new asset. Google and Apple own more of the loop than any specialist and disclose less.

04

Credit is not cause, and the tests are gated.

Across 663 randomized test-and-control comparisons on Meta, statistical methods run without control groups overstated median lift three to thirteen times [CL-012](#). The lift tests we reviewed at the big platforms run through account teams [CL-014](#). Our searches found no independent causal study of streaming TV's effect on app installs [CL-026](#).

05

The store's cut is now a lever.

At Stripe's published card rate, web checkout keeps about 34% more of a first-year subscription payment than Apple's 30% tier. Against the 15% tier it keeps about 11% more. Both figures come before the cost of running your own billing [CL-022](#). How much Apple may charge on US link-outs was still being set in court at the cutoff [CL-021](#).

06

"Apps DSP" is a checklist, not a proven category, and agents move the control point.

On our own screen, one of eleven products documents the full set of capabilities, and none shows the bundle wins [CL-029](#). Meta's and TikTok's agent servers can create and change campaigns. Meta's creates new campaigns paused. Neither documents an approval step for edits to live campaigns, so buyers have to set limits, logs and data rights themselves [CL-030](#).

HOW TO READ THIS

Twelve chapters and an evidence appendix. In the HTML edition, the **Skim** control collapses the paper to its summaries, figures and instruments, and / searches the text, claims and sources. Claim chips such as [CL-012](#) open the claim's record; numbered references open the source. Chapter 9 holds the scored vendor matrix with full profiles. Chapters 5, 7 and 10 contain interactive models. They are teaching tools with synthetic numbers, and they are labelled that way.

INTERESTS TO DECLARE

The author, Evgeny Popov, has worked at **Samba TV** since May 2025 as Global Head of Enterprise. Samba TV sells TV viewing data and measurement, and it appears in chapter 8 as the data source in one CTV-to-app method. From September 2022 to August 2024 he was EVP and GM for international growth at **Verve Group**, whose Verve Dataseat app DSP is scored in chapter 9. He co-leads a working group within the **Ad Context Protocol**, an agentic-advertising standard discussed in chapter 11. He founded **No Fluff Advisory**, which publishes this paper and sells a Mobile App Growth advisory playbook that chapters 3, 10 and 11 test and in places correct. Earlier he held senior roles at Hearts & Science (Omnicom) and Lotame.

No company evaluated in this paper received a prepublication draft or exercised editorial control. Samba TV, Verve and the Ad Context Protocol were held to the same evidence rules as everyone else.

LIMITS TO KEEP IN MIND

Most causal evidence here comes from web, search, retail and audio advertising, not app installs; the transfer is an inference. Market sizes rest on vendor panels, and the fastest growth in them is global and partly driven by China-based e-commerce budgets. Vendor capabilities are scored on public documents, which measure what firms disclose, not how well they perform. The research used AI agents for search and verification under the author's direction. ChatGPT (OpenAI) assisted with technical and editorial review and source verification. The author retains responsibility for the paper. This assistance does not constitute independent assurance or endorsement by OpenAI. It has not been peer reviewed or independently audited.

CHAPTER 1

Define growth before you count it

An install is an event. A customer is a person who stays and pays. Much of the muddle in app growth starts when the two are counted as one.

IN SHORT

App growth has six stages: discover, acquire, activate, retain, reactivate and monetize. Paid media sets bids, audiences, placements and ads. It does not set onboarding, price, the paywall or the store fee, and those decide whether an install turns into profit. This paper counts growth as extra contribution, in total and per new customer. It does not count installs as growth.


Six stages, five owners

Start with a plain map. A person hears about an app. They tap an ad or a search result, land on a store page and install. They open the app, find some value, and come back or leave. Some pay. Some return later. Six stages cover that path: discovery, acquisition, activation, retention, reactivation and monetization.

Five groups hold the levers. Ad platforms set bids, audiences, placements and delivery. App stores run the listing, search, ranking, fees and the privacy rules on the phone. The product team sets onboarding, first value, price and ad load. Lifecycle teams send push, email and in-app messages. Measurement and finance set the rules for credit, the tests and the meaning of profit.

FIGURE 1

Paid media owns the levers for acquisition and reactivation, and reaches activation and retention only indirectly



Six stages of app growth against the five groups that own the levers. Shading shows how directly each group controls each stage, as judged by the author from platform and store documentation; it is a framework, not a measurement. The dashed box marks where paid media has no direct lever. Feedback to the buying engines is delayed and thresholded on iOS [1]. Also: [2,3]

Limits. Control levels are the author's judgment. Paid media reaches activation and retention only indirectly, through whom it reaches, what the ad promises and where a deep link lands the user.

Package files: figures/F01.svg · data/figures/F01.csv

The map frames the rest of the paper. Paid media acts on activation and retention mostly at one remove: through whom it reaches, what the ad promises and where a deep link lands the user [CL-001]. Reactivation is the exception, since remarketing buys it directly. A buying platform can choose who sees an ad. It cannot make the first session useful. It can only learn which people tend to stay, and that learning rests on the signal the app sends back.

What counts as growth

Growth here means more durable value than you would have had without the spend. That takes four numbers, and the industry often blurs them.

- **Installs.** A device event. The count includes reinstalls and redownloads, which are not new customers.
- **Activated users.** People who reach a first-value event the app defines, such as a finished level, a first order or a started trial.
- **Retained value.** Revenue or use that is still there at day 30, day 90 or a year, on a stated definition.
- **Contribution.** Revenue after store fees, refunds, payment costs, taxes where they apply, cost of goods and promotions. Contribution here comes before acquisition cost. Growth is contribution minus the full cost of winning the customer, counted once.

A lower cost per install does not prove growth. Nor does a higher return on ad spend in a platform report. Cost per install can fall and reported return can rise while contribution falls. That happens when cheap installs churn faster, when credit goes to users who would have come anyway, or when a report counts gross sales instead of net revenue [CL-002](#).

The units matter

Cost per install divides spend by installs. Cost per payer divides it by payers. A household that saw a TV ad is not a person. An iOS privacy postback is not a user at all. It is a report about one install with no user ID, sent a day or more later, with its detail capped by privacy thresholds ^[1]. The metric dictionary in the appendix defines each unit and says what it can be added to.

Six money pools that must never be added up

Market sizes in this field are noisy because different pools of money get quoted as one. This paper keeps six apart.

1. **Q1: ad spend that promotes apps.** What advertisers pay to win installs and bring users back.
2. **Q2: ads sold inside apps.** What any advertiser pays to reach people while they use apps. Much of it promotes things other than apps.
3. **Q3: app-store consumer spending.** What people spend through Apple and Google billing.
4. **Q4: subscription revenue.** It overlaps Q3 when a store bills it. It sits outside Q3 when billed on the web.
5. **Q5: commerce value.** Orders placed in shopping, delivery and travel apps. This is gross merchandise value, not the app's revenue.
6. **Q6: ad-platform revenue.** What the platforms report. It is often net of what they pay publishers, so it is not ad spend.

These overlap in places and describe different markets in others. Adding them into one total overstates the market. Chapter 2 puts the best public number for each pool side by side.

What the paper asks

The brief for this research set twelve questions. They come down to one: **what produces extra durable app value, for which buyers, under which limits, and at what cost?** Each chapter takes part of that question. Each ends with a test: the case for a thesis, the best case against it, other explanations, and the result that would prove it wrong.

THESIS TEST · WHERE THE BINDING BARRIER SITS (Q2)	
THESIS	For many apps the binding barrier comes after the install, at activation and early retention, not at the price of the install.
SUPPORTING EVIDENCE	In a panel of more than 16,000 games, the median title kept about 22% of new users on day 1 and under 1% on day 30 ^[4] . Subscription apps with a hard paywall turn about five times more installs into payers by day 35 than freemium apps ^[5] .
BEST COUNTEREVIDENCE	Both figures come from vendor panels. The game medians are per title, include games that buy no ads, and cover games only. Apps that choose hard paywalls differ in category and price, so the gap does not show that the paywall causes it.
OTHER EXPLANATIONS	Low retention can reflect installs bought from low-intent placements. That is a media problem that looks like a product problem.
WHAT WOULD PROVE IT WRONG	A randomized test in which better-targeted media, with the product held constant, moved day-30 retention more than product changes of similar cost did.

CHAPTER 2

Where the budgets and the power sit

The money is large and still growing. Profit sits with a few firms, and the firm that earns most on its revenue also owns an auction its model learns from. Whether the ownership explains the profit is still an open question.

IN SHORT

AppsFlyer's client panel puts 2025 app marketing spend at about \$109 billion worldwide: \$78 billion to win new users and \$31 billion to bring old ones back. The US holds about 42% of acquisition spend. Remarketing grew faster than acquisition, at 37%. Store spending is a larger, separate pool, and two data vendors disagree on its size by \$11 billion. Among specialist ad platforms, AppLovin reported \$5.5 billion of net revenue for 2025 at an 82% adjusted EBITDA margin, the highest in our table. Liftoff, which owns no mediation auction or attribution company, reports a 55% margin on the same net basis. Google and Apple own more of the loop than anyone, and they do not disclose app-ad revenue.

The pools, side by side

The most detailed public estimate of ad spend that promotes apps comes from AppsFlyer, an attribution company. Its report, first published in December 2025, says app marketers spent \$109 billion worldwide in 2025. User acquisition took \$78 billion, up 13%.

Remarketing took \$31.3 billion, up 37%. That lifted its share of the total from 25% to 29% ^[6] CL-003. The panel is large: 32 billion paid installs and 45,000 apps. But it is AppsFlyer's own clients. The report does not state its measurement period, which may include estimated months, or how it scales the panel to a world total. Treat it as a well-grounded estimate. It is not a census.

The same report says the US holds 42% of global spend, in the section on acquisition. That puts US acquisition spend near \$33 billion in 2025, before remarketing ^[6]. The report does not print that dollar figure. It is our arithmetic, and it assumes the share refers to acquisition spend.

FIGURE 2


Three money pools, three scales: none of them can be added to another

OBSERVED VENDOR PANEL FILINGS

Q1 · Ad spend promoting apps

global, 2025, USD bn — AppsFlyer client-panel estimate


■ global ■ US



Q3 · App-store consumer spend


2025, USD bn — two app-intelligence vendors disagree

■ global ■ US



Q6 · Ad-platform company revenue

latest fiscal year, USD bn, company filings (not ad spend)



Not shown as a bar: in-app advertising revenue (Q2) and subscription revenue incl. web billing (Q4) have no public, non-paywalled market total at the cutoff. Context only, not app-specific: US digital advertising revenue across all channels was \$294.6bn in 2025 (IAB/PwC). Verve Group reports in EUR (€550.9m, FY2025) and is omitted to avoid mixing currencies.

Each panel has its own scale. Q1 is AppsFlyer's client-panel estimate of 2025 app marketing spend; the US bar is our derivation from its stated 42% share ^[6]. Q3 shows two app-intelligence vendors' estimates for the same year ^[7,8]. Q6 shows company revenues from filings, which mix net and gross revenue recognition ^[9,10,11,12,13,14]. Sensor Tower's figures include paid-app downloads and exclude third-party Android stores; its US figure is 'nearly \$60 billion'.

Limits. Vendor estimates are not audited; company revenue is not ad spend; fiscal years differ (Digital Turbine to March 2026).

Package files: figures/F02.svg · data/figures/F02.csv

Store spending is a different pool. Sensor Tower estimates \$167 billion of in-app purchase and paid-app spending on iOS and Google Play in 2025, up 10.6% ^[7]. Appfigures estimates \$155.8 billion, up 21.6% ^[8]. The two vendors describe the same year and the same two stores. They differ by about \$11 billion on the level and 11 points on the growth rate ^[CL-004]. Part of the gap is scope, and neither publishes a method a buyer could check. Both agree on one thing: non-game apps now earn more in the stores than games do.

Ads sold inside apps have no public market total at the cutoff that we could read without a paywall. We leave that bar empty rather than fill it with a guess. Liftoff's prospectus cites a commissioned estimate of a \$79 billion market in 2025 for independent in-app ad technology ^[15]. That is a different quantity again, prepared for a share sale.

Where growth came from in 2025

Three shifts stand out in AppsFlyer's data, all of them global. First, acquisition growth came entirely from iOS, up 35%, while Android slipped 1% ^[6]. Second, non-gaming apps drove it. Their acquisition spend rose 18% to \$53 billion, while games rose 3% to \$25 billion. Shopping apps grew 70%, largely on budgets from China-based e-commerce apps. Third,

money moved toward existing users. Remarketing grew almost three times as fast as acquisition.

The first two shifts matter for this paper's argument, and they cut against the simple version of it. Chapters 3 and 5 show that the bidding machines learn best where signal is dense and fast. Yet in 2025 money grew fastest on iOS, where signal is thinnest, and in non-game apps, whose value often arrives late. So the machines' view does not decide where budgets go. Value per user, category mix and a wave of China-based e-commerce spending explain the shift at least as well. AppsFlyer does not publish a US-only category split, so we cannot say how much of the shopping surge reached the US.

The third shift is the one to watch. Remarketing reaches people who already have the app, and some would have come back anyway. Randomized studies of web retargeting find real but modest lift (chapter 5). We found no independent randomized study of app remarketing. So the faster-growing of the two budget lines is also the one we found no independent app test for [CL-005](#).

Who earns the money

Company filings give harder numbers. They measure a different thing: the revenue of ad platforms, not ad spend.

COMPANY	YEAR	REVENUE (BASIS)	PROFIT MEASURE	SOURCE
AppLovin	2025	\$5,481m, up 70% (net of publisher payouts)	Adjusted EBITDA \$4,512m, 82% of revenue	[9,16]
Mobvista (Mintegral)	2025	\$2,047m, of which Mintegral \$1,961m (gross)	Adjusted EBITDA \$191m, 9% of revenue	[10,17]
Unity (all segments)	2025	\$1,850m (mixed; includes its engine business)	Adjusted EBITDA \$409m, 22%	[11]
Liftoff	2025	\$686m, up 32% (net)	Adjusted EBITDA \$374m, 55%; net loss \$23m	[12,15,18]
Digital Turbine	year to March 2026	\$565m, up 15% (mixed)	Adjusted EBITDA \$122m, about 22%	[13,14]
Verve Group	2025	€551m (mixed after a Q3 change)	Adjusted EBITDA €134m: 24% of reported revenue, 22% of like-for-like revenue	[19,20]

Three cautions apply before anyone ranks these. First, firms book revenue differently. AppLovin and Liftoff act as agents and report revenue net of what they pay publishers ^[21,22]. Mobvista books gross, as principal, with a gross margin of 21.2%; its own non-IFRS ad-tech net revenue for 2025 adds up to about \$519 million ^[17,23]. Digital Turbine books its exchange net and its brand and performance business gross ^[14]. Verve moved some revenue from net to gross in the third quarter of 2025 ^[19]. A margin on net revenue looks far higher than the same business on gross revenue [CL-006](#). Second, adjusted EBITDA is each firm's own measure. Third, Unity's margin covers its game-engine business as well as ads, so it is not an ad-business margin.

So the table supports a narrow claim. In this table, AppLovin's net revenue exceeded the next four specialists combined, even with Mobvista counted gross. That is a comparison of selected firms, not a measure of market concentration. Its margin is the highest in the table. Liftoff, the one peer that also reports net, is the fairest comparison.

The latest quarter at the cutoff kept the pattern. AppLovin reported \$1,924 million of revenue for the second quarter of 2026, up 53%, at an adjusted EBITDA margin of 83.9% ^[24]. Unity's Grow revenue rose 35% to \$389 million. Unity credits its Vector-driven ad network, partly offset by the older ironSource network ^[25]. Liftoff listed on 4 June 2026. In its first quarterly report as a public company, for April to June, revenue rose 35% to \$220 million. One customer made up about 10% ^[22].

The mix differs too. Game advertisers supplied three quarters of Mintegral's 2025 revenue ^[10]. At Liftoff, slightly more than half of advertiser revenue came from outside gaming before its listing ^[12]. At Unity, the Unity Ad Network, now running on the Vector model, made up 56% of Grow revenue in the fourth quarter of 2025. The older ironSource network fell to 11% ^[11].

Stacks, and who owns the loop

Common ownership is the key structural fact. AppLovin owns four pieces. Its ad network and bidding engine was sold as Axon Ads Manager and renamed AppLovin Ads by mid-

2026. MAX is the mediation auction many game publishers use to sell their inventory. Adjust, bought in 2021, is an attribution company. Wurl, bought in 2022, is a streaming TV business [26,21]. Unity owns the Vector and ironSource ad networks and the LevelPlay mediation auction [11]. Digital Turbine combines phone preloads with its own exchange [13]. Mintegral buys on its own software kit inside publisher apps [10].


The biggest owners of the loop are not specialists. Google owns Android, the Play store, the Install Referrer, the AdMob mediation platform and the Firebase analytics kit. Its app campaigns expect Firebase-sourced events for target-return bidding [27,28]. Apple owns iOS, the App Store, the privacy thresholds and Apple Ads, which has its own direct attribution interface (chapter 3). Neither discloses app-ad revenue, so neither appears in the table, though their app-ad businesses are likely larger than any specialist's.

FIGURE 3

Money flows one way; feedback flows back to whoever owns the auction and the measurement

SCHEMATIC

FILINGS



Common ownership stacks (company filings and disclosures)

■ AppLovin	Axon ads (demand) · MAX (mediation) · Adjust (MMP) · Wurl (CTV)
■ Unity	Vector / ironSource Ads (demand, network) · LevelPlay (mediation)
■ Digital Turbine	On-device distribution · exchange and app-growth platform
■ Mobvista	Mintegral (demand plus its own SDK supply)

Solid arrows: money and bids.
Dashed arrows: feedback data.
Orange: data an auction owner can use in its own bidding (disclosed by AppLovin).

A simplified map of roles in app advertising. Solid arrows carry money and bids; dashed arrows carry feedback data. AppLovin discloses that its Axon model uses MAX auction outcomes [29]. Ownership stacks are from company filings [26,11,13,10]. A shared owner creates a potential conflict of interest; it is not evidence of misconduct.

Limits. Schematic only; flows are not measured and many buyers use several paths at once.

Package files: figures/F03.svg · data/figures/F03.csv

Why does ownership matter? A mediation auction sees who bid what for every impression, and who won. AppLovin's disclosure on its Axon model says the model uses win and loss notices from the MAX auction, along with device and engagement signals and data advertisers share [29] [CL-007](#). A buyer who uses AppLovin media, measures it with Adjust and monetizes through MAX deals with one company at three points. AppLovin says Adjust data is not shared with it unless the customer directs [26]. That is the company's own

assurance, and it has not been independently checked. The overlap is a potential conflict of interest. It is not evidence of misconduct.

AppLovin has also faced public allegations. In February and March 2025, four research outfits, three of them short sellers who profit if the stock falls, published reports on the company. They claimed silent app installs pushed through the MAX software kit and use of persistent identifiers against partner platforms' terms. They also claimed e-commerce results came more from retargeting than from new customers. One report, by Muddy Waters, said log files from a demand-side platform covering five advertisers showed about 52% of AppLovin's e-commerce conversions came from retargeting. It put incrementality at only about 25% to 35%. Its data and method are not public ^[30,31,32]. The stock fell 12% on 26 February and 20% on 27 March 2025. AppLovin's chief executive and technology chief denied the claims in blog posts ^[32,33,34]. A securities class action followed. In January 2026 the plaintiffs asked to add reported SEC and state investigations to their complaint. In April the court took that request and AppLovin's motion to dismiss under submission. No ruling appeared on the public docket as of mid-August, and related shareholder suits were on hold ^[35,36,37]. In October 2025 Bloomberg reported, and Reuters relayed, that the SEC was probing AppLovin's data collection. AppLovin said it does not comment on potential regulatory matters ^[38]. Its filings name no specific investigation, though each says in general terms that the company is involved in regulatory investigations ^[39,40,41,35]. In September 2026 a separate securities suit was filed over statements about AI products ^[42]. The short-seller reports are claims by parties with a financial stake. The complaints are allegations: the public docket showed no ruling on the Brownback motion to dismiss as of mid-August 2026, and the September suit had just been filed. That suit does not itself confirm the reported SEC inquiry. We record all of this because it bears on data use, not because this research could test it.

The market keeps reshaping

Ownership has moved fast. Unity merged with ironSource in November 2022, after rejecting an all-stock approach from AppLovin that valued Unity at \$20 billion ^[43,44]. AppLovin sold its game studios to Triplicated in June 2025 for \$430.6 million in cash and about a fifth of Triplicated's equity, valued at \$285 million ^[45,16]. Liftoff listed on Nasdaq in June 2026 at \$23 a share. A first attempt, filed in January, was withdrawn on 17 February 2026 ^[46,47]. LoopMe bought Chartboost from Zynga in December 2024 ^[48]. Aarki had rebranded as RZR by March 2026 ^[49,50]. Affle runs at least four app-marketing brands: Jampp, YouAppi, RevX and mediasmart. In June 2026 it agreed to buy the AdColony technology and brand from Digital Turbine ^[51,52,53,54,55]. The tools around the market move too. OpenAI bought the testing firm Statsig in 2025, and in May 2026 Amplitude bought Statsig's customer contracts, brand and a licence to its technology ^[56,57].

The pattern is consolidation around a few stacks, with specialists and new listings in between. It is consistent with the playbook's view that a US buyer compares any challenger with AppLovin, Moloco and Liftoff first, but it does not test that view. This research did not study firms that failed or closed, so the sample leans toward survivors.

THESIS TEST · WHAT COMPOUNDS FOR SCALE (Q5)

THESIS (HYPOTHESIS)	Owning an auction whose outcomes feed the bidding model gives a lasting edge in model quality.
SUPPORTING EVIDENCE	AppLovin's 70% revenue growth and 82% adjusted EBITDA margin in 2025 ^[9] , and its disclosure that the Axon model uses MAX auction outcomes ^[29] .
BEST COUNTEREVIDENCE	Liftoff has no mediation auction or attribution company and still reports a high margin on net revenue ^[15] . Unity owns an auction, yet its Grow revenue picked up only after a new model, Vector, arrived ^[11] , which points to model quality more than ownership.
OTHER EXPLANATIONS	Model quality, revenue accounting, category mix (high-value game genres) and AppLovin's e-commerce push could each explain part of the gap.
WHAT WOULD PROVE IT WRONG	Independent tests in which platforms without owned supply delivered equal incremental value per dollar across categories, or no change in AppLovin's performance where MAX signals were restricted.

CHAPTER 3

What the signal can tell you

On iOS the operating system decides what an advertiser may learn, when, and in how much detail. On Android the old identifiers still work, and Google's privacy replacements are being wound down.

IN SHORT

In one attribution vendor's panel, 35% to 38% of iOS users who see the tracking prompt accept it. For each install credited to an ad network, Apple's frameworks send up to three postbacks with no user ID. Each arrives one to six days after its window closes, and the windows run up to 35 days from first launch. How much detail each carries depends on crowd size, so bigger campaigns tend to learn more. SKAdNetwork and its successor, AdAttributionKit, run side by side, and Apple has set no end date for the old one. On Android the advertising ID and the Play Install Referrer still carry attribution. Google announced in October 2025 that it would phase out its Privacy Sandbox ad tools on Android. At the cutoff that phase-out was scheduled, not done.

iOS: consent first, then thresholded postbacks

Since April 2021, an iOS app must ask before it tracks a user across other firms' apps ^[58,59]. Most users who see the prompt say no. Adjust, an attribution vendor owned by AppLovin, reports on its own client base. There, opt-in among users shown the prompt rose from 35% to 38% between early 2025 and early 2026 ^[60,61,62,63]. That is a panel figure, not a rate for all users. A survey asked 11,000 US and UK iPhone users which prompts they would accept. About 25% said yes to Apple's own prompt for personalized ads, against 13% for the tracking prompt other firms must show ^[64]. Two of that study's five authors work at Meta, which has a stake in the answer, and a third has consulted for it.

Apple's frameworks report ad-driven installs without a device ID, whether or not the user opted in ^[65]. For users who decline, an app may not link its data with other companies' data to measure ads ^[66]. For ads outside the App Store, the postbacks are then the only attribution Apple itself provides. Under SKAdNetwork 4, an install can produce up to three postbacks. They cover days 0 to 2, 3 to 7 and 8 to 35 after the user first opens the app. Each is sent after a random delay once its window closes, or sooner if the app locks its value early. The delay is 24 to 48 hours for the first and 24 to 144 hours for the others. So the last can arrive about 41 days after first launch ^[1]. Apple gives each download one of four privacy tiers. The tier reflects crowd size across the app or site showing the ad, the advertised app, the install country and the campaign code. Apple does not publish the thresholds. In the lowest tier the advertiser gets one postback, a two-digit campaign code and no conversion value at all ^[1].

That design has a commercial side effect. Detail rises with crowd size. Installs that cluster in a few publishers, countries and campaign codes are more likely to reach the top tiers. Even there, the fine-grained value comes only in the first postback; later ones carry a coarse value ^[1]. A small advertiser, or one that splits spend thinly, is more likely to get coarse values or none ^{CL-008}. Apple itself says smaller developers may get limited data back because of these thresholds ^[67]. Rules written to protect users also make signal a scale advantage.

What else fills the iOS gap

Postbacks are not the only tool in use. Attribution companies document probabilistic matching, which links an ad click to an install by device and network traits inside a short window. Platforms also report modeled conversions, which estimate what they cannot observe ^[68,69]. That sits close to a line Apple draws. Its rules bar apps from deriving data from a device to identify it uniquely. Apps or kits caught doing so may be rejected from the App Store ^[70]. Adjust says such modelling applies only where Apple's tracking policy allows it ^[71]. AppsFlyer ranks it below exact matching and says it produces campaign-level reports

rather than identifying devices^[72,73]. Buyers should know which part of a reported number is observed, which is matched by probability, and which is modeled.

AdAttributionKit has not replaced SKAdNetwork

AdAttributionKit arrived with iOS 17.4 in March 2024. It first added attribution for alternative app stores in the EU. Re-engagement followed in the iOS 18 cycle, then a country code, adjustable windows and cooldown periods in the iOS 26 cycle^[65,74]. Apple's SKAdNetwork page tells developers to use AdAttributionKit for app ad campaigns^[75]. Its guide says only one impression can win an install across the two frameworks. It bridges them one way: a value update sent through SKAdNetwork is mirrored into AdAttributionKit. The same guide still tells apps that work with both to call both^[76]. The release notes list no SKAdNetwork 5, and no document we opened sets a date to retire SKAdNetwork^[77] CL-009. The changelog has no entry after June 2025. One later addition, a way to register view-through impressions in iOS 26.2, appears in Apple's reference pages but not in the changelog. The iOS 27 release notes list no attribution changes^[74,78].

Two rules shape re-engagement. For installs, the winning network gets postbacks and up to five others get one losing postback each. For re-engagement, only the winner hears anything^[65]. A network that helped bring a user back but lost the last touch learns nothing.

The big attribution companies document live support for both frameworks, including copies of postbacks sent to them^[79,80,81]. The plumbing exists. The limit is what the postbacks say.

Apple measures its own ads differently

Apple Ads, Apple's own App Store ad business, has its own attribution interface. It covers only Apple Ads download campaigns, in search results, the search tab, the Today tab and product pages. The app gets a token that must be traded with Apple's server within 24 hours^[82,83]. Apple credits a tap up to 30 days before the install and, since March 2025, a view up to 24 hours before^[84]. In September 2026, the month of the cutoff, Apple said campaigns that target by age or gender will get no attribution through that interface^[84]. Apple Ads also registered with AdAttributionKit in April 2025, for now through SKAdNetwork versions 1 to 3 and for click-throughs only^[85,67].

So Apple provides two kinds of attribution. Its own ads get a direct interface as well as postbacks. For other networks, Apple provides only postbacks. Those networks may add attribution from data they own, but matching users across companies needs the user's consent^[67,66]. Buyers should read Apple Ads results with that in mind. It is a difference in what can be seen, and it says nothing about which channel works better.

Europe is pushing on the consent screen

Competition authorities in three EU countries have acted against the way Apple runs its tracking prompt. France fined Apple €150 million in March 2025. It found the idea of the prompt acceptable but the way Apple ran it abusive^[86]. Italy fined Apple €98.6 million in December 2025 over the double consent the prompt adds on top of privacy-law consent^[87]. In August 2026 Germany's competition authority made Apple's commitments binding. Apple must align the wording of its own and other firms' prompts, drop discouraging design, and let apps combine the prompt with their privacy-law consent. It has four months to comply^[88].


Apple's developer documents now describe a fuller, text-rich consent sheet for the EU. It is part of iOS 27.2, which was in beta at the cutoff. It will be required for users in France, Germany, Italy, Poland and Romania. The same release notes add a yearly re-prompt in the EU^[89,90]. Apple does not link these changes to the rulings, and they had not shipped at the cutoff. If they ship and lift opt-in, European iOS measurement will improve. US rules do not change.

Android: the old tools still carry the load

On Android, attribution still rests on two long-standing tools. The advertising ID can be reset or deleted. Apps that target Android 13 or later must declare a permission to read it, and a user who deletes it hands apps a string of zeros ^[91,92]. The Play Install Referrer tells an app which referral link led to its install, with click and install times ^[93].

In October 2025 Google said it would phase out the Privacy Sandbox ad tools on Chrome and Android. They include Attribution Reporting, Topics, Protected Audience, the SDK Runtime and several others ^[94]. At the cutoff, Google's own status page listed the Android pieces as "scheduled for phaseout", with no removal date ^[95] CL-010. The author's playbook called these tools "retired" on Android in October 2025 ^[96]. The exact position is that retirement was announced, and the identifiers they were meant to replace are still in service. Chapter 12 lists this as a correction.

FIGURE 4
Legal and technical changes to app measurement, 2021 to the cutoff



Dated changes to app measurement, store rules and privacy law, April 2021 to October 2026, from official documentation, regulators and court records ^[1,65,74,84,86,88,87,95,94,92,97,98,99,100,101,102]. Month-only dates are shown on the first of the month. The Android Privacy Sandbox phase-out was announced, not completed, at the cutoff.

Limits. Selected events only. The EU business terms effective 1 October 2026 and Google's link-out reporting are shown after the cutoff line as scheduled. Dates shown as year and month only are approximate release cycles.

Package files: figures/F04.svg · data/figures/F04.csv

OS measurement matrix				
What each platform lets an advertiser see, with the status of each mechanism at 27 September 2026. The full policy and product status register is in the appendix.				
FUNCTION	IOS	ANDROID	STATUS AT CUTOFF	SOURCES
Device identifier	IDFA only after the user allows tracking	Advertising ID; apps targeting Android 13+ must declare AD_ID; zeros once the user deletes it	Live	C-S07 C-S16 C-S18
Install attribution (paid)	SKAdNetwork 4 / AdAttributionKit postbacks: up to 3 windows (days 0-2, 3-7, 8-35), 24-144h delay, crowd-anonymity tiers	Play Install Referrer (referrer, click and install times) plus Advertising ID	Live; SKAN and AAK bridged	C-S02 C-S05 C-S17
Re-engagement attribution	AdAttributionKit re-engagement; winner-only postbacks	Attribution company deep-link and ID matching	Live (iOS 18 cycle onward)	C-S04 D-S14
Platform's own ads	Apple Ads via AdServices token (24h to exchange); 30-day tap and 24-hour view windows (view-through since March 2025); no attribution for age/gender-targeted campaigns from Sept 2026	Google, Meta, TikTok self-attribution via device ID	Live	D-S20 D-S21 C-S10 D-S13
Aggregate privacy APIs	n/a (postbacks are Apple's aggregate path)	Privacy Sandbox Attribution Reporting, Topics, Protected Audience, SDK Runtime: phase-out announced Oct 2025, 'scheduled for phaseout' at cutoff	Announced retirement	C-S14 C-S15
Consent screen	ATT prompt; beta expanded EU consent sheet documented, mandatory in five EU countries once shipped	No OS-level tracking prompt; user can reset or delete the ID	Live; EU beta not shipped	C-S08 C-S12 C-S16

What this means for the buyer

The two platforms now differ less in whether you can measure and more in what you can see. On Android, exact install attribution is still normal. On iOS, most ad-driven installs are reported late, without a user ID, and with detail that depends on crowd size. Measured growth can move for reasons unrelated to real growth. AppsFlyer reports that US iOS paid installs rose 31% in 2025 while Android rose 8% ^[6]. Part of that gap may be real demand and part may be better attribution coverage. The report does not separate them.

THESIS TEST · PRIVACY LIMITS FAVOUR SCALE (Q5, Q7)	
THESIS	Apple's privacy thresholds give larger advertisers richer feedback, which adds to their optimization edge.
SUPPORTING EVIDENCE	Postback detail is tied to crowd size through the privacy tiers ^[1] , and Apple says smaller developers may get limited data back ^[67] ; value bidding on major platforms needs dense conversion signal (chapter 5).
BEST COUNTEREVIDENCE	Networks can pool learning across advertisers, so a small advertiser on a big network may still gain from the network's scale.
OTHER EXPLANATIONS	Large advertisers may simply have better data teams and bigger creative budgets.
WHAT WOULD PROVE IT WRONG	Evidence that small iOS advertisers with coarse postbacks get similar incremental return per dollar to large ones on the same networks.

CHAPTER 4

From attribution to causality

Every system that counts app installs counts them differently, on purpose. None of them counts what would have happened without the ad.

IN SHORT

Ad platforms, attribution companies, Apple's postbacks and finance each use their own clock, window and rules. Their totals will not match, and they are not meant to. Attribution assigns credit. It does not measure cause. Across 663 randomized test-and-control comparisons on Facebook, statistical methods run without the control groups overstated the lift the experiments measured by three to thirteen times. The gap does not always run one way: on iOS, privacy limits also hide some installs that ads did cause. Lift tests exist on the big platforms, but the ones we reviewed run through an account team.


Why the numbers never match

Put a campaign's results from Meta, TikTok, an attribution company, Apple's postbacks and the finance ledger side by side. They will differ. That is not a bug in any of them. Each answers a different question with different rules.

- **Different clocks.** Meta's Ads Manager files a conversion under the time of the ad view or click, not the install. Meta warns this causes gaps against attribution-company reports ^[103].
- **Different windows.** AppsFlyer's defaults credit a click up to 7 days before the install and a view up to 1 day before ^[68]. Meta's windows depend on where you look. App-install ad sets optimize on a 1-day click, with a 1-day view optional, and iOS app-event campaigns use a 1-day or 7-day click. Ads Manager can show 28-day click counts, but only as a comparison. Meta's older App Events API, no longer recommended for new work, counts clicks within 30 days ^[104,105,106,103]. TikTok offers 1, 7, 14 or 28 days for clicks, up to 7 days for views, and a separate window for "engaged views" of 6 seconds or more ^[107].
- **Different ideas of a click.** In Singular's aggregate reports on Meta campaigns, engaged-view conversions are filed as click-throughs ^[108].
- **Separate claims.** Meta, Google and TikTok are "self-attributing": each decides for itself which installs it claims, under its own rules. Since March 2025 TikTok requires this mode for all apps, and its help page says it works "without affecting your MMP's final attribution logic" ^[109]. Two platforms can claim the same install. The attribution company then picks one winner.
- **Reinstalls.** If someone reinstalls inside the reattribution window, AppsFlyer records no new install unless the user engaged with a retargeting campaign. After the window, the same person counts as new ^[110].
- **Thresholds and delay.** Apple's postbacks arrive days later, with no user ID and detail capped by privacy thresholds. Only one network wins each install (chapter 3).
- **Finance.** The ledger books revenue net of store fees, refunds and taxes, when earned. Ad reports usually book gross value on the day of the event.

FIGURE 5

Five ways to count the same campaign, and what changes between them



Each system uses its own clock, window and identity basis, drawn from platform and attribution-company documentation [68,110,103,107,109,108,1]. Apple sends up to three postbacks for each install credited to the winning network, one per conversion window, and at most one to each of up to five losing networks; privacy tiers limit what each carries. Arrows show what each step removes or adds relative to the one before. The steps are not a single funnel with measured drop-offs.

Limits. Default windows are configurable; platforms change reporting rules often.

Package files: figures/F05.svg · data/figures/F05.csv

The practical rule follows. Never add platform-reported conversions and privacy postbacks into one total. Summing claims from several self-attributing platforms counts some installs twice, so the total is not a count of installs [CL-011]. Reconcile each system on its own terms. Then ask a separate question about cause.

The error does not always run one way. On iOS, installs in the lowest privacy tier carry no conversion value [1]. An install that no method can match to an ad is counted as organic [68]. Paid volume may also lift an app's store ranking and so its organic installs, which no report credits to the ads. All of that makes paid media look weaker than it is. The net direction varies by channel and platform, which is one more reason to test instead of trusting any single report.

Attribution reconciliation register								
Each reporting system's rules, from its own documentation. Use it to explain gaps, not to merge totals.								
SYSTEM	EVENT	CLOCK	WINDOW	IDENTITY BASIS	MODELED PART	LAG	WHY IT DIFFERS	SOURCES
Meta Ads Manager / App Events API	App install or in-app event logged via Facebook SDK, App Events API, or Conversions API	Ads Manager: impression or click time, not install time (Meta flags this as a cause of gaps with MMP reports)	App-install ad sets optimize on a 1-day click (1-day view optional); iOS app-event and value campaigns use a 1-day or 7-day click; Ads Manager shows 28-day click only as a comparison; the legacy App Events API endpoint counts clicks within 30 days	Device/user match for consented traffic and Android; aggregated for iOS AEM	Not independently confirmed in this research for Ads Manager 'estimated results' wording; App Events API confirms only the impression/click-time clock basis	Numbers can change after the fact because insights are keyed to impression/click time rather than install/conversion time	Clock basis (impression/click-time vs. install-time) plus a documented reclassification of engaged-view conversions into the 'click-through' bucket in partner (Singular) aggregate reporting	D-S16 D-S24
AppsFlyer (MMP)	Install/re-engagement attributed via self-reporting-network (SRN) device-ID query, referrer match, device-ID matching, or probabilistic modeling	Click/impression time opens the lookback window; attribution is finalized at first app launch	Click: 7-day default (1-30 configurable); view-through: 1-day default (0-24h); engaged-click: 2-day default (1-7 days)	Deterministic methods (referrer, device-ID matching) prioritized over probabilistic modeling within the lookback window; SRNs use MMP-API device-ID query subject to ATT consent plus SKAdNetwork on iOS	Probabilistic modeling used as a fallback method for CTV, PC/console and some non-SRN networks	Not documented in the pages accessed	Choice of window length, deterministic-vs-probabilistic priority order, and reattribution-window rules (installs inside the reattribution window get no postback and are classified organic) produce a different organic/non-organic split than a platform's own native reporting	D-S13 D-S14 D-S15
TikTok Ads Manager (SAN)	In-app event or install claimed via TikTok's mandatory Self-Attributing Network (all apps since 31 March 2025)	Click, view, or engaged-view event time	Click-through: 1, 7, 14 or 28 days; view-through: off, 1 or 7 days; engaged view-through (6s+ watched) for app campaigns: 1 or 7 days	SAN device-level claim reported directly to TikTok Ads Manager; SKAN 4.0 is documented as a separate, additional iOS mechanism	Not detailed in the pages accessed beyond the SKAN 4.0 reference	Not documented in the pages accessed	TikTok's own documentation explicitly states SAN integration does not affect an MMP's 'final attribution logic' -- the two systems are designed to be allowed to disagree	D-S22 D-S23
SKAdNetwork 4 (iOS, network-agnostic)	Install/re-download plus up to three staged conversion-value updates	Three fixed postback windows keyed to days since install/re-download (0-2, 3-7, 8-35), each delivered 24-144 hours after window close	Up to 35 days total across three sequential windows; developers can 'lock' a conversion value early to accelerate postback delivery	No device ID; aggregated postback gated by a 4-tier crowd-anonymity threshold determining fine (6-bit, 0-63) vs. coarse (none/low/medium/high) conversion-value detail	Coarse-value bucketing is itself a privacy-driven substitute for granular data when crowd size is insufficient	24-144 hours per postback, on top of the multi-day window itself	Delayed, tiered, aggregated, device-ID-free postback logic is structurally incompatible with MMPs' deterministic click/device-ID attribution, producing a persistent, by-design reconciliation gap on iOS	D-S11
Apple AdServices / Apple Ads Attribution API	Apple Ads app-download campaigns on the App Store (search results, search tab, Today tab, product pages)	24-hour attribution-token TTL; the token must be exchanged server-side within that window to retrieve the attribution record	Tap within 30 days; view within 24 hours (view-through since March 2025); no attribution for age/gender-targeted campaigns from Sept 2026	Server-side token tied to device, scoped only to Apple Search Ads	None documented	None documented beyond the 24-hour token TTL	AdServices is structurally narrower in scope (Apple Search Ads only) than SKAdNetwork/AdAttributionKit (all networks), so it is not a substitute reconciliation source for non-Apple-Ads traffic	D-S20 D-S21
Google Ads (Data-Driven Attribution)	Conversion path across Search, YouTube, Display and Demand Gen ads	Not addressed in the page accessed	Not stated in the page accessed	Multi-touch statistical credit across the observed conversion path; the default for most conversion actions and available at any volume; Google recommends 200 conversions and 2,000 interactions in 30 days	Not confirmed in this research -- the accessed page did not address modeled conversions for privacy-limited traffic despite that being specifically sought (see open_questions)	Not documented in the page accessed	Accounts below the documented data-sufficiency threshold cannot use data-driven attribution and default to a different (rules-based) model, creating cross-account inconsistency in how Google's own reporting allocates credit	D-S10

Credit is not cause

Attribution tells you which touch came last, or which system claims the install. It does not tell you whether the install would have happened anyway. Only a valid comparison group can do that. The cleanest is a group randomly kept from the ads, analyzed as assigned. Comparing people who happened to see an ad with people who did not keeps the selection bias.


The strongest public evidence on the size of the gap comes from Meta's own test platform, studied by academics with Meta ties. Gordon, Moakler and Zettelmeyer studied 563 US ad experiments on Facebook, run from November 2019 to March 2020. Some tested several ads, so there were 665 test-and-control pairs, which the authors count as experiments. Together they served about 38 billion impressions. The median lift the experiments measured was 29% for upper-funnel outcomes, 18% for mid-funnel outcomes and 5% for purchases. They then ran two statistical methods on the same campaigns without the control groups. A double machine-learning method estimated median lifts of 83%, 58% and 24%. Propensity-score matching estimated 173%, 176% and 64% ^[111] [CL-012](#). Both

overstated lift at every stage. As a multiple, the error was largest at the bottom of the funnel: about five times the measured 5% for machine learning and thirteen times for matching. One author works at Meta, and two were part-time Facebook employees to reach the data. An earlier study of 15 experiments reached the same conclusion ^[112].

This study tests statistical corrections, not attribution rules. Two older studies speak more directly to attribution-style reports. At eBay, a standard regression put the return on paid search above 4,000%. Adding controls for region and day still left it above 1,600%. A 60-day test then switched off non-brand search ads in about 30% of US regions and found a return of minus 63% ^[113]. Ads on eBay's own brand name showed no measurable short-term benefit ^[114] **CL-013**. Across 25 retail and finance experiments at Yahoo, the median 95% confidence interval on return on investment was more than 100 percentage points wide. An informative test could need more than 10 million person-weeks ^[115]. Real effects exist. They are smaller and noisier than dashboards suggest, and hard to measure without large tests.

FIGURE 6

The strongest causal evidence is randomized, large and mostly not about apps



Each dot is a study used in this paper, placed by design and by the independence of its authors. Orange dots study apps or app users. Several of the largest experiments were run on Meta's platform with Meta co-authors ^[112,111,116]. The study table below gives sample, outcome and limits for each. Also: ^[115,114,117,118,64,119,120,121,122,123]

Limits. Placement on the independence axis is the author's reading of disclosed affiliations; abstract-only studies are marked in the study table.

Package files: figures/F06.svg · data/figures/F06.csv

Privacy changes show how much depends on data access. Aridor and co-authors found that after Apple's tracking prompt arrived, click-through rates on Meta's conversion-optimized ads fell 37%. Stores more exposed to the change saw revenue fall 8% to 40% relative to less exposed stores, mostly among smaller ones. Exposure was measured by reliance on Meta ads or on iPhone shoppers. The revenue data come from Grips Intelligence, a firm one co-author works with ^[117]. A randomized test on Meta in fall 2021 covered more than 70,000 advertisers. It estimated that losing off-Meta data would raise the median cost per

incremental customer from \$38.16 to \$49.93, a 31% rise ^[118,124]. Two of the four authors worked at Meta at the time. Meta could screen the paper for confidential data but had no right to block it over its results. It was peer reviewed, posted online in September 2024 and printed in the March 2025 issue of Marketing Science. So it is platform-affiliated evidence, though built on a large randomized test.

Tests exist, but someone has to switch them on

Each large platform offers a lift test with test and control groups. Meta's Conversion Lift, Google's Conversion Lift and TikTok's Conversion Lift Study each document a test-and-control design that can cover app outcomes. In the pages we reviewed, all three run through an account team or eligibility rules, not open self-service ^[125,126,127] [CL-014](#). In August 2026 Google's Ads API added read-only access to Conversion Lift results, so software can now pull them. The release notes do not describe creating studies through the API ^[128]. Among specialists, Moloco documents built-in ghost-bid tests for streaming TV and says it offers incrementality testing for re-engagement, without naming the method ^[129,130]. Jampp advertises always-on lift measurement for acquisition and retargeting at no extra cost, without documenting the method ^[69]. Remerge describes ghost bids for retargeting and an econometric method for installs and re-engagement ^[131]. AppLovin describes running geographic holdouts for clients as a managed service ^[132]. These are vendor descriptions. We found no independent audit of their results.

The "ghost ads" idea needs a plain explanation. Users are first split at random into test and control groups. The system then logs each moment a control-group user would have seen your ad, and records what they did next. That gives a clean comparison without paying for filler ads in the control group ^[133]. The random split does the work; the log only finds the right people to compare. The logging must work the same way in both groups, or the comparison is biased ^[133,134]. It needs the ad platform to run it, which leaves the vendor marking its own homework unless it shares the method and raw results.

Where models fit

Marketing mix models estimate each channel's share of outcomes from spend and results over time. Google's open-source Meridian and Meta's Robyn both model diminishing returns and both accept test results to calibrate the model ^[135,136]. Neither claims independent accuracy. Robyn's guide says its outputs are not guaranteed and should be checked before use ^[136]. PyMC-Marketing cites its own benchmark claiming faster fitting and 40% lower error on channel shares than Meridian ^[137]; that is a vendor's self-test. For apps, a mix model works best at weekly or monthly grain across a few large channels, calibrated with experiments. Google's Meridian guide says the model works at channel level and does not recommend campaign-level runs ^[138]. With the data most apps have, estimates for single campaigns are unstable. So a mix model is a poor guide to which of fifty campaigns to cut.

Can a buyer reconcile without inventing one number?

Yes, if the buyer stops asking for one number. A workable setup keeps four ledgers, each with a stated job. The attribution company's ledger sets day-to-day credit and pays partners. Apple's postbacks and the platforms' own reports feed bidding. Experiments and a calibrated mix model set budget between channels. The finance ledger decides whether the business made money. The reconciliation register above sets out each system's clock, window, identity basis and modeled share.

THESIS TEST · REPORTED GAINS VERSUS CAUSAL GAINS (Q3)

THESIS	Much of the reported gain in app marketing reflects credit rules, customer mix and changes in measurement coverage rather than causal lift.
SUPPORTING EVIDENCE	Statistical methods without control groups overstated median measured lift three to thirteen times ^[111] ; a regression-based search return above 1,600% against a tested minus 63% ^[114] ; attribution windows and view credit that vary by platform ^[68,103,107] .
BEST COUNTEREVIDENCE	The same studies show positive lift at every stage, so ads do cause outcomes. On iOS, privacy limits also hide installs that ads caused, so some reports understate paid media.
OTHER EXPLANATIONS	Most of the causal studies cover web and retail, not app installs; app-specific randomized results are rarely public.
WHAT WOULD PROVE IT WRONG	Published app-install experiments showing attributed and incremental conversions agreeing within a narrow band.

THESIS TEST · RECONCILIATION WITHOUT ONE COUNT (Q7)

THESIS	Buyers can reconcile platforms, attribution, postbacks, experiments and finance only by giving each a defined job, not by merging them.
SUPPORTING EVIDENCE	Each system's clock, window and identity basis are documented and differ by design ^[103,109,110] .
BEST COUNTEREVIDENCE	Some attribution companies sell modeled "single source of truth" dashboards that blend these inputs.
OTHER EXPLANATIONS	Smaller buyers may lack the volume for experiments, which leaves them with blended models by default.
WHAT WOULD PROVE IT WRONG	A blended model checked against randomized holdouts across many apps, with small and stable error.

CHAPTER 5

What the acquisition engines optimize

The bidding machines are built to learn from what arrives fast and often. Value that shows up late, or rarely, gets replaced by a stand-in.

IN SHORT

Value bidding works only when the machine sees enough value events, soon enough. Google's app campaigns need at least 10 conversions a day to bid for a return target. TikTok's Smart+ app campaigns offer value bidding on Android only. Ad-funded apps earn most of their first 60 days of ad revenue in the first week, so they feed the machines well. Subscription renewals arrive months later. That shapes how the engines work, but not where the money goes. In 2025 budgets grew fastest on iOS and in non-game apps, where signal is thinner. A buyer who splits spend too thinly can starve every campaign. And the average return a dashboard shows is not the return on the next dollar.

The machine needs dense signal

Platforms state how much signal their bidding needs. Google's app campaigns need at least 10 conversions a day, or 300 in 30 days, before a buyer can bid for a return target. The events it bids on should come from Google's Firebase kit ^[27]. Google also warns that editing a campaign before its first 100 conversions may disrupt learning. It suggests daily budgets of 10 to 50 times the bid, by goal ^[139] [CL-015](#). For its data-driven attribution model, Google recommends 200 conversions and 2,000 ad interactions in 30 days, though the model runs at any volume ^[140].

Other platforms set softer rules. TikTok lists three goals for its Smart+ app campaigns: installs, in-app events and value. In Smart+, value bidding is Android-only ^[141,142]. TikTok's standard app campaigns offer it on Android and iOS, and TikTok treats 50 conversions as the main sign that value-bidding learning is done ^[143,144]. Meta says an ad set leaves its learning phase once delivery is stable, which usually takes about 50 results in the week after the last big edit. It also advises combining similar ad sets, because many small ones each teach the system less ^[145]. AppLovin recommends a daily budget that buys at least 15 to 20 conversions a day. Its ad-revenue and blended return goals also require the app to use its MAX mediation ^[146].

Put these together and a pattern appears. The engines learn best from events that are common, early and valued. Installs and first opens pass that test. So do early ad views and early purchases in games. A yearly renewal does not [CL-016](#).

Three clocks for value

Value reaches an app on three clocks. Ad revenue is fast but shallow. AppsFlyer's panel finds in-app ad revenue reaches 89% of its day-60 total by day 7, partly because most players leave early ^[147]. That describes a 60-day horizon in one panel; it does not show that early revenue predicts lifetime profit. First payment can be fast too, since an annual plan is often paid up front. Renewal is slow. In RevenueCat's 2026 report, the median app kept only 28% of annual subscribers past their first year ^[148]. The first year is paid up front; most renewals never come. More than half of three-day-trial cancellations happened on the first day ^[5]. The report does not state the cohort window.

So an ad-funded game can send a reliable 7-day value to its bidding partners. A subscription app can send its first payment, but must guess the renewal or send a stand-in such as "started trial". Mintegral takes three quarters of its revenue from game advertisers ^[10], which fits the signal story.

The 2025 spending data do not fit it neatly. Acquisition spend grew 35% on iOS, where signal is thinnest, and slipped 1% on Android. It grew 18% in non-game apps and 3% in games. The fastest growth, 70%, came in shopping apps, largely on budgets from China-

based e-commerce apps ^[6]. Higher value per iOS user, the category mix and those shopping budgets explain the shift at least as well as signal speed does. The machines' view shapes how campaigns are run. It does not decide where the money goes.

Stand-ins are where optimization goes wrong

When the real outcome is slow, teams optimize a stand-in: a trial start, a finished tutorial, a day-3 purchase. The machine then finds people who do the stand-in. That works only while the stand-in predicts value, and it can fail quietly. A cheaper trial start can come from people who cancel on day one. Chapter 7 shows the public evidence on predicted lifetime value is weak, so few buyers can prove their stand-in still works.


Average return is not marginal return

Dashboards report average return: total revenue divided by total spend. Budget choices need marginal return: what the next dollar brings. On a curve that flattens, the two diverge. The first dollars reach the most responsive people. Later dollars reach the rest. The mix-model tools from Google and Meta both build diminishing returns into their models ^[135,136].

The model below lets you see the gap. It is a teaching model with made-up numbers, not an estimate for any channel.

FIGURE 7
A channel can average a healthy return while its last dollars lose money

SYNTHETIC INTERACTIVE MODEL



A teaching model with made-up parameters. Revenue follows a saturating curve of the kind used in open-source mix models ^[135,136]. Move the sliders to see average and marginal return diverge. The break-even spend is solved exactly and rounded to the nearest \$100. Nothing here estimates any real channel.

Limits. Synthetic. Real response curves must be estimated from experiments or calibrated models.

Package files: figures/F07.svg · data/figures/F07.csv

A channel can show a healthy average return while its last dollars lose money ^{CL-017}. That is why "scale what works" is risky advice when "works" is measured on averages.

Diversification is a claim to test

Two sound arguments pull in opposite directions. Spreading a budget across channels spreads risk and can reach people one platform misses. But each extra channel shares the budget, and each needs enough signal to learn. A RevenueCat post argues that adding channels early can dilute learning, slow optimization and add work ^[149]. Meta's advice to

combine small ad sets makes the same point inside one platform^[145]. The case against piling into one channel rests on overlap: people reached on two channels at once cost twice and add little. The example we found comes from a vendor selling a de-duplication tool and is stylised^[150]. Neither side has a public, controlled test for apps. Diversification is a hypothesis to test with holdouts. It is not a rule.

Remarketing and search: paying for intent you already had?

Remarketing is now 29% of app marketing spend^[6]. It targets people who already installed the app, and some would return without an ad. Randomized studies of web retargeting find real but modest lift. For one apparel retailer, a two-week ghost-ads experiment on Google's display network found retargeting raised site visits by 17.2% and purchases by 10.5%. The test system was built at Google^[151,152]. At another, switching retargeting on brought 14.6% more users back within four weeks, and the effect faded fast^[153]. Those are web shops, not apps.

Search ads on an app's own name raise a related question. At eBay, ads on its own brand name showed no measurable short-term benefit, because people searching for eBay found it anyway^[114]. A study of Bing searches in January 2014, by researchers then at Microsoft, found brand ads added only 1% to 4% of clicks when no rival bid on the name. For brands that faced rivals and chose not to advertise, rivals took 18% to 42% of their clicks, though the authors warn those brands may differ^[154,155]. So brand bidding can be defensive, and its value depends on who else is bidding. We found no public study, academic or vendor, of how often paid App Store ads on an app's own name take installs that would have come anyway.

None of this makes remarketing or brand search worthless. It means their attributed returns are especially likely to include demand that already existed CL-018. Holdout tests, not attribution reports, should set these budgets.

Promotions raise the same issue from another side. First-order discounts, referral credits and promo codes drive installs in delivery, ride-hailing, marketplace and fintech apps. A report can credit those installs to the ad that ran alongside the offer, and the discount cuts contribution. This research did not find public causal evidence on app promotions, so we flag it as an open question.

THESIS TEST · WHAT THE MARKET OPTIMIZES (Q1)

THESIS	Bidding engines are built to optimize events they can see early and claim; they optimize for profitable customers only where value is dense and quick to arrive.
SUPPORTING EVIDENCE	Minimum-signal rules for value bidding ^[27,139,145] ; value bidding on Android only in TikTok's Smart+ ^[142] ; 89% of 60-day ad revenue by day 7 in ad-funded apps ^[147] .
BEST COUNTEREVIDENCE	In 2025 acquisition spend grew on iOS (+35%) and in non-game apps (+18%), where signal is thinner, while Android and games barely grew ^[6] . Annual plans are often paid up front, which gives slow-value apps a fast first signal.
OTHER EXPLANATIONS	Value per user, category mix, China-based e-commerce budgets, promotions and changes in measurement coverage.
WHAT WOULD PROVE IT WRONG	Evidence that slow-value apps reach contribution payback as reliably as fast-value apps at similar scale using value bidding.

THESIS TEST · RETARGETING, SEARCH AND INCENTIVES (Q8)

THESIS	The attributed value of remarketing and brand search usually exceeds its causal value, because some of the credited demand already existed.
SUPPORTING EVIDENCE	Brand-name search ads showed no measurable short-term benefit in a large randomized test ^[114] . Web retargeting experiments found lifts of 10% to 17% ^[151,153] , and brand ads added 1% to 4% of traffic where no rival bid ^[154] .
BEST COUNTEREVIDENCE	For brands facing rivals that chose not to advertise on their own name, rivals took 18% to 42% of clicks. So brand search can protect real demand, though selection may explain part of the gap ^[154] . Vendors also report large gains for app re-engagement. Jampp's blog reports an 86% rise in new riders and a 30% campaign lift for a ride-hailing app under its always-on ghost-bid measurement. It gives no period, baseline or uncertainty ^[122] . Those results are self-published and unaudited.
OTHER EXPLANATIONS	Remarketing may matter more for apps with high early churn, where users would not come back without a prompt. Promotions can inflate both.
WHAT WOULD PROVE IT WRONG	Independent app remarketing holdouts showing incremental conversions close to attributed conversions.

CHAPTER 6

Creative, discovery and conversion

Many creative "tests" in app marketing are spend decisions made by an algorithm. The store tests are closer to real experiments, and the store's own rules are now in play.

IN SHORT

Apple's and Google's store-page tests split visitors at random and state their statistics. Custom product pages and most in-platform creative tools do not. Even a platform split test randomizes who is eligible, not who sees the ad, so a "winning" ad may just have reached better prospects. Public evidence on creative fatigue is thin. The US rules on paying outside the stores changed in 2025 and are still in court. That opens a web-to-app path with different costs and better measurement.

Store tests are the cleanest experiments most teams run

Apple's Product Page Optimization shows a random share of App Store visitors up to three variants of the original page ^[156]. It uses Bayesian statistics and shows results once there are at least five first-time downloads. That is a display rule, not a sample-size rule. A variant may be labelled better or worse once it reaches 90% confidence against the original ^[157]. Google Play's store listing experiments let the developer set the smallest effect worth detecting and the confidence level. They stop on their own after six months ^[158].

Custom product pages are a different tool. Apple allows up to 70 extra versions of a product page, each with its own link for use in ads. Its documents describe no random split and no significance test against a baseline ^[157,159] [CL-019](#). Apple also says developers see a lift of 2.5 percentage points in conversion, on average, when they send people to a custom page. The average default page converts 1.6%, so Apple calls this a 156% increase ^[159]. That comparison shows the problem. People sent to a custom page came through an ad or link aimed at them. People on the default page include everyone browsing the store. The gap mixes the page's effect with the traffic's intent. A custom page routes traffic. On its own it creates no random split, though a team can use one inside a designed test.

Store conversion benchmarks depend on the definition. Apple's own figure is the 1.6% average above, with no period or definition given ^[159]. Vendors that sell store optimization publish much higher page-view-to-install rates from their own panels. The 25% to 33% range often quoted in the industry could not be traced to the source we opened.

Ad creative: test or allocation?

On the ad side, two very different tools share the word "test". Meta's A/B test shows each version to a separate part of the audience, so nobody sees both ^[160]. Its Dynamic Creative mixes uploaded assets into many ads and reports results for all versions together. Meta says it should not replace a split test, and since June 2024 it may not be offered for new app-promotion ad sets ^[161].

Even a real split test has a catch. It randomizes who is eligible to see each version, not who actually sees it. The delivery system then shows each version to the people it predicts will respond. Braun and Schwartz, writing in the Journal of Marketing in 2025, call this divergent delivery. Each version reaches a different mix of people. So a platform split test mixes the effect of the ad with the effect of the targeting. Their evidence is one Meta test with a web lead form, plus a formal argument ^[162,163]. The winner is the winner for the audience it was given. It may not win with everyone [CL-020](#).

A team that wants to know which version will perform best on that platform can use the split test as it is. The authors say that is what the tool is built for. A team that wants to learn what message works cannot get that from a platform split test. Each version is measured on a different mix of people, even with a holdout ^[162,163]. Isolating the message needs a design that holds delivery fixed. A team that only wants spend to flow to what performs can use adaptive allocation, but should not call its output a finding.

FIGURE 8

Only some creative and store tests randomize, and fewer state their statistics

OBSERVEDOFFICIAL DOCUMENTATION

Tool	Assignment	Statistics	Unit	What it can identify
Apple Product Page Optimization	random split	stated (Bayesian, 90%)	store visitors	which page variant converts better
Google Play store listing experiments	random split	stated (MDE, confidence)	store visitors	which listing variant converts better
Apple Custom Product Pages	none (routing)	none stated	traffic from a linked ad	performance of a page for its own traffic only
Platform A/B test (e.g. Meta Experiments)	random split	platform-reported	people, then delivery	difference, but delivery can still diverge
Adaptive creative allocation	algorithm	none (prediction)	impressions	where the algorithm sends spend
Platform lift study	random holdout	stated, account-gated	people	extra conversions from ads
Vendor-reported holdout	vendor-stated	vendor-reported	people or devices	not independently verifiable: assignment and analysis unpublished

randomized

algorithmic allocation

assignment stated by vendor, not verified

no assignment

Seven tools that teams call tests, by how they assign users and what they can identify. Store tests randomize and state their statistics ^[157,158]. Custom product pages route traffic without a test, and Apple's own lift figure for them compares different traffic ^[157,159]. Platform splits can still diverge in delivery ^[162]. Lift studies are account-gated ^[125]; vendor holdouts are self-reported, and their assignment and analysis are not published ^[122]. Also: ^[160,161]

Limits. Meta's A/B test and Dynamic Creative are described from Meta's help pages read through a renderer; how delivery diverges inside a test is not documented in detail.

Package files: figures/F08.svg · data/figures/F08.csv

Creative fatigue is real in practice, thin in evidence

Every performance team manages creative decay. Yet we found no rigorous public estimate of its size for app ads. A September 2026 preprint by a data scientist at the agency group WPP proposes a way to screen for fatigue. It was tested only on 264 synthetic cases, and the author disclaims validation on real campaigns. It reports no real-world effect size ^[164,165]. Widely repeated figures on how fast click rates fall after repeated exposure could not be traced to any source we could open. We do not use them.

Formats and the rules on rewards

Rewarded video, playable ads and interstitials dominate game supply. Store rules draw a line between rewarding a user for something inside an app and paying them to install or rate other apps. Apple bars apps from making users install other apps, rate or review to unlock features, but lets apps reward in-app actions such as watching an ad ^[166]. Google Play's current policy bars attempts to manipulate an app's placement, including paid or incentivized ratings and reviews. It also bars apps whose main purpose is paying users to install other apps ^[167]. In 2017 Google also called incentivized installs a legitimate channel for some developers, if not used to game rankings ^[168]. The two stores draw the line in different places.

Deep links and the path back into the app

A deep link sends a user straight to the right screen. A deferred deep link does the same for someone who does not have the app yet. It carries the context through the store install and restores it on first open ^[169]. For commerce and content apps, this is often the difference between an install and a first order. It can also carry campaign context from a web page or an ad to an in-app event, within consent and platform limits. Routing a user is not the same as attributing the install.

Web-to-app: the store rules are moving

In April 2025 a US court barred Apple from charging commission on purchases that apps link out to on the web. It also barred Apple from steering users with discouraging screens ^[97]. In December 2025 the Ninth Circuit upheld the contempt finding and most of the rules against steering. It held that a total ban on any commission went too far and sent the fee question back to the trial court ^[97]. On 30 June 2026 the Supreme Court agreed to hear one question from Apple's appeal: whether a company can be held in contempt for breaking the spirit of an injunction ^[170,99]. The fee itself is being set by the trial court. In August 2026 that court refused Apple's request to pause the work ^[171]. Justice Kagan paused it for a day on 12 August, then denied Apple's application on 13 August ^[172,99] CL-021. At the cutoff no US link-out fee had been approved.

Google's US changes followed its own case. From October 2025 developers could tell US users about outside prices, link out and offer other billing. From July 2026 third-party stores gained access to Play's catalog unless a developer opts out. From October 2026 developers using outside billing or links must report those sales and pay service fees ^[98]. Epic and Google had asked the court to soften the original order as part of a settlement. They withdrew that request on 14 July 2026, so the original order still applies ^[173].

Web-to-app journeys matter for two reasons at once. They can change unit economics, because a web payment can cost far less than a store commission (chapter 7). They can also change measurement, because a web checkout sees the user, the campaign and the purchase in one place. Claims that web-to-app flows convert several times better than store installs circulate widely; we could not find a source to verify them.

New discovery surfaces

AI assistants are starting to surface apps. At its May 2026 developer conference, Google said app discovery in its Gemini assistant would roll out "in the coming weeks", alongside "Ask Play", a chat layer on Play search ^[174]. We could not confirm the rollout was complete by the cutoff. OpenAI opened app submissions and an in-ChatGPT app directory in December 2025, with apps called up by name or from a tools menu ^[175]. These are early surfaces to watch, not yet channels with measurable demand.

THESIS TEST · WINNING CREATIVE OR WINNING AUDIENCE (Q3)

THESIS	Many reported creative wins on ad platforms reflect delivery to different audiences rather than a controlled difference in what the ad does.
SUPPORTING EVIDENCE	Divergent delivery in platform split tests ^[162] ; Apple's own custom-page lift figure, which compares pages with different traffic ^[159] ; store tests that do randomize show how a clean comparison is built ^[157,158] .
BEST COUNTEREVIDENCE	Large creative differences can swamp audience effects, in which case adaptive allocation still finds the right winner.
OTHER EXPLANATIONS	Seasonality and fatigue can change results between test waves.
WHAT WOULD PROVE IT WRONG	Evidence that adaptive-allocation winners usually also win when compared across all assigned users.

CHAPTER 7

Retention, monetization and payback

Payback is decided by who stays, how they pay and what the store keeps. The first two are hard to predict. The third changed in 2026.

IN SHORT

Most installs do not last. In a panel of 16,000 games, the median title keeps about 22% of users to day 1 and under 1% to day 30. Subscription apps turn few installs into payers, and in one panel the median app keeps only 28% of annual subscribers past the first year. The first year is paid; most renewals never come. Predicted lifetime value fills the gap, but the public proof that those predictions hold up on app cohorts is thin. Meanwhile the store's cut has become a lever. Moving a first-year payment from Apple's 30% tier to web checkout keeps about a third more of it, before the cost of running your own billing. Measure payback on contribution, and keep observed value apart from forecast value.

Retention: a steep cliff and a long tail

GameAnalytics' 2026 benchmarks cover more than 16,000 mobile games with at least 1,000 monthly users across 2025. The median game kept about 22% of new users on day 1 early in the year, a figure that fell through 2025. It kept just under 4% on day 7 and 0.7% to 0.8% on day 30. The top 1% kept 64% to 68% on day 1 and 13% to 15% on day 30 ^[4].

Three limits apply. The report does not say whether "day 30" means users active on exactly that day or on any day since. It reports medians across games, not across players, so it describes the typical title, not the typical user, and many of those titles buy no ads. And the games in it are those that chose GameAnalytics. Compare your own day-30 figure to these only on the same definition and the same cohort age.


FIGURE 9

Game retention and subscription renewal measure different kinds of loss

OBSERVED


VENDOR PANELS

A · Mobile games: share of new users active on day N (GameAnalytics panel, 2025)



Log scale. Definition of 'day N' (exact day or since) not disclosed; medians across games.
 Median D1 is the early-2025 level; it declined during the year.
 Panel = games using GameAnalytics.

B · Subscription apps: conversion and churn (RevenueCat, Adapty panels)



Each bar has its own denominator. Do not compare bars as one funnel.

Panel A: GameAnalytics' 2025 benchmarks across more than 16,000 games, medians and top-1% ranges, on a log scale ^[4]. Panel B: subscription conversion and churn from RevenueCat's and Adapty's client panels, each with its own denominator ^[5,176]. None of these is a market average. Also: ^[148]

Limits. Vendor panels; retention definition not disclosed; medians across apps, not across users.

Package files: figures/F09.svg · data/figures/F09.csv

Subscriptions: few convert, many leave

RevenueCat's 2026 report draws on more than 115,000 apps and about \$16 billion of tracked revenue, mostly from 2025 ^[148]. In its panel, apps with a hard paywall turn a median 10.7% of downloads into payers within 35 days, against 2.1% for freemium apps. Trials of 17 to 32 days convert a median 42.5% of trial starts, against 25.5% for trials under four days. The median app kept 28% of annual subscribers after one year, for subscriptions started in 2024, down from 31% a year earlier ^[148]. RevenueCat's blog restates this as about 72% cancelling in year one. Its comparison with 56% the year before mixes two reports' methods, so it is not a trend ^[5]. The first year is paid up front, so the loss falls on renewal revenue. More than half of three-day-trial cancellations happened on the first day ^[5]. Adapty's panel of about 16,000 apps, with no period stated, puts average install-to-trial conversion at 10.9% and trial-to-paid at 25.6% ^[176].

Two findings matter for acquisition. First, failed payments. In RevenueCat's report, billing errors caused 32.2% of Google Play subscription cancellations, against 15.2% on the App Store; its summary gives 31% and 14% ^[148,5]. That is a billing problem, not a media problem, and fixing it can raise lifetime value without buying a single install. Second, early value and retention can pull apart. Across RevenueCat's categories, "AI apps" realized a median \$30.16 per paying user in the first year, against \$21.37 for other apps. But after 12 months they kept 6.1% of monthly subscribers, against 9.5% ^[148,5]. That is a comparison between categories, not evidence about how a bidding model ranks users. It is still a warning: a model that learns from early value alone can miss faster decay.

These panels describe each vendor's own customers. They are not market averages. Paywall design, price and category explain much of the spread, and the panels say nothing on their own about media quality.

Ad-funded apps get paid fast

AppsFlyer's monetization panel finds that in-app ad revenue reaches 89% of its day-60 total by day 7 ^[147]. Part of that speed is simply churn: most players are gone by then. It is still why ad-funded games suit today's bidding machines (chapter 5). The same report drew on \$7.2 billion of in-app ad revenue across its measured apps from January 2025 to March 2026. It also counted \$0.9 billion of store-verified one-off purchases and \$0.8 billion of store-verified subscriptions ^[147]. Those are panel totals, not market sizes.

Ad load links acquisition and monetization. More ads raise revenue per user today and can push users away tomorrow. Unity's study of eight apps found that users who watched rewarded ads were 4.5 times more likely to buy in-app ^[177]. That is a correlation among users who chose to engage. It cannot show the ads caused the purchases.

The best randomized evidence on ad load comes from outside games. Pandora randomized about 34 million listeners into nine ad-load levels for 21 months. After 21 months, listeners were about three times as sensitive to ad load as a one- or two-month test would have shown. Estimates from observational data were biased and sometimes pointed the wrong way. Heavier ad load also pushed more listeners to buy the ad-free plan, though more stopped listening altogether ^[120,178] CL-032. One author works at Sirius XM Pandora, which allowed publication on condition that the authors not discuss its policy implications. The paper is a working paper. The lesson for app teams is about method: short ad-load tests can understate the long-run cost, and correlations mislead.

Cancellation rules shape renewal too. The FTC's 2024 click-to-cancel rule would have made online cancellation as easy as sign-up. A federal appeals court struck it down in July 2025 on procedural grounds ^[179]. In February 2026 the FTC restored its narrower pre-2024 rule, and in March it asked for comment on changes. No proposal had followed by the cutoff ^[180,181,182]. So no federal click-to-cancel rule was in force.

The store's cut is now a lever

Store fees used to be fixed background. They are now a choice with real money attached.

PATH	RATE AT THE CUTOFF, UNLESS MARKED	SOURCE
Apple App Store, standard	30%; subscriptions fall to 15% after a subscriber's first paid year	[3]
Apple Small Business Program	15% for enrolled developers with up to \$1 million in prior-year proceeds, counted across associated accounts; crossing \$1 million in the year restores the standard rate	[183]
Apple, US link-out purchases	No approved commission: the appeals court recommended none until the trial court sets a fee; Apple proposed one in August 2026; Supreme Court review of a separate contempt question pending	[97,99]
Apple, EU, in force at the cutoff	Standard rates, unless a developer opted into Apple's EU alternative terms: 17% plus a 3% payment fee through Apple's in-app purchase, 10% plus 3% for small businesses and subscriptions after year one, and €0.50 per first annual install above one million	[184]
Apple, EU, announced, in force 1 October 2026	App Store: 26% through Apple's in-app purchase, 20% through other in-app payment, 15% on link-out sales within seven days; lower rates for small businesses and subscriptions after year one. Apps outside the App Store pay a 5% core technology commission instead	[185]
Google Play, US/UK/EEA from 30 June 2026	Auto-renewing subscriptions 10%. Other purchases 20% for users who installed from 30 June, 25% for earlier installs (20% through an external web link). A developer's first \$1 million a year: 10%. Add 5% when Google Play's billing is used. Program rates open only on 30 September 2026	[100,186]
Google Play, markets not yet moved	15% on the first \$1 million a year for developers enrolled in the 15% tier, 30% above; subscriptions 15%. Australia and Japan move on 30 September 2026	[186,100]

Take a \$9.99 monthly subscription. At Apple's standard first-year rate, the developer keeps \$6.99. At 15%, it keeps \$8.49. On a web checkout at Stripe's published US card rate of 2.9% plus 30 cents, it keeps about \$9.40 ^[187]. That is about 34% more than the 30% tier and about 11% more than the 15% tier CL-022. Put the other way, the web route can lose about a quarter of the store route's paying customers and still match its receipts against the 30% tier, at the same price. Against the 15% tier it needs to keep about 90%. Those figures count card fees only. Running your own billing adds subscription software, tax handling, chargebacks, fraud and support, which the store fee covers. And a web checkout adds steps, so fewer people may finish it. The net effect is something each app has to test.

Adoption is early and skewed to the largest apps. In RevenueCat's panel, web purchases were 3.2% of subscription revenue worldwide and 4.9% in North America. Some 41% of its top-tier apps earned web revenue, against 1.3% of the smallest ^[188].

Forecast value is a forecast

Because real value arrives late, many app teams bid on predicted lifetime value. The public evidence behind those predictions is weak. Google's published zero-inflated lognormal method was tested by its authors on a retail shopper dataset and a charity donor dataset, not on app install cohorts ^[189,190] [CL-023](#). A 2021 blog post from one attribution vendor, written before SKAdNetwork 4, quoted its technology chief. The post relayed customers' claims of reaching "the same level of scale" as under the old device ID, with no accuracy figures, sample or method ^[191]. Practitioners have long warned that backtests on old cohorts miss changes in competition and cohort size ^[192].


The published game evidence shows where models break. A game-data firm tested models on one freemium role-playing game, using about 2,500 paying players who had stopped playing. Its neural network's error averaged 5.7% of the largest single player's spend, against about 9% for a classic statistical model. That scaling flatters both. On a standard per-player measure, the errors were 74% and 96%. For the top fifth of spenders, the first measure gave 15.6% against 33.4% ^[121,193]. The classic models badly underestimated the biggest spenders, who in this game can bring in up to half of all revenue. Two recent preprints, neither peer reviewed, claim better game models. One, by Tencent staff, used the first week of play of about 3 million new players of one very large game to predict the next month's spend. Adding uncertainty estimates cut error on the top 5% of spenders from 48% to 19% for one model. Google's published method did worst on that measure, at 75% ^[194,195]. The second preprint was removed by arXiv's administrators because the submitter did not have the right to agree to its license. No full text is available, so we do not use it ^[196].

The minimum standard is simple. Backtest any model on mature holdout cohorts it never saw, with no future data leaking in. Report error by channel, platform, country, price point and value band, with ranges, not just an average. Track drift every month. When the product, price or channel mix changes, expect the model to break and re-check it.

FIGURE 10

A cohort can return 129% of its media cost in gross revenue in year one and still take 21 months to pay back

SYNTHETIC INTERACTIVE MODEL



A teaching model with made-up numbers, per install. Solid line: observed months. Dashed line: forecast, with a scenario band that widens with time; the band is an assumption, not a confidence interval. The red line is the fully loaded cost per install; every install is treated as an acquired user. The default payment path is Apple's: 30% in year one, 15% after a subscriber's first paid year ^[3]. Other paths use published rates ^[183,100] and Stripe's US

card rate of 2.9% plus 30 cents per payment for web checkout ^[187]. Change the observed window to see how much of payback rests on the forecast.

Limits. Synthetic. Retention and revenue curves are assumptions, not benchmarks.

Package files: figures/F10.svg · data/figures/F10.csv

Contribution, not ROAS

Return on ad spend hides several choices. It can use gross sales or net revenue. It can count ad revenue or not. The cost side can be media only, or fully loaded with fees, creative, measurement and staff. The same campaign can return more than 100% of media cost in gross revenue and still lose money on contribution after fees, refunds and overheads. The model above shows both views side by side. The metric dictionary in the appendix defines each term.

FIGURE 11

Business models differ in how fast value becomes visible to the bidding engines

SYNTHESIS

Model	Value to count	When value is visible	Store fee exposure	Main measurement risk
Games, ad-funded	ad revenue per user	● fast: 89% of day-60 ad revenue by day 7	low (ads not billed by stores)	retention cliff; ad-load trade-off
Games, in-app purchase	payer revenue; top spenders	● medium: payers appear in days, whales over months	high: 15-30% store fee	top-spender prediction error
Subscription	renewals net of fees and refunds	● slow: renewals at month 1, 12	high; web checkout now possible in the US	trial proxy; median app keeps 28% of annual subscribers past year 1
Commerce and marketplace	contribution from repeat orders	● medium: first order fast, repeat slow	low for physical goods	GMV is not revenue; discount buyers
Fintech	funded, active accounts	● medium to slow	low	consent-sensitive data; long verification
Streaming and media	subscribers or listening time	● slow; ad load shifts behavior over months	high if store-billed	short tests can understate long-run effects

● value visible fast

● medium

● slow

Synthesis; timing cells for commerce and fintech are the authors' inference.

App business models compared on the value to count, when it becomes visible, store fee exposure and the main measurement risk. Ad-funded timing is from AppsFlyer's panel ^[147]; subscription churn from RevenueCat ^[5]; top-spender error from a game LTV study ^[121]; long-run ad-load effects from Pandora's experiment ^[120].

Limits. A synthesis; commerce and fintech timing are the author's inference.

Package files: figures/F11.svg · data/figures/F11.csv

THESIS TEST · WHEN EARLY EVENTS PREDICT VALUE (Q4)

THESIS	Early events predict long-term value only while the product, price and acquisition mix stay stable; forecast value fails quietly when any of them shift.
SUPPORTING EVIDENCE	Categories where early value and retention diverge ^[5] ; Google's published LTV method was validated on retail and donor data, and did worst on top spenders in one game preprint ^[189,194] ; large forecast errors for top spenders in a game study ^[121] ; backtest limits ^[192] .
BEST COUNTEREVIDENCE	Ad-funded apps realize most of their first 60 days of ad revenue in the first week ^[147] , so their early signal is a better guide; annual plans paid up front give a fast first signal too.
OTHER EXPLANATIONS	Forecasts may be accurate inside large firms that do not publish them.
WHAT WOULD PROVE IT WRONG	Published backtests on app cohorts showing stable, low error across channel and price changes.

THEESIS TEST · ACQUISITION AND MONETIZATION TOGETHER (Q10)

THEESIS	Acquisition and monetization are easiest to improve together when one owner sees contribution per cohort, including the retention cost of ad load and price.
SUPPORTING EVIDENCE	Long-run ad-load effects three times the one-month estimate ^[120] ; the ad-engagement evidence from games that we verified is correlational ^[177] ; fee paths move net revenue per payment by 11% to 34%.
BEST COUNTEREVIDENCE	Platforms that run both demand and mediation can balance the two sides for the publisher.
OTHER EXPLANATIONS	In ad-funded games the same user is acquired and monetized within days, so the trade-off shows up fast.
WHAT WOULD PROVE IT WRONG	Randomized evidence that higher ad load or higher prices raise contribution without measurable retention loss across app categories.

CHAPTER 8

Quality, leakage and the expansion bets

Fraud takes credit it did not earn. Streaming TV can earn credit it cannot prove. Both are measurement problems before they are media problems.

IN SHORT

The fraud mechanisms that attribution companies describe include theft of attribution credit, through fake clicks and spoofed installs that claim users who came anyway, as well as fabricated activity. Attribution companies describe their defences. But none of the big ones appears on the industry's accreditation list for in-app fraud detection, and the pages we checked describe no standard refund process. Streaming TV is the favourite expansion story. Every method we found links TV exposure to app installs by household, IP or modelled matching. We found no independent randomized test of the effect on app installs. Commerce is a real bet for the biggest specialist, but its results are still reported only by the company.

Fraud that steals credit

Many of the fraud types the attribution companies describe take credit rather than fake impressions. Click spamming fires masses of clicks so that some land just before real installs. Click injection fires a fake click in the seconds before an install completes, to win last-click credit. SDK spoofing sends fake install or event messages with no real phone behind them. Device farms produce installs at scale. Faked in-app events claim cost-per-action payouts ^[197,198]. All of them exploit the credit rules of chapter 4.

FIGURE 12

How app fraud steals attribution credit, and when it is caught

OBSERVEDVENDOR DOCUMENTATION

Fraud type	Exploits	Detection (vendor-described)	Who bears the cost
Click spamming	last-click credit rules	● post-hoc: click-to-install timing	advertiser pays for organic installs
Click injection	last-click credit rules	● pre- and post-attribution: timing anomalies	honest network loses credit
SDK spoofing	trust in SDK messages	● pre-attribution: signed messages	advertiser pays for installs that never happened
Device and install farms	volume-based payouts	● heuristics: anonymous IPs, device reuse	advertiser; store ranking distortion
Attribution hijacking	deduplication gaps	● deduplication and timing review	honest network; advertiser misallocates
Fabricated in-app events	cost-per-action payouts	● signatures, behavior checks	advertiser; bidding models learn from fakes
Incentivized-install abuse	install and engagement metrics	● hard: users are real	advertiser; inflated user and session counts

● blocked before credit

● found after credit

● hard to detect

No major attribution company appears on MRC's in-app invalid-traffic accreditation list; no public standard refund process was found.

Fraud types, the credit rule each exploits, detection timing and who bears the cost, from attribution-company documentation and one academic study ^[197,198,199]. MRC's accreditation list is at ^[200].

Limits. Detection methods are vendor-described and unaudited; no market-wide fraud rate is implied.

Package files: figures/F12.svg · data/figures/F12.csv

Attribution companies describe detection by the timing between click and install, and by signed messages from their software kits ^[197,198]. Some block traffic before credit is given;

some flag it afterwards. What happens to the money is less clear. AppsFlyer's product page speaks of helping advertisers "reclaim wasted spend". But it and the other pages we checked describe no standard process, such as fixed dispute windows or automatic refunds, between advertisers and ad networks ^[197].

Accreditation tells a similar story. The Media Rating Council lists several firms accredited to detect sophisticated invalid traffic in mobile apps. They include DoubleVerify, HUMAN, Integral Ad Science, Pixelate, Protected Media and some of Google's and Meta's own services. AppsFlyer, Adjust, Singular and Kochava, the attribution companies app marketers rely on for install fraud, do not appear on it, checked on 27 September 2026 ^[200] **CL-024**. That does not mean their tools fail, or that no audit has ever taken place. It means the list we checked did not show them as accredited for this.

The one large public legal test ended without a court ruling on fraud. In 2017 Uber sued its agency Fetch Media. Uber said it had paid more than \$82.5 million for mobile ads. It alleged that some were not viewable or took credit for installs they did not cause. That federal case was dismissed later that year ^[201,202]. In a separate state case, Uber, the ad network Phunware and four individuals settled in October 2020 for \$6 million, all denying wrongdoing. The court had earlier struck Phunware's pleadings as a sanction ^[203] **CL-025**.

Incentives sit on a line. Paying users to install and open an app can inflate user counts and session numbers. An academic study found paid-install services on Google Play doing just that, with weak enforcement ^[199]. Rewarded ads inside an app the user already chose are a different thing, and both stores allow them (chapter 6).

Supply accountability is still partial

The industry's supply-chain files let buyers check who may sell an app's inventory. In a May 2023 post, the security firm HUMAN reported on bid requests its own product screened. For app and streaming TV requests, 25% came from apps with no known app-ads.txt file, against 3% on the web. About 9% came from sellers missing from the seller registry, against 8% on the web. One check favoured apps: unauthorized sellers were 5% of app requests, against 8% on the web ^[204,205]. Those are one vendor's figures, three years before the cutoff. On most of these checks, supply controls covered apps less well than the web.

Streaming TV: household credit is not lift


Streaming TV ads for phone apps are the expansion most vendors now pitch. The question is how an ad on a living-room screen gets tied to an install on a phone. We found four kinds of method in vendor documents.

- **IP matching.** AppsFlyer's cross-platform attribution links a TV ad view to later installs by probabilistic modelling on IP address alone. It can credit several installs by several people to one TV impression ^[206,207].
- **Viewing data plus an identity graph.** Kochava links Samba TV's automatic content recognition (ACR) data to a household and device graph. Samba TV says the data come from about 50 million smart TVs whose owners opted in. It measures lift with synthetic control groups built after the campaign, which are modelled rather than randomized ^[208]. ACR data carry privacy risk. In 2017 a TV maker paid \$2.2 million to settle US charges that it collected viewing data from 11 million TVs without consent ^[209]. The author works at Samba TV; see the disclosures.
- **A household graph.** MNTN credits a later website visit, not an app install, to a household that watched a full TV ad. It uses a graph it says covers 99% of available US households and does not disclose the matching method ^[210].
- **Platform conversion feeds.** Roku's ads manager records outcomes, including app installs, through a pixel and a server-to-server feed. It steers delivery toward households likely to convert, without saying how it matches them ^[211]. Moloco's streaming TV page describes household and IP matching, with credit running through the advertiser's attribution company ^[129].

Each method makes a modelled link between a TV or household exposure and a later event. It does not show who in the household watched. Matching alone cannot say how many installs the TV ad caused. That needs a control group: randomized, or a geographic or

synthetic comparison with its assumptions stated. Some vendors now build one in: Moloco describes a ghost-bid holdout for its streaming TV product ^[129]. But our searches found no independent, published randomized study of streaming TV ads on app installs at the cutoff. Every lift claim we located was designed and reported by the vendor ^[CL-026]. Moloco's launch release, for example, reports up to 1.5 times higher return on TV than mobile in early campaigns, with two-thirds of installs within six hours of exposure. It does not describe the test behind those numbers ^[212]. We did not examine every CTV vendor; Wurl, tvScientific and iSpot are among those not profiled here.

FIGURE 13
Four ways to link a TV ad to a phone install; matching alone does not show the ad caused it



*The author works at Samba TV. No independent randomized study of streaming TV effects on app installs was found at the cutoff.

The steps from a streaming TV impression to a credited app event, and the identity methods vendors document ^[206,208,210,211,129]. Only a randomized or geographic control shows effect. The author works at Samba TV.

Limits. Vendor descriptions; no vendor publishes an audited match rate.

Package files: figures/F13.svg · data/figures/F13.csv

Two things share the "CTV and apps" label. One is TV advertising aimed at a phone app outcome. The other is winning users for apps that run on the TV itself. The bridge above applies only to the first. The second is a closed loop on one device, which is easier to measure and smaller.

The best general causal study of TV advertising, covering 288 consumer brands, measured effects on store sales, not apps ^[123]. It shows such studies can be done. We found none for app installs.

Commerce: the biggest bet is still self-reported

AppLovin opened its self-serve Axon Ads Manager to e-commerce advertisers by referral on 1 October 2025, with a few hundred advertisers live. An earlier invite-only phase was limited to brands with at least \$10 million of gross merchandise value ^[213]. Management promised a public opening for June 2026. On 22 June the chief executive said any business could now sign up without a referral code, under the new name AppLovin Ads ^[214]. On its first-quarter call the chief executive offered a hypothetical: if 100,000 new customers joined in the first year, that would imply about \$7 billion of ad spend ^[215]. That is arithmetic on an assumption, not a target or a result. AppLovin's filings through its second-quarter 2026 report mention e-commerce as an area of growth but do not break out e-commerce revenue ^[21,24] ^[CL-033].

Moloco sells a separate commerce-media product that powers onsite ads for retailers, and lists Costco among its customers ^[216]. That is retail media, not app acquisition, and a different buyer.

For commerce apps, the value to count is contribution from repeat orders, not gross merchandise value. A marketplace's GMV is not its revenue, and its revenue is not its profit. A campaign that lifts first orders but fills them with discounted, one-time buyers can grow GMV and lose money.

THESES TEST · CTV AND COMMERCE EXPANSION (Q9)

THESES	Credible evidence that streaming TV creates extra app value needs a stated identity bridge and a randomized or geographic control; household matching alone shows exposure, not effect.
SUPPORTING EVIDENCE	Every CTV-to-app method we found uses household, IP or modelled matching ^[206,208,210] ; our searches found no independent causal study.
BEST COUNTEREVIDENCE	Vendors report fast install responses after exposure ^[212] , and some now build holdouts into their TV products ^[129] .
OTHER EXPLANATIONS	Matching can overstate TV's effect, because delivery is steered toward households already likely to convert. It can also understate it, if TV lifts app installs through search and organic discovery that no household match captures.
WHAT WOULD PROVE IT WRONG	A published geo or randomized test, run by a party without a stake in the media, showing app-install lift consistent with matched-household credit.

CHAPTER 9

The market instruments

A matrix can show what vendors document. It cannot show who performs best. This one is built to keep those two things apart.

IN SHORT

The vendor universe covers 37 companies and 64 products found during the research. Eleven execution products and five benchmark channels were scored on eleven dimensions, against anchors fixed before scoring, by two separate model passes. Bidding toward value is well documented. Fees are not: only one product publishes a fee schedule. An App Quotient sums the matrix up in bands of documented capability, but its order moves too much under other weights to read as a ranking. No product documents both of Apple's attribution frameworks in technical detail, and no specialist documents a standard lift test with its method. The big self-attributing channels document lift tests that, in the pages we read, run through account teams, and their automated campaigns choose placements for the buyer.

How the sample was built

The research started from the playbook's list of app platforms, used as search seeds, not as a verified peer list. It added companies found in filings, partner pages and product documents. Every row in the universe carries its parent company, roles, ownership and status as found, with sources. Parent companies and products are counted separately, so a company with four brands counts once as a company. The universe is a discoverable sample, not a census. Firms that publish little are under-represented by design, because the method only scores what can be seen, and firms that failed or closed were not studied.

Vendor universe

Every company and product encountered, with roles and ownership as found. A discoverable sample, not a census.

PRODUCT	PARENT	ROLES	OWNERSHIP	STATUS	SCOPE GROUP	SOURCES
Axon Ads Manager (AppDiscovery)	AppLovin Corporation	DSP, ad network	public:APP	active	1	V1-S01 V1-S06
MAX	AppLovin Corporation	mediation	public:APP	active	1	V1-S01 V1-S02
Adjust	AppLovin Corporation	MMP	subsidiary of AppLovin Corporation	active	2	V1-S01
Wurl	AppLovin Corporation	CTV ad network/distribution, SSP	subsidiary of AppLovin Corporation	active	1	V1-S01 V1-S08 V1-S09
Moloco Ads	Moloco, Inc.	DSP	private	active	1	V1-S17 V1-S18 V1-S20
Moloco Commerce Media	Moloco, Inc.	retail media platform, exchange	private	active	3	V1-S18
Moloco Performance CTV	Moloco, Inc.	CTV DSP	private	active	1	V1-S13 V1-S19
Liftoff Accelerate / Direct	Liftoff Mobile, Inc.	DSP	public:LFTO	active	1	V1-S21 V1-S28
Liftoff Monetize (Vungle Exchange)	Liftoff Mobile, Inc.	SDK, exchange, mediation-adjacent	subsidiary of Liftoff Mobile, Inc.	active	1	V1-S24
Unity Ads / LevelPlay ('Unity Grow')	Unity Software Inc.	ad network, mediation, DSP-like UA	public:U	active	1	V1-S29 V1-S30 V1-S31
ironSource Ads (direct-demand network)	Unity Software Inc.	ad network	subsidiary of Unity Software Inc.	unconfirmed (reported sunset April 2026)	1	
Mintegral	Mobvista Limited (HKEX: 1860)	DSP, ad network, ADX, SSP	subsidiary of Mobvista Limited (public parent)	active	1	V1-S37
DT Exchange / App Growth Platform (AdColony/Fyber lineage)	Digital Turbine, Inc.	ad network, exchange, mediation	public:APPS	active	1	V1-S43 V1-S44
DT ignite / On Device Solutions (ODS)	Digital Turbine, Inc.	OEM/on-device distribution	public:APPS	active	1	V1-S43 V1-S44
Pangle	ByteDance Ltd.	ad network, SDK, exchange	subsidiary of ByteDance Ltd. (private)	active	1	V1-S47
InMobi Advertising (DSP + Omnichannel Exchange + Ad Monetization)	InMobi Pte. Ltd.	DSP, ad network, exchange, mediation-adjacent	private	active	1	V1-S48 V1-S49 V1-S50
Kayzen	Ioniq Group (stake held by Shackleton Ventures' Victoria Fund since 2024)	DSP	subsidiary of Ioniq Group	active	1	V2-S01 V2-S02 V2-S03 V2-S04 V2-S05
Smadex	Entravision Communications Corporation (NYSE: EVC)	DSP, managed service	subsidiary	active	1	V2-S08
Adwake	Entravision Communications Corporation (NYSE: EVC)	managed service, ad network	subsidiary	active	1	V2-S08
Jampp	Affle 3i Limited (NSE/BSE: AFFLE; formerly Affle (India) Limited, renamed April 2025)	DSP	subsidiary	active	1	V2-S21 V2-S22 V2-S23
Bidease	none disclosed (independent, founder-led)	DSP	not-publicly-disclosed	active	1	V2-S17
RZR (formerly Aarki, rebranded 2026-03-17)	none disclosed (independent since 2023 Skillz spin-off — carried from prior public record, not re-verified in this research)	DSP	private	rebranded—RZR	1	V2-S18 V2-S33
Remerge	none disclosed (independent, founder-led)	DSP, MMP-adjacent (incrementality testing)	not-publicly-disclosed	active	1	V2-S06 V2-S07
YouAppi	Affle 3i Limited	DSP, retargeting	subsidiary	active	1	V2-S16
Adikteev	none disclosed (independent, founder/management-controlled; a 2025 management buyout was reported by third-party aggregators but not independently confirmed in this research)	DSP, retargeting	not-publicly-disclosed	active	1	V2-S14 V2-S15
Verve / Verve Dataseat	none (public: dual-listed VER [Nasdaq First North Premier, Stockholm] / VRV [Frankfurt]; renamed from Verve Group SE to Verve Group Media SE at the 2026-06-05 AGM, alongside approved relocation of the registered office from Sweden to Ireland)	exchange/SSP, DSP	public	active	1	V2-S09 V2-S10 V2-S11 V2-S12 V2-S13
InMobi DSP	InMobi Group (private; SoftBank invested 2011)	DSP, exchange/SSP, mediation	private	active	1	V2-S19 V2-S20
Zoomd	none (public: TSX Venture Exchange, ticker ZOMD)	ad network, mediation	public	active	1	V2-S25
Mobupps (incl. MAFO, MOBUPPSX, IRTB)	none disclosed (private)	ad network, DSP, exchange/SSP	not-publicly-disclosed	active	1	V2-S26
Persona.ly	none disclosed (private, independent)	DSP	not-publicly-disclosed	active	1	V2-S28
Chartboost	Zynga (2021-2024) → LoopMe (from 2024-12-10)	DSP, exchange/SSP, mediation, SDK	subsidiary	acquired by LoopMe	1	V2-S27
Applier Ad Cloud (incl. Retargeting, AIBID, AIXPERT, AdCreative.ai)	Applier Group	DSP, retargeting	validate	active	1	V2-S29
mediasmart	Affle 3i Limited (via Affle Iberia, S.L.)	DSP	subsidiary	active	1	V2-S30

RevX	Affle 3i Limited	DSP	subsidiary	active	1	V2-S31
Appnext	Affle 3i Limited (inferred; not an explicit on-page statement)	ad network, OEM/on-device distribution	subsidiary	active	1	V2-S32
Advantage+ App Campaigns	Meta Platforms, Inc.	ad network, self-serve platform	public:META	active	3	V3-S05 V3-S06
Google App Campaigns (incl. AI Max)	Alphabet Inc. / Google LLC	ad network, self-serve platform, DSP-adjacent (cross-network buying: Search, Display, YouTube, Play, Discover)	public:GOOGL	active	3	V3-S03 V3-S11
TikTok App Campaigns / Smart+ App Campaigns	ByteDance Ltd.	ad network, self-serve platform	private	active	3	V3-S01 V3-S08 V3-S09
Apple Ads (Apple Search Ads)	Apple Inc.	owned-inventory ad platform (App Store search/browse)	public:AAPL	active	3	V3-S02 V3-S12
Amazon DSP (app promotion, incl. events manager)	Amazon.com, Inc.	DSP	public:AMZN	active	3	V3-S04 V3-S10 V3-S13
Adjust	AppLovin Corporation	MMP	subsidiary of AppLovin Corporation (public:APP)	acquired by AppLovin	2	V3-S14 V3-S15
AppsFlyer	AppsFlyer Ltd.	MMP	private:VC-backed (General Atlantic, Qumra Capital, Pitango, Eight Roads, Salesforce Ventures)	active	2	V3-S21
Singular	Singular Labs, Inc.	MMP	private:unconfirmed detail	active	2	V3-S22
Kochava	Kochava Inc.	MMP	private:unconfirmed detail	active	2	V3-S23
Branch	Branch Metrics, Inc.	MMP, deep linking	private:unconfirmed detail	active	2	V3-S24
Airbridge	AB180 Inc.	MMP	private	active	2	V3-S25
Amplitude (product analytics + Feature/Web Experimentation)	Amplitude, Inc.	product analytics, experimentation	public (implied by market presence; ticker not independently reconfirmed on pages fetched in this research)	active	2	V3-S26
Mixpanel	Mixpanel, Inc.	product analytics, experimentation/feature flagging	private:unconfirmed detail	active	2	V3-S27
PostHog	PostHog Inc.	product analytics, feature flags/experimentation	private:VC-backed (YC W20); explicit no-sale statement	active	2	V3-S28
Firebase A/B Testing + Google Analytics for Firebase (GA4)	Alphabet Inc. / Google LLC	product analytics, experimentation	subsidiary product of Google	active	2	V3-S29
Statsig	Amplitude, Inc. (formerly OpenAI, formerly independent)	experimentation/feature flagging	subsidiary of Amplitude, Inc.	acquired by Amplitude (previously acquired by OpenAI)	2	V3-S18 V3-S19 V3-S20
RevenueCat	RevenueCat, Inc.	subscription ops / IAP infrastructure	private:unconfirmed detail	active	2	V3-S30
Adapty	Adapty	subscription ops / paywall A-B testing	private:VC-backed (Irrvnt, F1V)	active	2	V3-S31
Superwall	Nest 22, Inc.	subscription ops / paywall infrastructure	private:unconfirmed detail	active	2	V3-S32
Braze	Braze, Inc.	lifecycle/CRM	public (investor-relations program confirmed; ticker/exchange not captured on pages fetched in this research)	active	2	V3-S33 V3-S34
OneSignal	OneSignal, Inc.	lifecycle/CRM	private:unconfirmed detail	active	2	V3-S35
Airship	Airship	lifecycle/CRM	private:unconfirmed detail	active	2	V3-S36
CleverTap	WizRocket, Inc.	lifecycle/CRM	private (legal entity: CleverTap Private Limited)	active	2	V3-S37
Iterable	Iterable, Inc.	lifecycle/CRM	private:unconfirmed detail	active	2	V3-S38
Sensor Tower (incl. data.ai / formerly App Annie)	Sensor Tower, Inc.	ASO / market intelligence	private:unconfirmed detail (Sensor Tower itself); data.ai is a subsidiary brand of Sensor Tower	active; data.ai status: acquired by Sensor Tower	2	V3-S16 V3-S17
AppTweak	AppTweak	ASO	private:unconfirmed detail	active	2	V3-S39
Appfigures	Appfigures, Inc.	ASO	unconfirmed	unconfirmed	2	
AppLovin MAX	AppLovin Corporation	mediation	subsidiary product of AppLovin Corporation (public:APP)	active	2	V3-S14
Unity LevelPlay	Unity Software Inc.	mediation	public:U (NYSE, per general market knowledge; not independently reconfirmed in this research)	active	2	V3-S40
Google AdMob	Alphabet Inc. / Google LLC	mediation, ad network	subsidiary product of Google	active	2	V3-S41

Deep profiles were chosen to cover roles, not size. Five are larger platforms, most with owned supply or measurement: AppLovin, Moloco, Liftoff, Unity and Mintegral. Digital Turbine combines device distribution with an exchange. Kayzen and Jampp are self-serve

DSPs, Smadex is a managed-service DSP and Remerge is a retargeting specialist. Verve Dataseat is the mobile DSP of an omnichannel group. The author was an executive at Verve Group from 2022 to 2024. Verve Dataseat was scored under exactly the same rules as every other product, by both passes. Its n/e cells mark documents neither pass could find, such as fee terms behind a "contact us" form.

How scores work

Each product was scored on eleven dimensions, each with observable anchors written before scoring. A 3 needs documented, live capability of the strongest kind: for example, a campaign-management API rather than a reporting API, or log-level export rather than dashboards. A 0 means documented absence. "n/e" means the public evidence was not enough to score, and it is never treated as zero. "n/a" means the dimension does not apply. Each cell also records availability: live, beta, limited, managed-only or announced. It also records an evidence grade. A means technical documents, API references or filings; B means vendor marketing or press releases; C means third-party reporting. The grade records the type of source, not the strength of a performance claim: an API reference is strong proof that an endpoint exists and says nothing about results. The matrix marks cells that are beta, limited or managed-only with an asterisk. That matters most for lift tests, where a well-designed product may still sit behind an account team.

Four rules kept the scoring even. A parent company's product does not count for a sibling product. A partner-directory listing does not prove a feature. A reporting API is not campaign control. And the same standard applies to incumbents, challengers and any firm connected to the author.

Two model passes scored the panel. The first built the profiles. The second scored the same products blind, from sources it opened itself, without seeing the first scores. Across 176 product-dimension cells, the two passes gave the same value in 94 (53%). Where both gave a number, they were within one point in 93 of 97 cells. 50 disagreements were between a score and n/e, which usually means one pass found a document the other did not. Disagreements were settled by the published rules in the reconciliation script, and every disagreement is listed in the package. Fact-check corrections changed four cells directly: Meta's and Google's publisher visibility, Jampp's lift-test score and Amazon's buyer control. Meta's and TikTok's campaign-API scores gained primary sources. This is two passes by AI agents under the author's direction. It is not independent human review, and the paper does not claim inter-rater reliability.

FIGURE 14 · INSTRUMENT

App growth execution matrix: documented capability, not performance

OBSERVED

SCORED BY RUBRIC

11 products, 11 dimensions · click a row for the profile

PRODUCT	OBJECTIVES	RE-ENGAGE	IOS PRIVACY	CREATIVE	VISIBILITY	FEES	LIFT TESTS	EXPORT	CONTROL	CTV-APP	NON-GAMING	AQ
AppLovin · Axon Ads Manager INTEGRATED	3 ^A	n/e	1 ^A	2 ^B	2 ^A	1 ^A	2 ^{B*}	2 ^A	3 ^A	2 ^B	2 ^B	71 DOCUMENTED 2/3 · 3/3 · 5/5
Moloco Ads INTEGRATED	3 ^A	3 ^A	2 ^A	2 ^A	3 ^B	1 ^A	2 ^{B*}	3 ^A	3 ^A	3 ^B	2 ^B	80 HIGH 3/3 · 3/3 · 5/5
Liftoff Accelerate INTEGRATED	3 ^B	2 ^B	1 ^B	2 ^A	2 ^A	1 ^A	1 ^{B*}	2 ^A	3 ^{A*}	n/e	3 ^A	70 DOCUMENTED 3/3 · 2/3 · 5/5
Unity Ads (Grow) INTEGRATED	3 ^A	n/e	1 ^A	3 ^A	2 ^A	1 ^A	n/e	2 ^A	3 ^A	n/e	1 ^A	72 DOCUMENTED 2/3 · 2/3 · 4/5
Mintegral INTEGRATED	3 ^A	2 ^B	n/e	2 ^B	3 ^A	1 ^A	n/e	2 ^A	3 ^A	n/e	3 ^A	79 HIGH 3/3 · 1/3 · 4/5
Digital Turbine · DT Exchange / AGP DEVICE + EXCHANGE	3 ^A	1 ^B	2 ^A	n/e	3 ^A	1 ^A	1 ^A	2 ^A	2 ^{A*}	n/e	2 ^B	66 PARTIAL 2/3 · 2/3 · 5/5
Kayzen SELF-SERVE DSP	2 ^B	2 ^A	1 ^A	2 ^B	3 ^A	3 ^A	1 ^A	2 ^A	3 ^A	2 ^A	2 ^B	69 DOCUMENTED 3/3 · 3/3 · 5/5
Smadex MANAGED DSP	2 ^A	2 ^A	n/e	2 ^B	2 ^A	2 ^A	n/e	n/e	1 ^A	1 ^A	3 ^A	66 PARTIAL 3/3 · 2/3 · 3/5
Remerge RETARGETING	3 ^A	3 ^A	2 ^A	3 ^{A*}	2 ^A	1 ^A	2 ^{B*}	2 ^A	1 ^B	n/e	2 ^B	76 DOCUMENTED 3/3 · 2/3 · 5/5
Jampp SELF-SERVE DSP	3 ^B	2 ^B	1 ^B	2 ^B	2 ^B	n/e	2 ^B	3 ^B	2 ^B	2 ^B	2 ^B	71 DOCUMENTED 3/3 · 3/3 · 4/5
Verve Dataseat ♦ OMNICHANNEL	3 ^B	2 ^B	2 ^B	n/e	2 ^B	2 ^B	1 ^B	n/e	2 ^B	2 ^B	2 ^B	71 DOCUMENTED 2/3 · 3/3 · 4/5

3 3 documented, strongest anchor
 2 2
 1 1
 0 0 documented absence
 n/e n/e not enough evidence
 n/a n/a

Superscript = evidence grade (A technical doc or filing, B vendor page, C third party). * = beta, limited or managed-only.

AQ App Quotient, bars in order: █ Depth █ Reach █ Standing. Numbers under the bars: fields scored per group. Read the band, not the rank. ♦ author has a declared interest.

Scores follow the fixed anchors in chapter 9 and measure what public documentation shows at the cutoff, not how well a product performs. n/e means we could not find enough public evidence; it is not a zero. The full rubric, both scorers' results and every source are in the profile rows and the CSV.

CORRECTION, 28 SEPTEMBER 2026

Four non-gaming evidence scores were raised after publication: Kayzen, Remerge and Verve Datasat from 1 to 2, and Smadex from 1 to 3. In each case the second scorer had documented enough named non-gaming cases, and for Smadex a filing, to meet the anchor, but the conservative reconciliation rule had kept the lower score. The same test was applied to every non-gaming cell that rule had lowered. The App Quotient and its bands were recomputed. The corrections ledger records this as COR-117.

What the matrix shows

With the sources from both passes combined, 17 of 121 product-dimension cells (14.0%) still lack enough public evidence to score and 104 are scored. Of the 104 scored cells, 62 (59.6%) rest on grade-A evidence: developer documents, API references or filings. The five benchmark channels are scored separately and sit outside these counts.

Optimization is the best-documented capability. 9 of 11 products document bidding toward value or return on ad spend.

Fees are the least documented. Only Kayzen publishes a fee schedule; 7 products state a pricing model and little more.

No product documents both SKAdNetwork 4 and AdAttributionKit at the level of technical documents. The best iOS scores are 2: Digital Turbine, Moloco, Remerge, Verve Datasat.

No specialist documents a standard lift test with its method at the top anchor. 4 document one at the middle anchor, on request, as a managed service or without a published method: AppLovin, Jampp, Moloco, Remerge. Meta, Google and TikTok document lift tests with test and control groups, though they run through account teams.

Buyer control and data export are uneven. Campaign-management APIs: AppLovin, Kayzen, Liftoff, Mintegral, Moloco, Unity. Log-level export: Jampp, Moloco. Documented publisher and placement visibility with block lists: Digital Turbine, Kayzen, Mintegral, Moloco.

The App Quotient

The matrix also rolls up into one summary number, the App Quotient (AQ). It is built the same way as the Contextual Quotient in the author's contextual research ^[217]. AQ is a quotient, not a sum. It has three groups, chosen by the author. Depth is what the engine can do for growth: optimization objectives, re-engagement and creative. Reach is channel and category coverage: iOS privacy measurement, streaming TV to app, and non-gaming evidence. Standing is whether a buyer can check and control it: placement visibility, fees, lift tests, reporting export and buyer control.

Each group is the mean of whatever is scored inside it, so an unscored dimension is not counted as a zero. That also means group scores can rest on different amounts of evidence, so each score shows how many fields back each group. The groups combine at 45/25/30. The result is then pulled toward the field mean, as if each product had four extra cells at that mean. The weights and the four extra cells are the author's assumptions, not a calibrated reliability adjustment. The pull reduces the effect of a thin record; it does not remove the bias from fields that are missing. AQ is scaled so that 100 would mean every cell at the strongest anchor with no pull; with the pull toward the current field mean, a product scoring 3 on all eleven fields would get about 92.2. The three bars under each score are Depth, Reach and Standing, in that order.

Method note, 28 September 2026. AQ's field average pools every group score equally, while AQ weights the groups 45/25/30, which can lift or lower a whole field: six identical products that are strong only on Depth would all move the same way. The corrected average is the mean of the products' own weighted scores. Recomputed that way, the field average rises from 2.12 to 2.17 on the 0–3 scale, every AQ moves by at most one point, and the top three (Moloco, Mintegral, Remerge) keep their order. Because AQ's cut-offs are fixed

points on the scale, one product would change band: Digital Turbine, from Partly documented to Documented. The published scores stand; the author's later papers use the corrected average.

Read the band, not the rank. Bands sit at anchor averages on the 0–3 scale: high documented capability from 77.8 (an average of 2.33), documented from 66.7 (2.0), partly documented from 55.6 (1.67), and thin below. The bands on these inputs are high documented capability: Moloco, Mintegral; documented: Remerge, Unity, Verve Dataseat ♦, AppLovin, Jampp, Liftoff, Kayzen; partly documented: Digital Turbine, Smadex. Mintegral clears the top band's cut-off by only 0.76 points, and its Reach score rests on 1 of 3 fields. The ♦ marks a product the author has a declared interest in. Under four reasonable reweightings (equal weights, then each group raised in turn), Moloco stays in the highest band under all four, though not always first; Mintegral drops a band under the Standing-led weights; Remerge joins it under the Depth-led weights. Mid-table positions move more: by up to 7 places in this eleven-product table, and by 1.1 places on average. Across the four, 6 of 44 product results change band. The order is not a league table and is not presented as one. AQ sums up public documentation under the author's chosen groups and weights; it does not measure campaign performance. It inherits every limit of the matrix, and because its inputs are documentation scores, it still favours the vendors that publish the most. Every input is in the CSVs in the appendix, so the index can be recomputed on different weights.

A buyer should use the matrix, not the index, to decide. It shows which claims a vendor documents, links the source, and marks what to ask for on the rest.

The benchmark channels

Meta, Google, TikTok, Apple Ads and Amazon answer the same buyer questions in different ways. All five document campaign-management APIs, not only reporting ^[218,219,104,220,221]. Meta, Google and TikTok document lift tests with test and control groups that run through account teams ^[125,126,127]. Their automated app campaigns choose placements for the buyer. Google's placement report names the apps, publishers, websites and videos that showed ads, but gives impressions only and folds low-volume placements into one line ^[222]. TikTok's Smart+ does not allow manual placement choice ^[142]. Meta lists placement among its reporting breakdowns but describes no breakdown by publisher app ^[223]. TikTok documents full SKAdNetwork 4 support ^[224]. Its Smart+ app campaigns offer value bidding on Android only; its standard app campaigns offer it on iOS too ^[142,143]. Apple Ads uses its own direct attribution interface alongside AdAttributionKit ^[85]. Amazon's February 2024 app-event tooling covered Android and Fire devices. By September 2026 AppsFlyer described a closed Amazon test for iOS, and we found no Amazon page stating current iOS availability ^[225,226].

Benchmark channel comparison

The same buyer questions asked of the big self-attributing channels. They do not share supply, controls or measurement with app DSPs, so they are compared in prose fields and a few scores only.

CHANNEL	OBJECTIVES	OS MEASUREMENT	TRANSPARENCY	CONTROL	EXPORT	EXPERIMENTS	SCORES	SOURCES
Meta Advantage+ App Campaigns	Install and app-event optimization documented, including a minimum-ROAS bid strategy in the Marketing API.	SKAdNetwork campaigns and Aggregated Event Measurement documented; SKAdNetwork 4 and AdAttributionKit specifics not found in the documents read.	Placement is a reporting breakdown at surface level (for example Feed or Audience Network); no breakdown by publisher app described. Setup targeting limited to OS, country and language; some limits sit in account settings.	Self-serve UI plus Marketing API that creates and edits campaigns, including app-promotion objectives; no placement selection.	Marketing API / Insights API provide aggregated, scheduled export; no impression-level export documented.	Conversion Lift: documented RCT methodology and API, but gated to advertisers with Meta account-team access, not self-serve.	D1:3 D3:1 D5:1 D7:3 D8:2 D9:3	V3-S05 V3-S06 FX1-S23 OA1-S06 RS-S128 RS-S129 RS-S130
Google App Campaigns	tCPI, tCPA, tCPre, Maximize Conversions, and tROAS are documented live bid strategies, i.e., value/ROAS-based optimization is available.	iOS measurement combines modelled conversions, data from App Attribution Partners and a SKAdNetwork report; the Ads API exposes SKAdNetwork installs, fine conversion value and user type. AdAttributionKit is not mentioned in the pages we read [RS-S136; FX1-S24].	App campaigns placement report names apps and publishers, websites and videos, with impressions only; Google-owned inventory grouped; low-volume placements folded into 'Total: other'. No results by placement.	Self-serve UI plus Google Ads API supporting campaign creation, CampaignBudget objects, and bidding-strategy objects; legacy installs-only subtype is UI-only.	The Google Ads API returns aggregated App-campaign reporting at campaign, ad-group and asset level, including SKAdNetwork segments, but has no placement view. The placement report (impressions only) is built in Google Ads' Report editor. No impression-level export [V3-S11; FX1-S24; OA1-S01].	Conversion Lift: documented user-based and geo-based RCT designs, gated to advertisers with Google account-team access, not self-serve.	D1:3 D3:2 D5:0 D7:3 D8:2 D9:3	V3-S03 V3-S07 V3-S11 OA1-S01
TikTok App Campaigns / Smart+	Smart+ documents install (MAIL), in-app event (AEO), in-app value (VBO) goals, with VBO Android-only in Smart+; standard app campaigns offer VBO on Android and iOS.	Official help page (last updated Feb 2025) confirms full SKAN 4.0 support (fine/coarse conversion values, three postback windows, crowd anonymity tiers, 4-digit source_id); no AdAttributionKit mention found.	Manual placement selection 'is not supported ... at this time,' defaulting to automatic placement selection; no placement-level reporting documented.	Self-serve UI plus Marketing API (CampaignCreationApi) supporting APP_PROMOTION objective, budget, and budget_mode.	TikTok Marketing API Reporting API supports synchronous/asynchronous aggregated reports; no log-level export confirmed.	Conversion Lift Study: documented RCT (test/control) methodology supporting app installs as an objective, but run as an 'exclusive managed service,' not self-serve, with unspecified minimum spend.	D1:3 D3:2 D5:0 D7:3 D8:2 D9:3	V3-S01 V3-S08 V3-S09 OA1-S12 OA1-S13
Apple Ads (Apple Search Ads)	Max cost-per-tap bidding or Maximize Conversions toward a target cost per tap-through install, with optional CPA caps; no post-install event or value optimization documented.	AdServices: the app's attribution token must be exchanged with Apple's server within 24 hours, and Apple credits a tap up to 30 days or a view up to 24 hours before the download (view-through since March 2025). From September 2026, campaigns using age or gender targeting get no attribution [D-S20; D-S21; C-S10]. Apple Ads registered with AdAttributionKit on 10 April 2025, through SKAdNetwork versions 1 to 3 for click-throughs only [V3-S02; RS-S151].	Four named placements (Today tab, Search tab, search results, product pages) reported per campaign and ad, plus keyword and search-term reporting and negative keywords.	Campaign Management API (OAuth2 REST) documented for programmatic creation/management of campaigns, ad groups, and keywords.	The Apple Ads Platform API pulls reports programmatically; reports cover the four placements per campaign and ad, plus keywords and search terms, as aggregated metrics. No impression-level export [V3-S12; RS-S153; RS-S150].	No incrementality or lift-test offering appears in the Apple Ads help pages we read, unlike Meta, Google and TikTok [RS-S151; RS-S152].	D1:3 D3:2 D5:3 D7:n/e D8:2 D9:3	V3-S02 V3-S12 RS-S146 RS-S149 RS-S150
Amazon DSP (app promotion)	Events manager (February 2024) documents CPI, CPA and ROAS optimization for Android and Fire devices; a closed iOS test using SKAdNetwork is described only by AppsFlyer (third party, September 2026).	Amazon's February 2024 announcement scopes events manager to Android and Fire TV/tablet devices, and no Amazon page documents SKAdNetwork or AdAttributionKit support. AppsFlyer (September 2026) describes a closed Amazon DSP iOS test using SKAdNetwork [V3-S10; OA1-S17].	Amazon's DSP page cites 'transparent reporting', an inventory hub and deals, but no site- or app-level reporting or block and allow lists are documented in the pages we read [V3-S04].	Self-serve and managed service (managed requires \$50,000 minimum spend); DSP Campaign Management API documents campaign creation and bid and budget changes.	Amazon Marketing Cloud lets advertisers query hundreds of event-level fields (ad-attributed impressions, clicks, conversions) in a clean room, but outputs are aggregated and anonymous only. DSP advertisers request access through their Ad Tech Account Executive [V3-S13].	Amazon DSP materials mention brand-lift studies, offline sales-lift insights and marketing mix modelling for eligible advertisers; no app-install lift option is documented in the pages we read [V3-S04].	D1:3 D3:0 D5:n/e D7:n/e D8:2 D9:3	V3-S04 V3-S10 V3-S13 OA1-S16 OA1-S17

The systems around the media

Attribution, testing, analytics, subscription, lifecycle, store-intelligence and mediation tools shape app growth without buying media. Ownership matters here, because independence is part of what a buyer pays for. Adjust is owned by AppLovin^[227]. Google owns the AdMob mediation platform and the Firebase testing and analytics tools alongside its app campaigns^[228,28]. Statsig was bought by OpenAI, and its customer business then passed to Amplitude within about eight months^[56,57]. Data.ai joined Sensor Tower in 2024^[229]. AppsFlyer, Kochava, Adjust and Singular all document incrementality products, and their methods differ^[230,231,232,233]. Adjust's InSight compares an app with a synthetic control built from other apps, so its baseline is modelled. Singular's Audience Incrementality holds back a control group from an audience; its FAQ does not say the split is random.

Complementary systems map

Attribution, experimentation, analytics, subscription, lifecycle, store intelligence and mediation tools. Roles and ownership only; these are not scored against the media platforms.

SYSTEM	CATEGORY	PARENT	OWNERSHIP	INDEPENDENCE	INCREMENTALITY PRODUCT	SOURCES
AppsFlyer	MMP	AppsFlyer Ltd.	private:VC-backed (General Atlantic, Qumra Capital, Pitango, Eight Roads, Salesforce Ventures)	Not owned by an ad network, mediation platform, or ad-tech DSP; investor list includes Salesforce Ventures (CRM-adjacent, not an ad platform).	Documented: an 'Incrementality' product is named under AppsFlyer's Measurement Suite; methodology detail not verified beyond the product listing.	V3-S21
Adjust	MMP	AppLovin Corporation	subsidiary of AppLovin Corporation (public:APP), acquisition announced February 2021	Owned by an ad platform: AppLovin also operates the MAX mediation platform and the AppDiscovery/Axon ad network. An MMP under the same parent as a mediation layer and ad network is a documented ownership structure that raises an independence question for advertisers using Adjust to arbitrate between AppLovin's own inventory and competitors' — flagged as a conflict-of-interest consideration, not proven misconduct.	InSight (sold as an Adjust Growth Solution): compares the app with synthetic control groups built from other apps, so the baseline is modelled, not randomized (OA4-S01)	V3-S14 V3-S15
Singular	MMP	Singular Labs, Inc.	private; ownership/investor detail not disclosed on pages fetched in this research	No ad-platform ownership found.	Audience Incrementality: holds back a control group from an audience; the FAQ does not say the split is random (OA4-S02)	V3-S22
Kochava	MMP	Kochava Inc.	private; ownership/investor detail not disclosed on pages fetched in this research	No ad-platform ownership found.	Documented: Marketing Mix Modeling ('Always-On Incremental Measurement'), a named 'Incrementality Testing' product, and 'Media Lift Studies.'	V3-S23
Branch	MMP / deep linking	Branch Metrics, Inc.	private; ownership/investor detail not disclosed on pages fetched in this research	No ad-platform ownership found.	Not documented on the About page fetched in this research (n/e, not a confirmed absence).	V3-S24
Airbridge	MMP	AB180 Inc.	private (South Korea)	No ad-platform ownership found; positioned specifically for subscription-app measurement.	Not documented on the homepage fetched in this research (n/e, not a confirmed absence).	V3-S25
Amplitude	product analytics / experimentation	Amplitude, Inc.	public company (implied by market presence and its acquisitive role in the Statsig deal; ticker not independently reconfirmed on pages fetched in this research)	Not owned by an ad platform; acquired the Statsig product/business on May 5, 2026, expanding its experimentation footprint (see Statsig row).	n/a (product-analytics/experimentation platform, not an ad-measurement incrementality product for media buying).	V3-S26 V3-S19
Mixpanel	product analytics	Mixpanel, Inc.	private; ownership/investor detail not disclosed on pages fetched in this research	No ad-platform ownership found.	n/a (product-analytics/experimentation platform).	V3-S27
PostHog	product analytics / experimentation	PostHog Inc.	private:VC-backed (Y Combinator W20); explicit stated intention not to sell the business	No ad-platform ownership found.	n/a (product-analytics/experimentation platform).	V3-S28
Firestore A/B Testing + Google Analytics for Firebase (GA4)	experimentation / product analytics	Alphabet Inc. / Google LLC	subsidiary product of Google	Owned by the same company that sells Google App campaigns — an app-experimentation/analytics product owned by an ad platform.	n/a (product-analytics/experimentation platform; Firestore A/B Testing integrates with GA4 event data, not an ad-media incrementality product).	V3-S29
Statsig	experimentation	Statsig, LLC (acquired by OpenAI, 2025); Amplitude holds its customer contracts, trade name and a non-exclusive technology licence (May 2026)	see parent_org	Two ownership changes within roughly eight months: OpenAI's acquisition (reported Sept 2025) was followed by OpenAI divesting the business, with Statsig 'joining the Amplitude family' on May 5, 2026 per Statsig's own blog; the original founding team moved to OpenAI, and new leadership (Chris Yu, VP of Product) now runs Statsig under Amplitude. A TechCrunch article dated Aug 5, 2026 refers to Statsig as OpenAI's 'once-subsiary.'	n/a (product-analytics/experimentation platform).	V3-S18 V3-S19 V3-S20
RevenueCat	subscription ops	RevenueCat, Inc.	private; ownership/investor detail not disclosed on pages fetched in this research	No ad-platform ownership found.	No dedicated incrementality product named; site references helping advertisers 'get more out of your MMP investment' via post-conversion analysis, which is MMP-adjacent but not itself an incrementality method.	V3-S30
Adapty	subscription ops	Adapty	private:VC-backed (Irrvnt, F1V)	No ad-platform ownership found.	Not documented (product is A/B testing of paywalls/pricing, not ad-media incrementality).	V3-S31
Superwall	subscription ops / paywall	Nest 22, Inc.	private; ownership/investor detail not disclosed beyond legal entity name	No ad-platform ownership found.	Not documented (product is paywall A/B testing, not ad-media incrementality).	V3-S32
Braze	lifecycle/CRM	Braze, Inc.	public company (investor-relations program confirmed via investors.braze.com; ticker/exchange text not captured in the excerpts retrieved in this research)	No ad-platform ownership found.	Not documented in sources opened in this research.	V3-S33 V3-S34
OneSignal	lifecycle/CRM	OneSignal, Inc.	private; ownership/investor detail not disclosed on pages fetched in this research	No ad-platform ownership found.	Not documented in sources opened in this research.	V3-S35
Airship	lifecycle/CRM	Airship	private; ownership/investor detail not disclosed on pages fetched in this research	No ad-platform ownership found.	Not documented in sources opened in this research.	V3-S36
CleverTap	lifecycle/CRM	WizRocket, Inc.	private (legal entity: CleverTap Private Limited)	No ad-platform ownership found.	Not documented in sources opened in this research.	V3-S37

Iterable	lifecycle/CRM	Iterable, Inc.	private; ownership/funding/IPO status not disclosed on pages fetched in this research	No ad-platform ownership found.	Not documented in sources opened in this research.	V3-S38
Sensor Tower (incl. data.ai / formerly App Annie)	ASO / app-store market intelligence	Sensor Tower, Inc.	private (Sensor Tower); data.ai is now branded 'a Sensor Tower company'	Confirmed acquisition via data.ai's own site ('We've Been Acquired by Sensor Tower'); exact acquisition date not confirmed via a primary source opened in this research (commonly reported in the trade press as 2023-2024, not independently verified here).	n/a (market-intelligence/ASO data provider, not an incrementality-testing product).	V3-S16 V3-S17
AppTweak	ASO	AppTweak	private; ownership/investor/HQ detail not disclosed on pages fetched in this research	No ad-platform ownership found; notably offers an Apple Search Ads management layer alongside organic ASO tools.	n/a (ASO intelligence tool, not an incrementality product).	V3-S39
Appfigures	ASO	Appfigures, Inc.	unconfirmed — official site returned HTTP 403 to WebFetch in this research	Not verified in this research; see open_questions.	not verified in this research	
AppLovin MAX	mediation	AppLovin Corporation	subsidiary product of AppLovin Corporation (public:APP)	Same parent as Adjust (MMP, see above) and AppDiscovery/Axon (ad network) — a vertically integrated mediation + MMP + ad-network stack under one owner.	n/a (mediation layer, not itself an ad-measurement incrementality product).	V3-S14
Unity LevelPlay	mediation	Unity Software Inc.	public company product (NYSE:U per general market knowledge; not independently reconfirmed on the page fetched in this research)	Owned by Unity, which also sells Unity Ads (a demand-side ad network) — mediation and network demand under one parent.	n/a (mediation layer).	V3-S40
Google AdMob	mediation	Alphabet Inc. / Google LLC	subsidiary product of Google	Owned by the same company that sells Google App campaigns and Firebase/GA4 — mediation, demand, and measurement all under one parent.	n/a (mediation layer).	V3-S41

CHAPTER 10

The buyer's operating model

Control is only worth what the buyer can do with it. A log file nobody reads is a cost.

IN SHORT

Buyers choose between automated walled gardens, self-serve specialist platforms, managed service, in-house teams working through APIs, and agencies. The big platforms offer campaign APIs, choose placements for the buyer in their automated app campaigns, and run their lift tests through account teams. Some specialists document log-level data and holdouts. Transparency has two uses. Cutting clearly bad supply is largely a policy choice. Choosing between partners does, and it pays only with enough volume, clean data and the ability to run a holdout. In the US the demand is already there. Proof and signal decide adoption more than the sales story does.

What each model gives and takes

Walled gardens. Google, Meta, Apple and TikTok document campaign-management APIs, so a buyer can create and change campaigns in code, and Amazon documents one for its DSP ^[218,104,219,220,221]. Their automated app campaigns choose placements for the buyer. Google's placement report names the apps, sites and videos that showed ads, but gives impressions only, not results by placement ^[222]. TikTok's Smart+ app campaigns do not allow manual placement choice ^[142]. Meta's Advantage+ app campaigns limit targeting to operating system, country and language. They list placement among their reporting breakdowns, but the page describes no breakdown by publisher app ^[223] [CL-027](#). The trade is simple: reach and automation in exchange for less visibility.

Self-serve specialists. Kayzen documents a campaign-management API that edits bids and daily and total budgets ^[234]. Jampp documents a reporting API and log-level data for data-science teams, and advertises always-on lift measurement without documenting the method ^[69]. These are the tools transparency advocates ask for. They are also work: someone has to take in the logs, keep the event pipeline clean and read the tests.

Managed service. Smadex, owned by Entravision, describes its teams as acting "as an extension of our client's marketing team" and setting bids and fraud controls for clients ^[235]. Remerge offers managed service alongside programmatic buying ^[236]. Managed service moves the work to the vendor, and moves some visibility with it.

In-house and agencies. In a 2023 survey of 162 members of the Association of National Advertisers, 82% had an in-house agency, up from 58% in 2013 ^[237]. In the association's 2026 survey, only 9% named cost savings as the main benefit, down from 30% in 2023 ^[238]. The two surveys asked different groups, and the 2026 question was worded slightly differently, so the drop is a pointer, not a measured trend. In-housing is common, and it is not mainly about saving money.

When control pays for itself

Transparency does two separate jobs. The first is cutting waste: blocking apps with no supply-chain file, sources with telltale click-spam timing, or placements that never convert. Some of that is policy, such as refusing unauthorized sellers, and needs no experiment. Whether a given cut raises profit is still a question a test can answer. Chapter 8 shows how much unverified supply and fraud there is to cut.

The second job is choosing between partners or strategies. That is where control has real costs: people, tools and the time a new setup takes to learn. On performance, it pays back only if the buyer uses it to beat what the platform would do alone, which needs three things:

1. **Volume.** Enough spend that a few points of efficiency cover the team. At small budgets, one analyst can cost more than any likely gain.
2. **Clean data.** A server-side event pipeline, a defined contribution metric, and someone who can read log-level data.


3. **A test.** The ability to run holdouts. Without them, the buyer is choosing between two attribution reports, not two platforms.

Control also has uses no performance test captures: brand safety, children's privacy, procurement and audit. Where the three conditions hold, self-serve and API control can beat a managed or fully automated setup. Where they do not, a managed or automated platform may produce a better total result. We found no public study that measures the net value of transparency to buyers after operating costs [CL-028](#). The tool below shows why the answer usually turns on the performance assumption, which only a test can supply. Its defaults are illustrative, not survey data.

FIGURE 15

Operating cost turns on the performance assumption, which only a test can supply

ILLUSTRATIVEINTERACTIVE MODEL



An illustrative cost model for four buying set-ups. Fees are charged on top of media; a walled garden's take sits inside the media price, so it shows as zero here and its difference is carried by the performance assumption. Defaults are illustrative, not survey data. Managed-service and self-serve descriptions come from vendor filings and documentation [235,234,69].

Limits. Illustrative. Staff costs, fee rates and performance gains vary widely; test before deciding.

Package files: figures/F15.svg · data/figures/F15.csv

US commercialization: localization is rarely the only constraint

The author's playbook argues that app-growth vendors fail in the US less for lack of demand than for localizing the wrong things: buyer sequence, positioning, measurement story and proof [96]. The evidence supports the first half. The US holds about 42% of global

spend in the acquisition section of AppsFlyer's report, and US iOS acquisition spend grew 25% in 2025 ^[6]. Demand exists.

It qualifies the second half. For a challenger, the binding constraint is often one of these instead:

- **Proof.** The big platforms run their own lift tests through account teams. A challenger that cannot show causal results through a third party competes on attribution reports that favour incumbents with more signal.
- **Signal.** Privacy thresholds give richer postbacks to bigger campaigns (chapter 3). A new entrant with little spend per campaign learns more slowly.
- **Supply and data.** Owned auctions feed some leaders' models (chapter 2). A challenger buying on open exchanges starts with less.
- **Economics.** Managed service carries staff cost. Self-serve moves that cost to the buyer.
- **Rules.** Children's privacy rules and state app-store laws add compliance work, with dates still moving (chapter 11).

The playbook's "representative outcomes" are anonymized patterns with no metrics ^[96]. They are illustrations, not measured client results, and this paper does not use them as evidence.

Proof gates

A practical sequence for a buyer, or for a vendor selling to one, follows from the evidence. Each gate has a pass condition.

1. **Event readiness.** Server-side events for the value metric, deduplicated, with refunds.
2. **Attribution parity.** The new partner's results sit in the same attribution company and the same windows as incumbents.
3. **Holdout.** A randomized or geographic holdout on at least one campaign, with the design written down before launch and enough volume to detect the expected effect.
4. **Contribution.** Payback on net revenue after fees, with observed and forecast value shown apart.
5. **Scale test.** The second budget step, judged on marginal return, not average return, at the higher spend.

THESIS TEST · WHEN CONTROL JUSTIFIES ITS BURDEN (Q6)

THESIS	Transparency pays for itself in two ways. Cutting clearly bad supply is largely a policy choice. On performance, choosing between partners pays only when the buyer has the volume, clean data and testing ability to act on it.
SUPPORTING EVIDENCE	Supply gaps and fraud patterns that visibility can catch without an experiment (chapter 8); walled gardens trade placement visibility for automation ^[223,239] ; specialists document logs and holdouts ^[69] ; in-housing is driven more by capability than by cost ^[238] .
BEST COUNTEREVIDENCE	A strong automated platform can outperform a skilled in-house team, in which case control adds cost without gain.
OTHER EXPLANATIONS	Buyers may want control for risk reasons, such as brand safety or children's privacy, whatever the performance effect.
WHAT WOULD PROVE IT WRONG	Evidence that buyers without testing ability also gain net value from self-serve access after costs, or that buyers with it do not.

CHAPTER 11

The next app-growth system

Bidding is already automated. The next contest is over who sets the objective, who approves the changes and what the buyer can take away.

IN SHORT

"Apps DSP" is a useful checklist and a weak category. One product documents the full set of capabilities it implies, and five document most of them. But documentation is not proof of better results, and no one publishes independent evidence that the bundle wins. AI already runs the bidding. The new step is agents that plan and change campaigns. Google, Meta and TikTok have each published servers that let AI agents work with their ad systems. Google's is read-only. Meta's and TikTok's expose the same write actions as their campaign APIs. Meta's creates new campaigns paused and relies on the AI client to confirm switching them on. Neither documents an approval step for edits to live campaigns, so the buyer has to set limits in its own tools and contracts.

Testing the "Apps DSP" idea

The author's playbook proposes that challengers sell an "Apps DSP": a decision layer across acquisition, re-engagement, lifetime value, streaming TV and commerce, in place of a mobile DSP ^[96]. Treated as a hypothesis, it makes three claims. There is a set of capabilities that belong together. Buyers see the set as a category. And platforms that have it produce better outcomes.

The capability matrix in chapter 9 can test the first claim. Define the core as documented capability of at least 2 on value optimization, re-engagement, iOS privacy measurement, lift tests, reporting export, buyer control and streaming TV. This is our own screen, not an industry definition. Commerce is in the playbook's list, but the matrix has no commerce field, so the screen leaves it out. On that definition, 1 of the 11 execution products documents the full set: Moloco. AppLovin, Jampp, Kayzen, Remerge, Verve Dataseat document at least five of the seven. The most common gaps are iOS privacy measurement, lift tests and streaming TV. The threshold matters: at a score of 1 on every field, 3 products pass; at 3, 0 do. For the second claim, we found no buyer survey that uses the term in the sources searched by the cutoff. For the third, we found no public test. The verdict: the label describes a real checklist that one product documents in full. No one has shown that buying the bundle beats buying the best parts CL-029.

What would prove the category? Independent lift results, across at least two non-game categories, showing an integrated platform delivers more contribution per dollar than a mix of specialists measured the same way. What would disprove it? The same tests showing the bundle does no better once the buyer's own measurement is held constant.

What AI does now, and what agents add

Automated bidding is old news. Every major platform prices impressions with machine-learning models. AppLovin's disclosure describes its model scoring impressions from device, auction and advertiser data ^[29]. Mintegral says smart-bidding products drove more than 90% of its revenue in the first half of 2026 ^[240]. Google keeps widening automatic features in search campaigns ^[241]. Generative creative is spreading too, though some of it is only announced. A June 2026 Unity Ads newsletter, hosted outside Unity's own site, lists generative creative tools as coming with "agentic Vector" later in the year ^[242].

Agents are different. They plan, create and change campaigns through the platforms' interfaces, not just set bids inside them. Three of the largest platforms have published official servers for the Model Context Protocol, the common standard AI assistants use to call outside tools:

- **Google** announced an open-source Google Ads API server in October 2025. It is read-only by design: its documents say it cannot change bids, pause campaigns or create

assets ^[243,244].

- **Meta** runs a hosted ads server. Users sign in and approve the permissions it requests ^[245]. Its tools create and edit campaigns, ad sets, ads and catalogs, and set up tests ^[246]. New campaigns, ad sets and ads start paused, and Meta says the AI client asks for confirmation before switching them on. The page describes no such step for edits to live campaigns, such as budget changes ^[247]. It also offers reporting and activity-log tools, and its overview page gives no launch date ^[246].
- **TikTok** exposes about 400 advertising functions, including campaign creation and budget changes. Its help page says agents can run whole workflows "without human intervention at every step" and describes no approval step ^[248] CL-030.

These servers wrap campaign APIs that have long let software make changes without a person approving each one. The new part is that a general-purpose AI assistant can now use them. Account permissions and budgets still apply. But a per-change approval limit usually sits in the agent the buyer runs, or in the contract, not on the platform's server. So buyers should ask three plain questions: is this read or write access, who approves changes above a set limit, and where is the log?

Standards for agents

Two families of standards aim to set how buying agents work. The Ad Context Protocol, governed by [AgenticAdvertising.org](https://agenticadvertising.org), reports its own first agent-to-agent media buy with real money in October 2025. Its latest stable release at the cutoff was 3.1.24, published on 23 September 2026 ^[249,250,251]. Its governance specification lets a buyer set approval thresholds, such as a budget cap above which a person must approve. It also requires an audit log entry for each governance check ^[252,253]. It also splits duties: the agent that runs buys cannot set its own spending limits or approve its own plans ^[254]. But this part of the protocol is experimental and opt-in. Its rules bind only when a plan has a governance agent, and approval-flow conformance tests apply only to sellers that declare governance support. A governance agent can also run in an audit mode that approves everything and only records findings ^[253,255]. IAB Tech Lab's Agentic Advertising Management Protocols (version 3.0; page last updated 22 September 2026) is an umbrella. It covers direct-deal workflows, a component for automated campaign planning, and the Agentic RTB Framework, which works inside real-time bidding systems ^[256,257,258]. Its 2.0 release in April 2026 added buyer and seller agent kits with configurable human approval gates ^[259]. In the pages we read, neither family refers to the other, and none of the app platforms we profiled documents support for either. The author co-leads a working group within the Ad Context Protocol; see the disclosures.

The rules tightening around app data

Whatever the agents do, the data they work with is getting harder to use. Four sets of rules moved in the year to the cutoff.

Children. The FTC's amended children's privacy rule took effect in June 2025, with a general compliance date of 22 April 2026 ^[260]. Apple's guidelines restrict third-party analytics and ads in apps aimed at children. Google Play requires child-directed apps to use only certified ad kits and bans interest-based ads and remarketing to children ^[59,261].

Age checks in the stores. Texas's app-store age law took effect after the Fifth Circuit paused the lower court's injunctions. It did so briefly on 28 May 2026 and for the whole appeal on 4 June. The appeal was argued on 4 August, and the docket showed no ruling on the merits by the cutoff ^[102,262]. Utah delayed its duties to 6 May 2027 and dropped state enforcement; only harmed minors or their parents can sue ^[263]. Louisiana stopped its 2025 law from taking effect, and its replacement duties start on 1 July 2027 ^[264]. California's Digital Age Assurance Act, signed in October 2025, takes effect on 1 January 2027 ^[265]. Apple's Declared Age Range and Google's Play Age Signals pass age bands to apps ^[266,267]. Google's API is still in beta. At the cutoff it returned signals only in Brazil and for eligible Texas accounts, and its terms bar use for advertising, marketing, profiling or analytics ^[267]. So some users now arrive with an age band, in a few places, under tight limits.

Location data. The FTC has finalized orders against two location-data brokers, X-Mode and InMarket ^[268,269]. It also settled with Kochava, an attribution and data company. The

court entered that order on 25 June 2026, after the FTC filed it in May ^[270,271,272]. It bars Kochava from selling or sharing sensitive location data, such as visits to medical or religious sites. The one exception is narrow: the consumer must deal with Kochava directly, give express consent and ask for the service. Kochava neither admitted nor denied the charges ^[271].

Consent on iOS in Europe. Chapter 3 covered the competition rulings and Apple's beta consent sheet.

None of these rules stops app advertising. Together they make the data behind targeting, attribution and model training more conditional. That raises the value of first-party events an app collects with clear consent.

What a buyer can take away

An app marketer's spend trains the platform's model, and the model then serves every advertiser. A pooled model cannot hand its weights back to one client. The real questions are about data use and access: what the platform may do with the buyer's data, and what the buyer can export. AppLovin's legal disclosure on its model lists the data it uses and what it excludes, such as other bidders' data. It says nothing about what rights advertisers have over the use of their campaign data ^[29] [CL-031](#). We did not review the full advertising terms of Meta, Google or TikTok, so we make no claim about them. In the EU, the Digital Markets Act gives some rights by law. Designated gatekeepers must give advertisers free access to their performance-measuring tools and the data needed to check ads independently. They must also give business users access to the data their use generates ^[273]. Those rights apply in the EU, not the US. A buyer who switches partners keeps its own event history, creative and test results. The pages we reviewed did not show any right to take the platform's learned state. That is a switching cost no invoice shows.

Three questions follow for any buyer adopting agents or automated buying:

1. **Who writes the objective?** The agent should optimize the buyer's contribution metric, not a platform default such as installs or attributed return.
2. **Who approves?** Changes above set limits should need a person, and each one should be logged where the buyer can see it.
3. **What leaves with you?** Event data, test results and the audit log should be exportable. The buyer's own test history is the learning that can move.

FIGURE 16
Three futures for app growth, with the signals that would tell them apart

SCENARIOFORECAST

THE STACK KEEPS WINNING Platforms that own an auction, a bidding model and a measurement feed keep taking share, and extend into e-commerce and streaming TV. ASSUMES - Owned feedback keeps compounding model quality - No regulator separates mediation or measurement from demand - Buyers keep accepting platform-run measurement WATCH - AppLovin discloses e-commerce or web-advertiser revenue - More specialists consolidate into portfolios - Platform lift tests stay gated behind account teams CASE AGAINST Unity owns an auction too and earns a fraction of AppLovin's margin ^[11], so model quality and execution may matter more than ownership. RULED OUT IF Independent tests show specialists matching the stacks' incremental value per dollar, or a regulator restricts data use across mediation and demand.	A MEASURED MARKET Buyers make holdout tests a condition of spend; causal results, not attributed ones, set budgets. ASSUMES - Lift tests become self-serve - Buyers write test designs into contracts - EU consent changes lift iOS opt-in and postback detail WATCH - Self-serve lift products at the big platforms - MRC accreditation for attribution companies' fraud detection - Buyer requests for proposals that require holdouts CASE AGAINST Informative experiments need very large samples ^[115], so most advertisers cannot run them often. RULED OUT IF Lift products stay gated and holdout use stays rare through 2027.	THE WEB PULLS VALUE OUT OF THE STORE Link-out rules and new store fees make web checkout normal for subscription and commerce apps in the US, and web-to-app paths become the best-measured route. ASSUMES - Courts leave any US link-out fee well below store rates - Web checkout conversion is close enough to in-app purchase - Deterministic web measurement is preferred by finance teams WATCH - The trial court's link-out fee ruling and the Supreme Court's contempt ruling in Apple v. Epic ^[99] - Google Play fee reporting from 1 October 2026 ^[98] - Web share of subscription revenue in North America rising above RevenueCat's 4.9% ^[188] CASE AGAINST Web checkout adds steps and may lose buyers; Apple may win the right to charge a substantial fee. RULED OUT IF Courts allow a link-out commission near the store rate, or web share stays flat through 2027.
Across all three: agents move the control point. Whichever future arrives, AI agents will plan and change more campaigns. The buyer who writes the objective, sets approval limits and keeps the test history keeps control. Watch for approval gates on write-capable agent servers ^[246,248] and for contracts that let buyers export their learning.		
WHO	OPPORTUNITY THE EVIDENCE SUPPORTS	
App advertisers	Fund a standing holdout budget; move payback to contribution; test web checkout where the law now allows it.	
Challenger platforms	Compete on third-party-verified incrementality and exportable data, not on the 'Apps DSP' label.	
Measurement companies	Earn independent accreditation; publish methods and error for predicted LTV.	
Investors and finance	Read ad-platform margins with revenue recognition in view; ask for cohort contribution, not ROAS.	

Scenarios for 2027-2028 written at the cutoff. Each lists its assumptions, the indicators to watch, the case against it and the result that would rule it out. They are judgments built on this paper's evidence, not forecasts with probabilities.

Limits. Scenario judgments by the author; no probabilities are assigned.

Package files: data/figures/F16.csv

THESIS TEST · IS "APPS DSP" A CATEGORY? (Q11)

THESIS	"Apps DSP" is a useful buyer checklist but not yet a proven category; its value rests on untested claims of better outcomes.
SUPPORTING EVIDENCE	On our screen, one of eleven products documents the full set of capabilities (chapter 9); we found no buyer survey that uses the term and no independent outcome comparison in the sources searched by the cutoff.
BEST COUNTEREVIDENCE	The most integrated specialist also has the strongest financial results ^[9] , which fits bundles winning.
OTHER EXPLANATIONS	Those results may come from owned supply, model quality and scale rather than a capability bundle a challenger can copy.
WHAT WOULD PROVE IT WRONG	Independent tests across categories showing integrated platforms deliver more contribution per dollar than best-of-breed specialists measured the same way.

THESIS TEST · WHAT AGENTS CHANGE (Q12)

THESIS	AI agents extend automation from bidding into planning and execution. Control of objectives, data use and accumulated learning stays with the platforms unless buyers set it in their own tools and contracts.
SUPPORTING EVIDENCE	Write-capable agent servers with no documented approval step for edits to live campaigns ^[247,248] ; a read-only alternative ^[243] ; silence on data-use rights in AppLovin's model disclosure ^[29] .
BEST COUNTEREVIDENCE	Open protocols define approval thresholds and audit logs ^[252] , and Meta's server is permission-scoped, creates new campaigns paused and keeps activity logs ^[246,247] .
OTHER EXPLANATIONS	Platforms may offer approval controls in product settings or contracts that their public help pages do not describe.
WHAT WOULD PROVE IT WRONG	Standard contract terms or platform settings that give buyers approval control, audit logs and rights over their data's use and export.

CHAPTER 12

Method, limits and corrections

This chapter says what was done, what it caught and what to distrust. It reports only checks that actually ran.

IN SHORT

The author ran the research with AI agents. Ten evidence streams and three gap-fill passes each opened their sources and recorded findings with an evidence class and confidence. Five fact-check passes re-opened the sources and tried to refute the findings. Two passes scored the vendor panel separately. A hostile reviewer attacked the argument, and an auditor rechecked the draft's numbers and citations. ChatGPT (OpenAI) reviewed three builds, and each of its factual points was checked against the sources. Every change is logged below. The main prose scores 57.6 on the Flesch Reading Ease scale. The original target of 85 was not met; the author accepts 57.6 as a disclosed exception for this technical paper. A separate plain-language edition carries the target, and its score is measured on its own page. None of this amounts to peer review or an external audit.

How it was built

The work ran on 27 September 2026 against a cutoff of the same day. The author set the questions, scope, definitions and scoring rules in a written brief before any research ran. Research agents then worked in separate streams. They covered market and platform economics, Apple and Android measurement, and attribution and causal evidence. Others took acquisition, creative and discovery; retention and unit economics; fraud and streaming TV; and buyer operations, AI and regulation. Three more covered vendor groups. A separate pass mapped the author's own writing. Three gap-fill passes followed up on questions the first round could not answer or the reviewers raised; the last of them answered the hostile reviewer's evidence requests.

Every stream worked from the same source rules. Filings and official documents came first; vendor material counted only as evidence of what the vendor says. Agents were required to open a source to cite it. Where only a summary, an abstract or a search result could be read, the source register says so, and confidence was lowered. The research consulted 621 distinct source addresses. That counts web addresses, not works: a study's abstract page and its PDF count twice. The research agents recorded 501 of these addresses as read in full. Records whose own notes say only an abstract, a navigation or index page, or headline figures were read are not counted as full. A full page is not always the full underlying report or docket. Part-way through, the session's web-search allowance ran out. Agents then worked from known official addresses, filing systems and academic indexes, and some gaps in the open-questions lists come from that limit. Some agents used a public read-only page renderer to read pages that blocked direct fetches, and the lead editor opened a few pages in a browser.

Verification ran as separate passes with a brief to refute. Five passes re-opened the sources behind the recorded findings and gave each one outcome: supported, qualified, contradicted, unresolved or withdrawn. At the time of this build the outcomes were 120 supported, 136 qualified, 10 contradicted, 2 unresolved and 4 withdrawn. Findings left unchecked: 0. Where a check found the numbers right but a label wrong, the editor regraded the finding as qualified and logged why. The paper uses corrected wording wherever a finding was qualified. It never uses the original claim of a contradicted or withdrawn finding; where a check supplied a corrected statement, the paper uses that instead.

A second scorer rated the eleven execution products and five benchmark channels without seeing the first scores. A hostile reviewer read the whole draft looking for other explanations, selection effects and places where matching was presented as cause. It returned 30 issues, five of them blocking. An auditor checked the draft's numbers, formulas and citations and returned 48 issues, four blocking. All blocking issues were addressed; the thesis itself was narrowed as a result. Changes from all passes are in the corrections ledger.

ChatGPT (OpenAI) assisted with technical and editorial review and source verification. The author retains responsibility for the paper. This assistance does not constitute independent assurance or endorsement by OpenAI. The author ran three review rounds with it. The first review returned about 100 numbered points, plus notes on each headline claim and figure. Five further passes checked its 74 factual points, plus two extra checks, against primary sources. They found 36 right and 17 partly right. Another 19 had already been dealt with, by earlier passes or by edits in the same round; the audited file was not kept, so those two cases cannot be told apart for every item. One point could not be resolved, and one had flagged a source it could not reach, which the docket then confirmed. The checks also found errors the review had missed, such as a study whose peer-reviewed version changed its figures. The second review returned 25 findings; the change log records what was done about each. The third round was a close-out: it confirmed the corrections and asked for four final fixes, which were made. Each check is kept with the evidence files.

These are passes by AI agents working under one person's direction. They reduce error; they do not make the paper independent. The paper has not been peer reviewed or externally audited, and it reports no inter-rater reliability statistic.

Sampling and scoring

The vendor universe started from the playbook's list and grew through filings, partner pages and documents. It is a sample of what was discoverable, tilted toward firms that publish and firms that survived. Scores measure documented capability against fixed anchors; they say nothing about performance, and "n/e" marks missing evidence, not absence. The one composite, the App Quotient, is reported in bands, with the number of scored fields shown for each group. It does not count unscored cells as zeros, pulls thin records toward the field mean by an author-chosen amount, and is shown under four reweightings. Its order is not a ranking, and it still rewards disclosure.

Readability

Main prose was scored with the Flesch Reading Ease formula: 206.835 minus 1.015 times words per sentence, minus 84.6 times syllables per word, using the script in the package. The score covers chapter paragraphs, lists, summaries, callouts and thesis tests. It excludes tables, figure captions, instruments, references and URLs. Syllables are counted with a documented vowel-group rule, with acronyms counted as spoken. The result is **57.6** over 19,791 words, at 13.38 words per sentence and 1.603 syllables per word. That is the score for the measured prose, not for every word in the PDF: tables, profiles, captions and references are left out. The target was 85. Sentences were shortened and plain words used wherever meaning allowed. Terms such as attribution, incrementality and subscription carry many syllables and could not be removed without changing meaning, so the target was not met. The author accepts the measured score as a disclosed exception for this technical paper. The target of 85 moved to a separate plain-language edition, written afresh in shorter words rather than cut down from this text. It is scored with the same script, and its build fails below 85. Chapter scores are in the package.

What to distrust

- **Transfer from web to apps.** The strongest causal studies measure web, search, retail and audio outcomes. Their lessons about attribution and ad load probably hold for apps, but that is an inference.
- **Vendor panels.** Spend, retention and subscription figures come from vendors' own client bases. They describe those clients, not the market.
- **Global figures read as US figures.** The spending shifts come from a global panel, and the fastest of them was driven by China-based e-commerce budgets. US-only category data were not available.
- **Absence claims.** Several findings say no public evidence was found. Those depend on what the research could reach after the search allowance ran out, and private studies may exist. Read each as "we did not find", not "none exists".
- **Documentation scores.** A vendor that documents less gets more n/e cells, and can look weaker than it is.

- **Fast-moving rules.** Court cases, store fees and consent screens were moving at the cutoff. Check dates before acting.
- **The author's position.** The author sells advisory services to firms in this market and works for a TV measurement company. The evidence rules were written to limit that bias; readers should still weigh it.

Corrections

Corrections ledger						
What the verification, scoring and argument passes changed, and why.						
ID	ORIGINAL	CHALLENGE	VERDICT	CHANGE MADE	AFFECTS	SOURCES
COR-01	Observational methods 'missed' randomized lift by a median 24% to 176% (read as error rates).	Fact-check pass FX3 re-read Gordon, Moakler & Zettelmeyer (2023).	contradicted	Reworded: the 83/58/24% and 173/176/64% figures are median estimated lifts, against true median lifts of 29/18/5%. Abstract, chapter 4 text, thesis test and claim CL-012 corrected.	Abstract; ch. 4; CL-012	D-S02
COR-02	An Android user who resets the advertising ID hands apps a string of zeros.	FX2 re-read Google Play's Advertising ID policy.	contradicted	Zeros are returned when the user deletes the ID; reset issues a new ID. Chapter 3, OS matrix and Figure 4 corrected.	ch. 3; OS matrix; Fig. 4	C-S18
COR-03	Google's data-driven attribution needs at least 200 conversions and 2,000 interactions in 30 days.	FX3 re-read Google Ads Help.	contradicted	Now stated as a recommendation; the model is available at any volume.	ch. 5	D-S10
COR-04	Verve Group adjusted EBITDA margin 22.3% of revenue.	FX1 re-read Verve's Q4 2025 release and annual report.	contradicted	22.3% is the margin on like-for-like gross revenue of EUR 601.8m; on reported revenue it is 24.3%. Table corrected and the Q3 2025 net-to-gross change noted.	ch. 2 table; CL-006	AB-S09
COR-05	Top-quartile iOS games keep 31-33% on day 1 against 25-27% on Android (GameAnalytics).	FX3 read the full GameAnalytics 2026 report.	withdrawn	Removed from chapter 7 and Figure 9; the report combines iOS and Android.	ch. 7; Fig. 9	GH-S11
COR-06	Meta's app-event reports use impression time and a 28-day click / 1-day view model.	FX3 re-read Meta's App Events API page.	qualified	Clock and 28/1 windows attributed to Ads Manager; the API endpoint uses 30-day clicks with optional 1-day views.	ch. 4	D-S16
COR-07	Meta files engaged-view conversions as click-through in its aggregate reporting.	FX3	qualified	Attributed to Singular's aggregate reporting of Meta campaigns.	ch. 4	D-S24
COR-08	TikTok click windows of 1, 7 or 28 days.	FX3	qualified	Added the 14-day option.	ch. 4	D-S22
COR-09	A reinstall inside the reattribution window generates no install postback and later events are organic.	FX3	qualified	Added the retargeting-reinstall exception.	ch. 4	D-S14
COR-10	Aridor et al. used Kantar Vivvix, Shopify and SimilarWeb data; no platform funding.	FX3 read the published paper.	qualified	Core data are an ad-analytics panel and a revenue panel from a firm linked to a co-author; the other sources are benchmarks. Chapter 4 and Figure 6 placement updated.	ch. 4; Fig. 6	D-S05
COR-11	Removing off-Meta data raised cost per incremental customer from \$43.88 to \$60.19.	FX3	qualified	Stated as an estimate under a median loss of effectiveness; two of four authors are Meta employees; working paper.	ch. 4	D-S06
COR-12	Survey experiment with 11,000 adults; two Meta co-authors.	FX3	qualified	Described as stated choices with hypothetical prompts; a third author has consulted for Meta.	ch. 3	D-S07
COR-13	Apple Ads covers App Store search download campaigns; view-through window inferred as 1 day.	FX2 and FX3	qualified	Windows now stated from Apple's reference (30-day tap, 24-hour view); placements beyond search added.	ch. 3; OS matrix	C-S10 D-S20
COR-14	No AdAttributionKit change after June 2025.	FX2 symbol crawl	qualified	Noted an iOS 26.2 view-through registration API missing from the changelog.	ch. 3	C-S06
COR-15	Google Play new model: 10% + 5% for new-install subscriptions; 20% + 5% one-off.	FX3	qualified	Corrected: 10% subscriptions; 20% new-install and 25% existing-install one-off; programs 15%/20%; 10% on first \$1M; 5% billing fee only on Google Play Billing in US/UK/EEA.	ch. 7 table; Fig. 4 label	GH-S06
COR-16	Apple EU alternative terms effective 1 October 2026 (presented as current).	FX3	qualified	Marked as announced in August 2026 and effective after the cutoff.	ch. 7 table	GH-S03
COR-17	Apple barred from any US link-out commission until a court approves a fee.	FX3 read the Ninth Circuit opinion and docket.	qualified	The 'no commission' is the appeals court's recommendation; Apple filed a fee proposal on 13 August 2026; no fee approved by late September.	ch. 7 table	EF-S06 EF-S17
COR-18	Singular says its SKAN-era predictive models perform comparable to IDFA-era measurement.	FX3	qualified	Dated as a 2021 post from the SKAN 3 era relaying customer claims.	ch. 7	GH-S19
COR-19	AppLovin owns Axon, MAX, Adjust and Wurl.	FX1 read the Q2 2026 10-Q.	qualified	Added the rename to AppLovin Ads, acquisition years, and AppLovin's stated assurance that Adjust data is not shared with it unless the customer directs.	ch. 2	AB-S01 G2-S01
COR-20	Gaming was slightly under half of Liftoff's advertiser revenue.	FX1	qualified	Restated as the filing words it: slightly more than half from outside gaming, excluding third-party programmatic buyers.	ch. 2	AB-S10
COR-21	PyMC-Marketing claims superior results versus Meridian.	FX3	qualified	Quoted as a vendor self-benchmark (faster fitting, 40% lower contribution error).	ch. 4	D-S19
COR-22	Author's playbook: Google 'retired' the Privacy Sandbox advertising APIs on Chrome and Android in October 2025.	Corpus audit and stream C against Google's status page.	qualified	The paper states the retirement was announced and listed as 'scheduled for phaseout' at the cutoff, not completed.	ch. 3; corpus map CP-01	C-S14 C-S15 CP-01
COR-23	Author's App DSP landscape file (last validated June 2026) describes AppLovin's product as Axon Ads Manager and carries mid-2026 facts.	Corpus audit and G2 filings check.	qualified	Marked 'update' in the corpus map; the paper uses Q2 2026 filings and the AppLovin Ads name.	corpus map CP-02	G2-S01 CP-02
COR-24	Adapty report compares 10,000+ paywalls; period implied as 2025.	FX3	qualified	Page says 105K paywalls and states no period; chapter 7 notes the missing period.	ch. 7	GH-S09
COR-25	AppsFlyer counted \$7.2B of verified in-app ad revenue.	FX3	qualified	Only store purchases and subscriptions are labelled verified; wording corrected.	ch. 7	GH-S10

COR-26	Supreme Court took Apple's appeal on the link-out commission; fee 'before the Supreme Court'.	FX2 read the question presented and docketed.	contradicted	Cert is limited to whether contempt may rest on an injunction's 'spirit'; the fee is set on remand, a stay was denied on 13 Aug 2026. Abstract, ch. 6, CL-021, scenario card and Fig. 4 label corrected.	Abstract; ch. 6; CL-021; Fig. 4; Fig. 16	FX2-S21 EF-S17
COR-27	Google released its Google Ads API MCP server in April 2026.	FX2	qualified	Dated to the October 2025 announcement; Meta and TikTok launch dates removed because their pages give none.	ch. 11	FX2-S09 KL-S06
COR-28	AAMP, ARTF and AdCP are three separate efforts that do not refer to each other.	FX2 read IAB Tech Lab's AAMP repository.	contradicted	AAMP is IAB Tech Lab's umbrella and lists ARTF as a component; ch. 11 now describes two families of standards.	ch. 11	FX2-S11 KL-S35
COR-29	Louisiana's app-store age law effective date unresolved (1 July 2026 vs 2027).	FX2 read Act 185.	contradicted	Act 185 repealed the 2025 law and sets new duties from 1 July 2027.	ch. 11; policy register	FX2-S18
COR-30	FTC finalized a Kochava order in May 2026.	FX2	qualified	Described as a proposed stipulated order; court entry not confirmed.	ch. 11	KL-S19
COR-31	Meta documents that Advantage+ app advertisers cannot see results by placement.	FX1 re-read Meta's help page in a browser.	contradicted	Meta's page now lists placement as a reporting breakdown; the paper states Meta reports by placement surface but not by publisher app. CL-027, ch. 9 and ch. 10 corrected; Meta's D5 score revisited in reconciliation.	ch. 9; ch. 10; CL-027; matrix	FX1-S23
COR-32	All five benchmark channels document campaign-management APIs.	FX1	qualified	TikTok's and Amazon's campaign-management APIs were not confirmed from primary docs; text now names Google, Meta and Apple.	ch. 9; ch. 10	FX1-S24 V3-S12
COR-33	Statsig was acquired by Amplitude on 5 May 2026.	FX1 read Amplitude's Q2 2026 10-Q.	contradicted	Amplitude bought Statsig's customer contracts, trade name and a non-exclusive technology licence, not the company.	ch. 2; ch. 9	FX1-S16
COR-34	Movvista's revenue basis not stated.	FX1 read the 2025 Annual Report.	qualified	Movvista books gross as principal, gross margin 21.2%; ch. 2 and CL-006 updated.	ch. 2; CL-006	FX1-S11
COR-35	Affle runs at least five app-marketing brands including Appnext.	FX1	qualified	Appnext link unconfirmed; now 'at least four' plus the June 2026 AdColony asset deal.	ch. 2	FX1-S18 FX1-S20
COR-36	Aarki rebranded to RZR on 17 March 2026.	FX1	qualified	Date not shown by the company; notice live by 12 March 2026.	ch. 2	V2-S18 FX1-S22
COR-37	Jampp documents an always-on ghost-bid holdout for every campaign, modelled on ghost ads.	FX1	qualified	Jampp advertises always-on lift with ghost bids for UA and retargeting at no extra cost; the method is not documented.	ch. 4; ch. 10	V2-S22
COR-38	Tripledot consideration: \$400M cash plus ~20% equity.	FX1 read the Q2 2026 10-Q.	qualified	Now \$430.6M cash after adjustments plus shares valued at \$285.0M (\$715.6M total).	ch. 2	FX1-S01
COR-39	Apple's iOS 27.2 expanded ATT sheet (beta).	FX2 read iOS 27.2 beta 2 release notes.	supported	Added that the same notes introduce a yearly EU re-prompt; both beta at cutoff.	ch. 3	FX2-S01
COR-40	Amazon DSP credits app installs to CTV exposure via logged-in identity on both screens.	FX3 re-read the cited Amazon page.	withdrawn	Removed from chapter 8, Figure 13 and the Q9 thesis test; the page covers Android and Fire devices only and says nothing about logged-in CTV-to-install.	ch. 8; Fig. 13	IJ-S26
COR-41	AppsFlyer links CTV exposure to installs by IP address and user agent.	FX3	qualified	Now 'probabilistic modelling on IP address alone'.	ch. 8; Fig. 13	IJ-S25
COR-42	Kochava matches deterministically only with device ID, user agent and IP together.	FX3	qualified	Removed; the cited post does not describe matching rules. Lift is by post-hoc synthetic controls.	ch. 8; Fig. 13	IJ-S10
COR-43	MNTN credits website or app visits, via deterministic matching.	FX3	qualified	Website visits only; matching method not disclosed.	ch. 8; Fig. 13	IJ-S12
COR-44	Roku integrates AppsFlyer, Adjust and Branch and offers Action Ads.	FX3	qualified	Now described as a pixel and server-to-server conversion feed; MMP names and Action Ads removed.	ch. 8; Fig. 13	IJ-S11
COR-45	Uber and Fetch settled in 2019; \$82.5M campaign.	FX3 read court-related sources.	qualified	The federal suit against Fetch was dismissed in Dec 2017; the \$6M settlement was in the separate Phunware state case after pleadings were struck; an Uber-Fetch settlement is not confirmed by a primary source.	ch. 8; CL-025	IJ-S04 IJ-S05 IJ-S06
COR-46	app-ads.txt on 66% of the top 1,000 Play apps and 24% of all apps.	FX3	qualified	Those two figures are not in the source; kept only the 25% of app bid requests without a known file (vs 3% on web) and the seller-registry gap.	ch. 8	IJ-S21
COR-47	AppLovin's e-commerce launch required \$10M GMV; CEO stated a ~\$7B target.	FX3	qualified	The \$10M bar applied to the earlier invite-only phase; the \$7B was a hypothetical sizing exercise, not a target.	ch. 8	IJ-S22 IJ-S23
COR-48	Store-page conversion benchmarks of 25-33%, Play 3-5 points above iOS.	FX3	contradicted	Removed; replaced with Apple's own 1.6% default-page average and its 2.5-point custom-page figure, used to illustrate traffic mixing.	ch. 6	MAIN-S04
COR-49	Unverifiable statistics (a '45% CTR drop after the fourth repetition'; web-to-app '2.8x' conversion).	FX3	withdrawn	No longer quoted; the text says such statistics could not be traced.	ch. 6	EF-S19
COR-50	Target-ROAS thresholds cited to the App campaign bid-strategy page; Meta learning phase rule unverified.	FX3 and the lead editor	qualified	Re-sourced to the correct Google page (MAIN-S03) and Meta's own page (MAIN-S05), which confirms ~50 results a week and advises combining ad sets.	ch. 5; CL-015	MAIN-S03 MAIN-S05
COR-51	A vendor analysis shows over-concentration wastes spend (10% overlap at ~2x frequency).	FX3	qualified	Described as a stylised example from a vendor selling de-duplication, not evidence for or against diversification.	ch. 5	EF-S23
COR-52	Meta Dynamic Creative described from third-party sources only.	FX3 read Meta's pages via a text proxy	qualified	Now states Meta's own guidance: A/B tests split audiences; Dynamic Creative should not substitute for split tests.	ch. 6	EF-S18
COR-53	Central thesis: 'the market optimizes what it can see early and claim; the edge belongs to whoever owns the feedback loop'.	Red-team RT-01 and RT-02 (blocking)	qualified	Narrowed: engines are built to learn from what they see, but 2025 money moved toward thinner-signal segments (iOS +35%, Android -1%; non-gaming +18%, games +3%). 'Owned feedback' restated as a hypothesis with Liftoff and Unity as counterevidence. Abstract, hero, ch. 2, ch. 5 and Q1/Q5 tests rewritten.	Abstract; hero; ch. 2; ch. 5	MAIN-S01 FX1-S05
COR-54	Platform table showed Liftoff's GAAP net loss while peers showed adjusted EBITDA; Unity's 22% used as an ad-business margin.	Red-team RT-03 (blocking); audit MS-09, MS-28, MS-30	qualified	Table now shows each firm's revenue basis and adjusted EBITDA on a like basis, with Unity's margin flagged as including its engine business; rounding fixed.	ch. 2 table	FX1-S05 AB-S05 FX1-S01
COR-55	Remarketing headline: attribution overstates value where intent exists (supported by Gordon et al. lower-funnel lift and eBay brand search).	Red-team RT-04 (blocking); RT-14	qualified	Now: randomized web retargeting studies find real but modest lift; no independent app-remarketing experiment was found; brand-search effects depend on competition for the term. CL-005 and CL-018 reworded.	Abstract; ch. 2; ch. 5; CL-005	D-S04

COR-56	Verification described as covering every finding; blind rescore described as covering the whole panel.	Red-team RT-05; audit MS-03 (blocking)	qualified	A fourth fact-check pass (FX4) covered the gap-fill and V1 findings; the rescorer was resumed for the benchmark channels; ch. 12 reports actual coverage and agreement statistics computed at build time.	Abstract; ch. 12; ch. 9	
COR-57	'57% of marketers use technical AI agents; 32% optimization agents' (AppsFlyer).	Audit MS-01 (blocking)	withdrawn	Sentences removed from ch. 11.	ch. 11	MAIN-S01
COR-58	Reconciliation instrument rows (Google DDA 'requires' 200 conversions; Meta, TikTok and AdServices windows).	Audit MS-02 (blocking)	contradicted	Rows overridden from corrected findings via data/register_overrides.json.	ch. 4 instrument	D-S10 D-S16 D-S22 D-S20
COR-59	Chapter 9 and 10 summaries: the big platforms offer 'no placement reporting'.	Audit MS-04, MS-05, MS-06	contradicted	Summaries now say the platforms choose placements; Meta lists placement among breakdowns but describes no publisher-app breakdown. CL-027 relinked to MAIN-F09.	ch. 9; ch. 10; CL-027	FX1-S23
COR-60	Contribution defined as net revenue minus the cost of getting the customer (ch. 1).	Audit MS-07	qualified	Contribution now excludes acquisition cost; growth = contribution minus acquisition cost.	ch. 1	
COR-61	Figure 10: 129% gross return and 23-month payback at a flat 30% fee.	Audit MS-17; MS-45	qualified	Model now applies 15% after month 12; payback at defaults is month 21; 'CAC' relabelled 'fully loaded cost per install'.	Fig. 10; models.py; paper.js	GH-S02
COR-62	Chen et al. game LTV error '5.7% mean error'.	Audit MS-15	qualified	Stated as error relative to the largest observed spend, the study's own normalization.	ch. 7	G1-S04
COR-63	SHORE preprint cited as support.	Audit MS-16	qualified	Noted as withdrawn by its authors.	ch. 7	G1-S08
COR-64	Web checkout 'about a third more' and 'more than most media optimizations'.	Red-team RT-24; audit MS-32, MS-33	qualified	Now +34% vs the 30% tier and +11% vs the 15% tier; card fees only; hidden billing costs listed; unsupported comparison with media tweaks removed.	Abstract; ch. 7; CL-022	G2-S15 GH-S02
COR-65	RevenueCat 72% 'cancel within the first year'; AI-app LTV used as evidence about bidding models.	Red-team RT-08, RT-09	qualified	Now 'turned off auto-renew'; cohort window not stated; the AI-app figure described as a category comparison.	ch. 5; ch. 7	GH-S08
COR-66	Paid media has 'no direct lever' on activation or retention.	Red-team RT-26	qualified	Now 'acts only indirectly', with reactivation as the exception.	ch. 1; CL-001	
COR-67	Attribution error described as running one way.	Red-team RT-07	qualified	Added sources of under-attribution on iOS; CL-011 reworded to 'can produce more installs than those platforms drove'.	ch. 4; CL-011	C-S02 D-S13
COR-68	Agent servers described as lacking approval gates; 'no public contract gives buyers rights to what models learn'.	Red-team RT-16, RT-17, RT-18	qualified	Now: servers wrap existing campaign APIs; approval usually sits in the buyer's agent and contract; Meta's server is permission-scoped with activity logs; learning claim scoped to AppLovin's disclosure and reframed as data-use and export rights; AdCP enforcement depends on seller implementation.	Abstract; ch. 11; CL-031	KL-S07 KL-S08 KL-S12 KL-S03
COR-69	Verve profile 'held to the same or a stricter standard'.	Red-team RT-20	qualified	Now 'exactly the same rules'; n/e cells explained by missing documents; both passes scored Verve.	ch. 9	
COR-70	Jampp lift-test score 3 (standardized method).	Red-team RT-21; audit MS-23; FX1 V2-F05	qualified	Set to 2 by reconciliation override: method not documented.	matrix	V2-S22
COR-71	Incremental ROAS 'undefined when incremental outcomes are zero or negative'.	Audit MS-19	contradicted	iROAS is undefined only when incremental spend is zero.	metric dictionary	
COR-72	Policy register rows for Louisiana, Utah, Declared Age Range and Kochava; study-table sponsor notes; Statsig ownership row.	Audit MS-24, MS-47	qualified	Rows overridden from corrected findings via data/register_overrides.json.	Appendix instruments	FX2-S18 FX2-S15 FX2-S19 FX1-S16
COR-73	Google's app campaign documents describe no placement-level reporting.	External AI audit item 1; re-check OA1	contradicted	Google's placement report names apps, publishers, websites and videos, with impressions only and low-volume placements grouped. Google D5 re-scored from n/e to 2 (grade A).	ch. 9, 10; CL-027; benchmark matrix	OA1-S01
COR-74	We found no record that the court had entered the Kochava order by the cutoff.	External AI audit item 2; re-check OA2	contradicted	Court entered the order on 25 June 2026; its narrow exception and no-admission clause described.	ch. 11; policy register	OA2-S01 OA2-S02
COR-75	SHORE preprint noted as withdrawn by its authors, and credited with naming late payments, sparse data and extreme spenders via another preprint's citation.	External AI audit item 3; FX4; re-check OA3	qualified	SHORE described as removed by arXiv's administrators over licence rights and not relied on; the other preprint (Tencent) now reported with its own figures, including Google's method doing worst.	ch. 7	G1-S08 G1-S07 OA3-S06
COR-76	Meta's Ads Manager uses 28 days after a click and 1 day after a view.	External AI audit item 4; re-check OA1	qualified	Optimization windows (1-day click for app installs; 1- or 7-day click for iOS app events) separated from the 28-day comparison view and the legacy App Events API.	ch. 4; reconciliation register	OA1-S04 OA1-S05 OA1-S06 D-S16
COR-77	Chen et al.: a year of spend forecast with 5.7% average error against 9.0%.	FX4; External AI audit item 12; re-check OA3	qualified	Error metric described as scaled to the largest single player's spend; per-player SMAPE of 74% against 96% added; sample of about 2,500 churned payers; top spenders can bring up to half of revenue.	ch. 7; F11	G1-S04 FX4-S03
COR-78	Gordon et al. took 663 randomized experiments.	External AI audit item 34; re-check OA3	qualified	563 US experiments (November 2019 to March 2020) giving 663 test-and-control pairs, which the authors count as experiments; 'true lift' replaced by 'the lift the experiments measured'.	Abstract; ch. 4; CL-012; F06; study table	D-S02
COR-79	At eBay, a simple regression put the return on paid search above 1,600%; test in about a third of regions.	External AI audit item 36; re-check OA3	qualified	Standard regression above 4,000%; with region and day controls above 1,600%; test switched off non-brand search in about 30% of regions.	ch. 4; CL-013; study table	D-S04 OA3-S05
COR-80	Off-Meta data study: \$43.88 to \$60.19; working draft; commissioned evidence.	External AI audit item 37; re-check OA3 (the audit did not catch the superseded figures)	contradicted	Published figures used: \$38.16 to \$49.93 (+31%) across 70,000+ advertisers; publication terms stated; labelled platform-affiliated, not commissioned.	ch. 4; F06; study table	OA3-S01 OA3-S02
COR-81	About 72% of annual subscribers turned off auto-renew within the first year, up from 56%.	External AI audit item 7; re-check OA4	qualified	Main report figure used: median app kept 28% of annual subscribers after one year (2024 starts), down from 31%; the blog's 72%-versus-56% comparison flagged as mixing methods; first-year payment kept separate from renewal loss.	ch. 5, 7; F09; F11	GH-S07 GH-S08
COR-82	Billing failures caused 31% of Google Play and 14% of App Store cancellations; AI apps' higher first-year value, 36% worse retention.	External AI audit items 60-61; re-check OA4	qualified	Report chart values (32.2% and 15.2%) given with the summary's 31% and 14%; AI-app value stated per paying user; retention gap stated as 6.1% against 9.5% of monthly subscribers after 12 months.	ch. 7; F09	GH-S07 GH-S08

COR-83	Apple EU unified terms row presented as the rate at the cutoff, with a 5% core technology commission added.	External AI audit item 8; re-check OA2	qualified	Split into terms in force at the cutoff (opt-in alternative terms) and terms announced for 1 October 2026; the 5% commission applies only to apps outside the App Store.	ch. 7 fee table; policy register; GH-S03 title	OA2-S03 GH-S03
COR-84	Google Play fee rows and Apple Small Business Program row.	External AI audit items 63-64; re-check OA2	qualified	Program rates open only on 30 September 2026; external-link rate for existing installs; 15% tier requires enrolment; Apple SBP requires enrolment and counts associated accounts.	ch. 7 fee table	GH-S05 GH-S06 GH-S01
COR-85	In August 2026 ... the Supreme Court also declined to stay it.	External AI audit item 9; re-check OA2	qualified	The trial-court refusal (11 August) was already cited; the Supreme Court step now names Justice Kagan's administrative stay (12 August) and denial (13 August). The audit's claim that the trial-court refusal was unsourced was not borne out.	ch. 6; CL-021	FX3-S04 FX3-S03 EF-S17
COR-86	Adjust and Singular: no incrementality product confirmed.	External AI audit item 10; re-check OA4	contradicted	Both documented; InSight uses synthetic controls (modelled baseline) and Singular's FAQ does not say its split is random.	ch. 9; systems map	OA4-S01 OA4-S02
COR-87	AppLovin's advertiser pages we opened state no minimum signal.	External AI audit item 11; re-check OA1	contradicted	AppLovin recommends a budget buying at least 15-20 conversions a day; ad-revenue and blended goals need MAX.	ch. 5	V1-S07
COR-88	Google, Meta and Apple document campaign APIs; TikTok's and Amazon's could not be confirmed. TikTok value bidding is Android-only.	External AI audit items 14, 42, 83, 84; re-check OA1	qualified	All five channels document campaign APIs (D9 sources attached; Amazon D9 raised to 3); Android-only scoped to TikTok Smart+; Amazon iOS status stated as unconfirmed.	ch. 5, 9, 10; benchmark matrix	OA1-S06 OA1-S12 OA1-S13 OA1-S16 OA1-S17
COR-89	Meta and TikTok agent servers: no server-side approval step described.	External AI audit item 88; re-check OA1	qualified	Meta's server creates new campaigns paused and relies on the AI client to confirm activation; no approval step documented for edits to live campaigns.	Abstract; ch. 11; CL-030	OA1-S10 OA1-S11
COR-90	Meta A/B test and Dynamic Creative cited to a bundled Meta source whose first URL returns 404.	External AI audit item 52; re-check OA1	qualified	Split into Meta's own A/B testing and Dynamic Creative pages; noted that Dynamic Creative may be unavailable for new app-promotion ad sets since June 2024.	ch. 6; F08	OA1-S08 OA1-S09
COR-91	iOS postbacks described as pooled; the only install signal Apple allows for users who decline; everyone else works through pooled postbacks.	External AI audit items 12, 23-28; re-check OA5	qualified	Postbacks described as per-install reports with no user ID; windows counted from first launch; fine values only in the first postback; tiers include country and campaign code; other networks may add consented first-party attribution; all five EU countries named.	Abstract; ch. 1, 3, 4; CL-008; F01; F05	C-S02 C-S04 C-S05 OA5-S01 RS-S151
COR-92	Custom product pages: a 2.5-point lift; a custom page does not test anything.	External AI audit items 49-50; re-check OA5	qualified	Stated as 2.5 percentage points (156%) over a 1.6% default rate; a custom page creates no random split on its own but can be used in a designed test.	ch. 6; CL-019	MAIN-S04
COR-93	Braun and Schwartz: compare everyone assigned to each version, or run a lift test per version.	External AI audit item 51; re-check OA3	qualified	Replaced with the authors' own guidance: platform split tests predict rollout on that platform; isolating the message needs a design that holds delivery fixed. Evidence scope (one Meta test, web lead form) stated.	ch. 6	G1-S01 FX4-S01
COR-94	AppLovin filings neither confirm the June 2026 public opening.	External AI audit item 77; re-check OA4; FX4	qualified	Opening on 22 June 2026 without a referral code, under the name AppLovin Ads, now stated; filings still do not break out e-commerce revenue.	ch. 8; CL-033	OA4-S06
COR-95	MRC list: none of the major attribution companies appears; methods have not been through that audit. HUMAN 2023 figures.	External AI audit items 68, 70; re-check OA4	qualified	Named the four attribution companies checked; absence from the list no longer read as no audit; HUMAN figures dated May 2023 with their CTV-and-app scope and the one measure that favoured apps.	ch. 8; CL-024	IJ-S27 IJ-S21 OA4-S07
COR-96	Age-check rules, Texas status, AdCP governance.	External AI audit items 89-92; re-check OA2	qualified	Exact dates (Utah 6 May 2027, Louisiana 1 July 2027, California 1 January 2027); Texas stay and pending appeal; Google Age Signals beta scope; AdCP governance described as experimental and opt-in at the v3.1.24 tag.	ch. 11; policy register	FX2-S15 FX2-S18 FX2-S26 FX2-S17 OA2-S05 OA2-S06 OA2-S07
COR-97	Google Play policy on paid installs cited to a 2017 blog as current.	External AI audit item 54; re-check OA2	qualified	Current policy cited; the 2017 statement labelled as history.	ch. 6	OA2-S04 IJ-S14
COR-98	Company table and quarter notes: Digital Turbine EBITDA \$123m; Unity 'Grow segment'; Liftoff's 'first quarter as a public company'; Liftoff's IPO withdrawal timing; Verve release date.	External AI audit section 2, item 20; re-check OA2, OA4	qualified	Digital Turbine \$122m; Grow described as a revenue category; Liftoff's listing date and partial quarter stated; first IPO attempt withdrawn 17 February 2026; Verve release dated 26 January 2026 (preliminary) with the audited report cited.	ch. 2	AB-S07 G2-S07 G2-S09 FX1-S08 AB-S09 FX1-S12
COR-99	Liftoff profit column showed only net loss; AppsFlyer called the largest attribution company; its '2025' figures treated as a completed year.	External AI audit section 2, items 15, 18; FX4	qualified	Liftoff adjusted EBITDA \$374m (55%) added; 'largest' removed; report's first publication (December 2025) and unstated period noted.	Abstract; ch. 2	FX4-S06 MAIN-S01
COR-100	Unsupported rankings and absolutes: fastest-growing and least checked; weakest guide; most fraud; none answers; cannot tell; starts cold; most app businesses; widely used; understate.	External AI audit front-matter and chapter items 2-3, 8-13, 31-33, 38-40, 44-47, 55, 65-67, 71-72, 76, 86, 93	qualified	Rewritten as bounded statements (for example 'grew faster than acquisition', 'especially likely to include demand that already existed', 'we found none'); deep links and routing added as direct media levers; valid comparison groups defined; control's non-performance uses added.	Abstract; ch. 1, 2, 4, 5, 6, 7, 8, 10, 11; CL-001, 005, 011, 014, 018, 031; F01, F04, F09, F11, F12, F13	
COR-101	Appendix lacked the full scoring anchors; Apps DSP screen presented without threshold sensitivity; Markdown edition implied working instruments.	External AI audit items 6, 78, 87, 95, 100	qualified	Anchors for scores 1-3 published in Appendix H; screen labelled as the author's own, commerce omission explained and threshold sensitivity reported; n/e share now shown to one decimal with the benchmark channels excluded; Markdown edition labelled as the text component with CSV pointers and a note on the reference list versus the full register.	ch. 9, 11; appendix; Markdown edition	
COR-102	Short-seller allegations, retargeting and brand-search evidence, fingerprinting policy, click-to-cancel status, DMA data rights, AdCP and AAMP governance were placeholders.	Gap-fill G3 (red-team requests)	qualified	Filled from G3 findings, stated as allegations where they are allegations; re-checked by FX5 (11 supported, 13 qualified, none contradicted).	ch. 2, 3, 4, 5, 7, 11	G3-S01 G3-S02 G3-S03 G3-S07 G3-S10 G3-S12 G3-S13 G3-S16 G3-S17 G3-S18
COR-103	Moloco documents built-in ghost-bid holdouts for re-engagement and CTV; Remerge describes three holdout designs; Jampp's 86% against a ~6% control; Google's 'value above zero' requirement.	FX4	qualified	Moloco's re-engagement page names no method; Remerge's methods restated; Jampp's figures given without period, baseline or uncertainty; the value rule belongs to Demand Gen, not App campaigns.	ch. 4, 5	V1-S13 V1-S14 G1-S05 G1-S06 MAIN-S03

COR-104	G3 fills: Muddy Waters said 25-35% of conversions were incremental; brand rivals take 18-42% when a brand stops bidding; ghost-ads study at 'one online retailer'; FTC 'started over'; no press report of an SEC inquiry could be confirmed; Google lift API 'does not create studies'.	FX5	qualified	Muddy Waters wording taken from its report (about 52% retargeting; incrementality about 25-35%); the 18-42% applies to brands that chose not to advertise, with the authors' selection caveat; ghost-ads study described as a Google-built two-week retargeting test; FTC restored its pre-2024 rule and issued an ANPRM with no proposal by the cutoff; the court took the motion to dismiss and a motion to add reported SEC and state investigations under submission in April 2026; Reuters' relay of Bloomberg's SEC report added with AppLovin's no-comment; a separate September 2026 securities suit noted.	ch. 2, 3, 4, 5, 7	FX5-S01 FX5-S02 FX5-S10 FX5-S11 FX5-S12 FX5-S13 FX5-S19 FX5-S20
COR-105	Vendor profiles rendered first-pass prose (AppLovin financials and Tripledot terms, an unsupported 'independent geo-holdout' risk line, Google 'no placement reporting', Apple '21-day retention', process language).	Round-2 review R2-01	qualified	Profiles now render a reconciled overlay (data/profile_overrides.json): 142 field changes across 16 profiles; unsupported assertions removed; first-pass profiles kept as history.	ch. 9 profiles; published data	
COR-106	DMA Articles 6(8) and 6(10) 'effective 2022-09-14'.	Round-2 review R2-03	contradicted	Adopted 14 Sep 2022; in force 1 Nov 2022; applies from 2 May 2023; six months to comply after designation.	Policy register	G3-S12
COR-107	F05: 'one postback per install'; platform window 'up to 28-day click'.	Round-2 review R2-04	contradicted	Up to three postbacks per winning install, one per conversion window; losing networks get at most one; platform window 'up to 30 days after a click'.	F05	C-S02
COR-108	F08: vendor ghost-bid holdout labelled 'random holdout'.	Round-2 review R2-05	qualified	Relabelled vendor-reported holdout; assignment and analysis not published.	F08	G1-S06
COR-109	LTV model: web checkout as a fixed 5.4% fee at any price; year-two fee chosen separately.	Round-2 review R2-06	contradicted	Fees modelled as percentage plus fixed fee per payment on linked payment paths; defaults unchanged.	F10 model	G2-S15
COR-110	Marginal-return model: break-even \$70,000 read off a grid; selected spend could fall off the chart.	Round-2 review R2-20	qualified	Break-even solved by bisection (about \$69,300 at defaults); chart domain includes spend and break-even; 'none' shown when marginal never reaches the target.	F07 model	
COR-111	Amazon DSP D8 scored 3 on 'event-level data'; Mintegral D9 kept at 2 despite a documented campaign API; Digital Turbine D9 rationale incomplete.	Round-2 review follow-up	qualified	Amazon D8 = 2 (AMC outputs aggregated only); Mintegral D9 = 3 (Open API); DT D9 = 2 with rationale covering the Offer Wall API.	Capability matrix; AQ	V3-S13 RS-S63 RS-S77
COR-112	Claim drawers showed first-pass passages, locators and limitations for corrected claims.	Round-2 review R2-14	qualified	Each claim renders a current evidence record (data/claim_support.json): 157 verbatim passages, 28 new sources, retired findings listed; CL-002, CL-014, CL-015 reworded to match their evidence.	Claim ledger and drawers	
COR-113	Platform-run or platform-data studies labelled 'independent measurement' (Gordon, Lewis & Rao, Blake, ghost ads, Pandora, Simonov); offsite-data study labelled 'commissioned'.	Round-2 review R2-15	qualified	New class 'affiliated measurement'; 'commissioned' reserved for documented commissioning.	Evidence ledger	
COR-114	Offsite-data study 'published in 2025'.	Round-2 review R2-24	qualified	Posted online September 2024; printed in the March 2025 issue.	ch. 4	OA3-S02
COR-115	'613 distinct sources'; 504 'read in full' by a prefix test.	Round-2 review R2-13	qualified	Counted as distinct source addresses (URLs, not works); read-in-full uses an enumerated access depth that excludes abstract, navigation and headline-only records.	Abstract; ch. 12; source register	
COR-116	Absolutes: 'unvalidated', 'no buyer survey', 'mostly steals credit', 'needs no test', 'only public evidence', 'most profitable'.	Round-2 review R2-21	qualified	Search-bounded wording throughout; 'highest adjusted EBITDA margin in our table'.	Abstract; ch. 5, 7, 8, 10, 11; CL-005, CL-029	
COR-117	Non-gaming evidence (D11) scored 1 for Kayzen, Remerge, Verve Datasheet and Smadex.	Post-publication author review, 28 Sep 2026	qualified	Corrected on 28 Sep 2026: Kayzen, Remerge and Verve Datasheet D11 from 1 to 2 (two or more named non-gaming case studies); Smadex D11 from 1 to 3 (named cases plus filing-level evidence of material non-gaming customers). App Quotient and bands recomputed.	ch. 9 matrix and App Quotient; capability-matrix.csv; aq.csv	RS-S89 RS-S114 RS-S126 RS-S97
COR-118	Readability: the main prose scores 57.6 against a target of 85. Chapter 12 kept 85 as the target for a public summary that did not yet exist.	Author decision, 28 Sep 2026, following the readability exception accepted at the round-3 close-out	qualified	Added a plain-language edition at plain/ (web only): a new, shorter text written for general readers with the same findings, counts filled from the same build data, and links to each chapter. Its build fails under Flesch 85 or if a stated number is not in the paper. Chapter 12, the footer and the README now say the paper's score stays a disclosed exception and the plain edition carries the target of 85.	new plain/ page; ch. 12 Readability; paper footer and top bar; README	
COR-119	App Quotient shrinkage target: the field average was the mean of all group scores pooled equally (2.1187 on the 0-3 scale).	Method review shared across the author's research papers, 28 Sep 2026	qualified	Method note added to chapter 9 after the AQ construction text. Scores and bands are unchanged as published.	ch. 9 App Quotient (note only); aq.csv unchanged	

APPENDICES

Evidence and reproduction

Every headline in this paper can be traced from the prose to a claim record, from the claim to findings and sources, and from any number to its formula or data file.

HOW TO TRACE A CLAIM

In the HTML edition, click a claim chip such as CL-012 to open its record: final wording, evidence class, review outcome, confidence, locator, passage and sources. Click a numbered reference to open the source with its date, access limits and funding interest. Every chart has a static SVG and a CSV; the Figure 14 matrix is a table with its own CSV. The scripts in the package rebuild all of it from the evidence files.

A. Claim ledger

Claim ledger: the paper's headline claims						
Each chip in the text opens its record here. Review outcome comes from the verification passes on the underlying findings.						
CLAIM	CH.	FINAL WORDING	EVIDENCE CLASS	REVIEW OUTCOME	CONFIDENCE	SOURCES
CL-001	1	Paid media acts on activation and retention mostly indirectly: through whom it reaches, what the ad promises and where a deep link lands the user. Reactivation, which remarketing buys directly, is the exception.	synthesis	AUTHOR ANALYSIS	MEDIUM	[2,1]
CL-002	1	A lower CPI or a higher attributed ROAS can coexist with falling contribution, for example when credit goes to users who would have come anyway, when reports use gross bookings, or when cheaper installs churn faster (a mechanism the author flags, not a measured rate).	inference	QUALIFIED	HIGH	[114,5]
CL-003	2	AppsFlyer estimates 2025 global app marketing spend at \$109B: \$78B user acquisition (+13%) and \$31.3B remarketing (+37%); remarketing's share rose from 25% to 29%.	vendor_assertion	QUALIFIED	MEDIUM	[6]
CL-004	2	Sensor Tower and Appfigures estimate 2025 global app-store consumer spend at \$167B and \$155.8B, a gap of about \$11B in level and 11 points in growth rate.	synthesis	QUALIFIED	HIGH	[7,8]
CL-005	2	Remarketing grew faster than acquisition in AppsFlyer's panel (+37% against +13%). Randomized web retargeting studies find real but modest lift, and we did not locate a public independent randomized app study that validates attributed app-remarketing returns.	inference	QUALIFIED	MEDIUM	[6,111,114]
CL-006	2	Ad-platform revenues are not comparable as reported: AppLovin and Liftoff book revenue net of publisher payouts, Mobvista books gross, Digital Turbine books part of its business gross, and Verve moved some revenue from net to gross in 2025.	direct_record	SUPPORTED	HIGH	[21,22,14,17,19]
CL-007	2	AppLovin's own disclosure says its Axon model uses win and loss notices from its MAX mediation auction, alongside device, engagement and advertiser-shared data.	direct_record	SUPPORTED	HIGH	[29]
CL-008	3	Under SKAdNetwork 4, Apple assigns each attributed download one of four postback data tiers based on crowd size across the app or site showing the ad, the advertised app, the install country and the campaign code; higher tiers get more campaign-code digits and, in the first postback only, a fine-grained conversion value, so installs concentrated in larger crowds receive richer feedback. Apple publishes no thresholds.	synthesis	SUPPORTED	HIGH	[1]
CL-009	3	Apple's release notes list no SKAdNetwork 5, and no Apple document opened sets a retirement date for SKAdNetwork; it is bridged with AdAttributionKit.	direct_record	SUPPORTED	HIGH	[77,76]
CL-010	3	At the cutoff Google listed the Android Privacy Sandbox eds APIs as 'scheduled for phaseout' with no completed-removal date; the retirement was announced, not finished.	direct_record	SUPPORTED	HIGH	[95,94]
CL-011	4	Summing conversions claimed by several self-attributing platforms, or adding platform reports to Apple postbacks, counts some installs twice; such a total is not a de-duplicated install count.	inference	SUPPORTED	HIGH	[68,109]
CL-012	4	Across 863 randomized test-and-control comparisons on Facebook (from 563 US experiments run November 2019 to March 2020), the median lift the experiments measured was 29%, 18% and 5% by funnel stage, while double machine learning estimated median lifts of 83%, 58% and 24% and propensity matching 173%, 176% and 64% on the same campaigns.	independent_measurement	SUPPORTED	HIGH	[111]
CL-013	4	In eBay's 60-day geo experiment, which switched off non-brand search ads in about 30% of US regions, regression estimates of paid-search return of 4,173% (no controls) and 1,632% (region and day controls) compared with an experimental estimate of minus 63%; brand-keyword ads showed no measurable short-term benefit.	independent_measurement	SUPPORTED	HIGH	[114]
CL-014	4	Meta, Google and TikTok each document a conversion-lift product with test and control groups (Meta's guide calls its test randomized); in the pages reviewed, each routes access through an account team or eligibility rules rather than open self-service.	direct_record	QUALIFIED	MEDIUM	[125,126,127]
CL-015	5	Google's app campaigns require at least 10 conversions a day (or 300 in 30 days), with bid-on events coming from Firebase, for target-ROAS bidding, and warn against edits before 100 conversions.	direct_record	QUALIFIED	HIGH	[27,139]
CL-016	5	The bidding engines learn best from events that are common, early and valued; in 2025 budgets nevertheless grew fastest on iOS and in non-game apps, where signal is thinner.	synthesis	QUALIFIED	MEDIUM	[27,145,147,142,6]
CL-017	5	On a saturating response curve a channel can show a healthy average return while the marginal dollar returns less than it costs.	illustrative_scenario	MODEL PROPERTY	HIGH	[135,136]
CL-018	5	Attributed returns for remarketing and brand search are especially likely to include demand that already existed, so holdout tests, not attribution reports, should set these budgets.	inference	QUALIFIED	MEDIUM	[114,111]
CL-019	6	Apple documents no random split or significance test for custom product pages, which route traffic but do not by themselves create random assignment; its own figure for them, a 2.5 percentage-point average lift (156% over a 1.6% default-page rate), compares pages that receive different traffic.	direct_record	SUPPORTED	HIGH	[157,159]
CL-020	6	Ad-platform delivery algorithms send split-test versions to different audiences, so a creative 'win' mixes the ad's effect with targeting (divergent delivery).	independent_measurement	SUPPORTED	MEDIUM	[162]
CL-021	6	On 30 June 2026 the Supreme Court agreed to hear Apple's appeal limited to whether contempt can rest on breaking an injunction's 'spirit'. The permissible US link-out fee is being set by the trial court, which refused Apple's request to pause that work on 11 August 2026; Justice Kagan denied Apple's stay application on 13 August. No fee had been approved at the cutoff.	direct_record	SUPPORTED	HIGH	[170,99,97]
CL-022	7	For a \$9.99 first-year subscription payment, web checkout at Stripe's 2.9% + \$0.30 keeps about \$9.40, about 34% more than Apple's 30% tier (\$6.99) and about 11% more than the 15% tier (\$8.49), counting card fees only.	synthesis	SUPPORTED	HIGH	[3,187]
CL-023	7	Google's zero-inflated lognormal LTV method was validated by its authors on a retail shopper dataset and a charity donor dataset, not on app install cohorts.	direct_record	SUPPORTED	HIGH	[189,190]
CL-024	8	MRC's digital accreditation listing (checked 27 September 2026) shows sophisticated invalid-traffic accreditation covering mobile in-app for several firms, including DoubleVerify, HUMAN, IAS, Pixalate, Protected Media and some Google and Meta services, but not AppsFlyer, Adjust, Singular or Kochava. Absence from the list does not show that no audit has taken place.	direct_record	SUPPORTED	MEDIUM	[200]

CL-025	8	Uber's 2017 federal suit against Fetch Media over allegedly fraudulent mobile ad credit was dismissed that year; in a separate state case Uber, Phunware and four individuals settled for \$6 million in October 2020, all denying wrongdoing.	direct_record	SUPPORTED	HIGH	[201,202,203]
CL-026	8	Our searches found no independent, published randomized or geographic study of streaming TV advertising effects on app installs at the cutoff; every CTV-to-app lift claim located was vendor-designed and vendor-reported.	inference	QUALIFIED	MEDIUM	[212,208,123]
CL-027	10	The big platforms' automated app campaigns choose placements for the buyer. Google's App campaigns placement report names apps, publishers, websites and videos but gives impressions only and groups low-volume placements; TikTok's Smart+ does not allow manual placement choice; Meta lists placement among reporting breakdowns but describes no breakdown by publisher app.	direct_record	SUPPORTED	HIGH	[222,223,142]
CL-028	10	No public study was found that measures the net value of transparency or self-serve control to app buyers after staff and tooling costs.	inference	QUALIFIED	MEDIUM	
CL-029	11	'Apps DSP' describes a real capability checklist. On the author's own seven-field screen, one of eleven execution products documents the full set and five document at least five of seven; scores describe documentation, not performance, and we did not find public comparative evidence that the bundle outperforms best-of-breed parts.	synthesis	AUTHOR ANALYSIS	MEDIUM	
CL-030	11	TikTok's agent server exposes about 400 advertising functions, including campaign creation and budget changes, and its help page says agents can work 'without human intervention at every step', describing no per-change approval step. Meta's server creates new campaigns paused and relies on the AI client to confirm activation, but documents no approval step for edits to live campaigns.	direct_record	SUPPORTED	HIGH	[248]
CL-031	11	AppLovin's public disclosure on its Axon model is silent on advertisers' rights over the use of their campaign data; the pages reviewed did not establish portable rights to any platform's learned state, and other platforms' full advertising terms were not reviewed.	direct_record	QUALIFIED	MEDIUM	[29]
CL-032	7	In a 21-month randomized experiment across about 35 million Pandora listeners, long-run sensitivity to ad load was about three times what a one-month test would show, observational estimates were biased, and heavier ad load raised paid-subscription conversion.	independent_measurement	SUPPORTED	MEDIUM	[120]
CL-033	8	AppLovin opened its self-serve platform, renamed AppLovin Ads, to all advertisers without a referral code on 22 June 2026 (chief executive's blog); its filings through Q2 2026 mention e-commerce as an expansion area but do not break out e-commerce revenue.	direct_record	SUPPORTED	MEDIUM	[21,24,215]

B. Source register

The full register of 720 source records is in the HTML edition and in data/tables/source-register.csv. The print edition lists every cited source, with its address, in the References at the end.

C. Evidence ledger

Evidence ledger: every research finding						
All findings from the research streams, with the outcome of the fact-check passes. Corrected wording is what the paper uses where a finding was qualified.						
ID	FINDING AS RECORDED	CLASS	REVIEW	CONFIDENCE	SOURCES	CORRECTED WORDING
AB-F01	AppLovin's continuing (advertising) business reported FY2025 total revenue of \$5,480.7 million, up 70% year-over-year.	direct_record	SUPPORTED	high	AB-S02	
AB-F02	AppLovin reported FY2025 Adjusted EBITDA of \$4,512.5 million, an 82% margin on revenue, up 87% year-over-year.	direct_record	SUPPORTED	high	AB-S02	
AB-F03	AppLovin completed the sale of its Apps (mobile gaming studios) business to Triplicated Studios on 30 June 2025, with announced consideration of \$400 million cash plus approximately 20% of Triplicated Studios' fully-diluted equity.	direct_record	SUPPORTED	high	AB-S03 AB-S01	
AB-F04	AppLovin's FY2025 10-K XBRL tagging shows the divested Apps business classified under discontinued-operations member tags across 2023-2025 periods, consistent with the company now reporting its advertising platform as effectively a single continuing-operations business line.	inference	SUPPORTED	medium	AB-S01	
AB-F05	Unity Software reported FY2025 total revenue of \$1,849.6 million (up from \$1,813.3 million in FY2024); Q4 2025 Grow Solutions segment revenue was \$338 million versus Create Solutions at \$165 million.	direct_record	SUPPORTED	high	AB-S05 AB-S04	
AB-F06	Unity's Grow Solutions segment includes both the ironSource ad network (11% of Grow Solutions revenue in Q4 2025, per company commentary, and described as declining) and the Vector machine-learning bidding model (56% of Grow Solutions revenue in Q4 2025), which Unity states rolled out in Q1 2025.	vendor_assertion	QUALIFIED	medium	AB-S05 AB-S04	In Q4 2025 the Unity / migrated to Unity's Ve (rolled out Q1 2025), r Grow Solutions revene IronSource Ad Networ 10-K; Q4 2025 releas
AB-F07	Unity reported FY2025 Adjusted EBITDA of \$408.8 million (22% margin) alongside a GAAP net loss of \$401.5 million.	direct_record	SUPPORTED	high	AB-S05	

73

AB-F23	The IAB/PwC Internet Advertising Revenue Report (30th edition) reports total US digital advertising revenue of \$294.6 billion in 2025, up 13.9% year-over-year; this figure spans all digital-advertising channels and formats and was not broken out by mobile app-install/in-app scope in the accessible (non-paywalled) portion of the report.	commissioned_measurement	SUPPORTED	high	AB-S14	
AB-F24	AppLovin's advertising platform integrates a demand-side ad network (Axon Ads Manager/AppDiscovery), a mediation layer (MAX), a mobile measurement partner (Adjust, acquired 2021), and a CTV/streaming distribution business (Wurl) under single corporate ownership.	direct_record	QUALIFIED	high	AB-S01 AB-S02	AppLovin owns its UA (AppLovin Ads, former Manager), the MAX mediation platform, the Wurl (acquired April 2021) CTV platform (acquired April 2021) in a single segment. App data is not shared with customer directs.
AB-F25	Unity's Grow Solutions segment combines the ironSource ad network (acquired via the 2022 merger) with the LevelPlay mediation platform under common ownership.	direct_record	QUALIFIED	medium	AB-S05 AB-S06	Unity's Grow Solutions is a mediation platform (LevelPlay), two owned Ad Network on Vector Network), an offerwall publishing under common ownership.
AB-F26	Digital Turbine's business combines On Device Solutions (OEM/device preload distribution) with App Growth Platform (an ad exchange/mediation business) under common ownership.	direct_record	QUALIFIED	medium	AB-S07	DT (formerly Digital Turbine) Device Solutions (OEM device preload distribution), Singl its App Growth Platform Wall and brand campaign inventory, plus SDK-based monetization for publishers common ownership. L assets were being sold
AB-F27	No accessible primary or market-research source in this pass provided a specific dollar figure for either 2025/2026 US app-install advertising spend (Q1) or in-app ad revenue delivered inside apps (Q2); EMARKETER maintains dedicated forecast series for both but the underlying figures are paywalled.	inference	QUALIFIED	high		No census-grade or in US app-promotion spend revenue (Q2) was found figure is AppsFlyer's v estimate of \$109B global spend in 2025, and Lit prospectus cites an A of a ~\$79B 2025 served independent in-app ad
C-F01	SKAdNetwork 4 (iOS 16.1+) supports up to three postback conversion windows -- days 0-2, 3-7, and 8-35 -- each sent after a random delay of 24-48 hours (first postback) or 24-144 hours (second/third).	direct_record	SUPPORTED	high	C-S02	
C-F02	SKAdNetwork 4 assigns each install a 'postback data tier' (crowd anonymity tier 0-3) that determines whether the postback contains a fine-grained conversion-value, a coarse-conversion-value only, a 2/3/4-digit source-identifier, and an optional country-code; Tier 0 apps receive only a single postback with a 2-digit source-identifier and no conversion value at all.	direct_record	SUPPORTED	high	C-S02	
C-F03	As of the research cutoff, Apple's official SKAdNetwork release notes list only versions 1, 2, 2.1, 2.2, 3, and 4 -- no SKAdNetwork 5 has shipped.	direct_record	SUPPORTED	high	C-S03	
C-F04	Apple's official documentation directs developers to 'Use AdAttributionKit for app ad campaigns on the App Store and alternative marketplaces,' but a separate interoperability document confirms the two frameworks currently coexist and are bridged (a SKAdNetwork conversion-value update call is mirrored into AdAttributionKit); no retirement or deprecation date for SKAdNetwork is stated anywhere in the reviewed official docs.	direct_record	SUPPORTED	high	C-S01 C-S05	
C-F05	AdAttributionKit (introduced iOS 17.4, March 2024) added re-engagement measurement and a re-engagement universal-link URL in June 2024 (iOS 18 cycle), and added postback country-code, configurable attribution rules, and cooldown windows in June 2025 (iOS 26 cycle, WWDC25).	direct_record	SUPPORTED	high	C-S04 C-S06	
C-F06	Apple's AdAttributionKit Updates changelog has no entries after June 2025; no new AdAttributionKit or SKAdNetwork functionality was found announced for WWDC26 (June 8, 2026, iOS 27 cycle) in official Apple documentation reviewed in this research.	direct_record	QUALIFIED	medium	C-S06	Apple's AdAttributionKit entries after June 2025.7.2 (beta) release no AdAttributionKit or SK the only post-June-202 API found is Apple's at iOS 26.2, which the record.
C-F07	A Beta API, requestTrackingAuthorization(usingExpandedInterface:additionalInformationAction:completionHandler:), available iOS/iPadOS/Mac Catalyst 27.2+, lets an app present a full-page, Markdown-formatted ATT consent sheet with an optional 'Additional Information' button; in France, Germany, Italy, Poland and Romania specifically, the full-page Markdown sheet is shown regardless of the app's own preference setting, while elsewhere in the EU it is shown only if the app opts in, and outside the EU the API has no effect (falls back to the standard system alert).	direct_record	SUPPORTED	high	C-S08	
C-F08	France's Autorite de la concurrence fined Apple EUR150 million on March 31, 2025 (decision 25-D-02) for abuse of a dominant position through its ATT implementation between April 26, 2021 and July 25, 2023, finding the ATT principle 'not problematic' but its implementation 'abusive within the meaning of competition law.'	direct_record	SUPPORTED	high	C-S11	
C-F09	Italy's AGCM fined Apple Inc., Apple Distribution International Ltd, and Apple Italia S.r.l. a combined EUR98,635,416.67 on December 22, 2025 for abuse of dominance, finding that ATT's 'double consent' requirement (Apple's prompt plus GDPR-compliant consent) harmed third-party developers and advertisers; the proceeding began May 2, 2023 after a complaint by Meta.	direct_record	QUALIFIED	high	C-S13	Italy's AGCM fined Apple Distribution International Italia S.r.l. a combined on 22 December 2025 dominance, finding ATT requirement disproportionate investigation (A561) w May 2023.
C-F10	Germany's Bundeskartellamt, on August 17, 2026, declared legally binding commitments requiring Apple to align consent-prompt wording/design for its own apps and third-party apps, remove discouraging symbols/wording from third-party prompts, and give app publishers more room to combine Apple's ATT consent request with GDPR-mandated consent requests; Apple has four months from service of the decision to implement the changes, with independent-trustee monitoring for seven years.	direct_record	SUPPORTED	high	C-S12	

C-F11	The Beta EU-specific expanded ATT interface (iOS 27.2, see C-F07) and the German Bundeskartellamt's binding commitments (see C-F10) describe closely overlapping remedies (more explanatory content, less discouraging framing, easier-to-combine consent flows), but no official Apple document reviewed in this research cites the Bundeskartellamt decision, the French decision, or the Italian decision as the reason for the new API.	inference	SUPPORTED	medium	C-S08 C-S12	
C-F12	On October 17, 2025, Google announced it will phase out Attribution Reporting, Topics, and Protected Audience (each across both Chrome and Android), plus IP Protection, On-Device Personalization, Private Aggregation/Shared Storage, Protected App Signals, Related Website Sets, SelectURL and SDK Runtime, while continuing to support CHIPS, FedCM and Private State Tokens.	direct_record	SUPPORTED	high	C-S15	
C-F13	As of the 2026-09-27 cutoff, Google's own Privacy Sandbox feature-status page lists Android's Attribution Reporting, Topics, Protected Audience, SDK Runtime, On-Device Personalization and Protected App Signals as 'Scheduled for phaseout' -- an announced-retirement status -- without giving a specific completed-removal date; Google states only that it 'will follow Chrome and Android processes for phasing out these technologies.'	direct_record	SUPPORTED	high	C-S14 C-S15	
C-F14	Google's Play Install Referrer API documentation states that devices with Google Play app version 8.3.73 or later automatically have API access, and that it returns the referrer URL, click/install timestamps (client- and server-side), first-install app version, and whether the user interacted with an instant experience in the past 7 days; a (now-outdated, 2023-era) Privacy Sandbox FAQ stated Privacy Sandbox changes would not affect Play Install Referrer functionality.	direct_record	QUALIFIED	high	C-S17 C-S23	Google's Install Referr (updated 2025-07-21) 8.3.73+ has automatic referrer URL, click/ins install version and 7-d experience flag; a Priv last updated 2025-03 2025 phaseout annou Privacy Sandbox woul Referrer.
C-F15	Google's Advertising ID (AAID) policy requires apps that target Android 13+ to declare the com.google.android.gms.permission.AD_ID manifest permission (rollout began late 2021 on Android 12, enforceable for all Play-supported devices from April 1, 2022), and returns a string of all zeros when a user resets the identifier -- this mechanism is distinct from, and unaffected by, the Privacy Sandbox retirement.	direct_record	CONTRADICTED	high	C-S18 C-S16	Google Play returns a a user deletes their ad out on Android 12 from Play devices from Apr targeting Android 13+ com.google.android.g permission. AAID is se Privacy Sandbox phas
C-F16	Apple's AdServices (Apple Ads Attribution API) changelog documents a 30-day tap-through/click attribution window (added November 2024) and introduced view-through attribution on March 27, 2025; a subsequent October 2025 changelog entry on pre-order attribution implies the standard view-through lookback window is 1 day (vs. 30 days for click-through/tap-through).	direct_record	QUALIFIED	medium	C-S10	Apple's AdServices re day tap-through and 2 through attribution wi introduced March 27, prioritized over view-t
C-F17	Apple's AdServices changelog entry for September 2026 (the same month as the research cutoff) states that attribution payloads will return attribution = false for campaigns using age or gender targeting.	direct_record	SUPPORTED	high	C-S10	
C-F18	AppsFlyer, Adjust and Singular each document, in their own official help/developer centers, live support for receiving copies of both SKAdNetwork and AdAttributionKit postbacks at a vendor-specified endpoint; Singular additionally documents a dedicated 'AdAttributionKit - Opt in for Reengagement Copies' Info.plist key specifically for AdAttributionKit re-engagement postbacks.	vendor_assertion	QUALIFIED	high	C-S19 C-S20 C-S21	AppsFlyer, Adjust and document how an app winning SKAdNetwork AdAttributionKit postb endpoint; Singular als opt-in key for winning engagement postback
C-F19	Apple's own Apple Ads help documentation describes AdAttributionKit as 'the free attribution solution Apple has provided to the mobile advertising industry since 2018' -- a continuity claim that conflates AdAttributionKit's actual March 2024 (iOS 17.4) launch with SKAdNetwork's original 2018 (iOS 11.3) launch.	direct_record	SUPPORTED	high	C-S22 C-S04	
C-F20	Adjust's vendor-panel Mobile App Trends 2026 report states the global average App Tracking Transparency opt-in rate rose from 35% in Q1 2025 to 38% in Q1 2026, based on data from 'thousands of apps' in Adjust's own client base.	vendor_panel	QUALIFIED	medium	C-S24 C-S25	Adjust's vendor-panel 2026 (published Feb 2 plus its tracked datae ATT opt-in among iOS prompt rising from 35 (early Q1 2026, a parti
C-F21	A market-research aggregator (Business of Apps, compiling multiple vendor panels without a single unified methodology) reports gaming apps have the highest ATT opt-in rate among categories at 39%, and the UAE has the highest opt-in rate by country at close to 50%.	vendor_panel	QUALIFIED	low	C-S26	Business of Apps (upc 2026; underlying pane visible without sign-up the highest ATT opt-in UAE the highest count 50%; the 39% matche figure.
C-F22	For install conversions, Apple's attribution frameworks send a single winning postback to one ad network and a single non-winning postback to up to five other qualifying ad networks; for re-engagement conversions, AdAttributionKit generates no 'runner-up' postbacks at all -- only the winning ad network receives postbacks, across multiple conversion windows.	direct_record	QUALIFIED	high	C-S04	For installs, one winni receives postbacks ac conversion windows a qualifying networks e winning postback; for only the winning netw postbacks (possibly a no runner-up postbac
C-F23	Germany's Bundeskartellamt proceeding against Apple's ATT implementation followed a multi-year regulatory timeline distinct from France's and Italy's: Apple's 'paramount significance for competition' designation issued April 2023 (confirmed by the Federal Court of Justice March 2025), a preliminary assessment in February 2025, a market test of Apple's offered commitments opened December 2025, and the commitments made legally binding August 17, 2026.	direct_record	SUPPORTED	high	C-S12 C-S27	
C-F24	No official Apple documentation opened in this research confirms an exact ship date (GA) for the Beta expanded-ATT-interface API; it is documented as Beta at iOS 27.2, and Apple's public release-note conventions elsewhere in this research (e.g. AdAttributionKit, AdServices) show Apple does date-stamp features once shipped -- the absence of such a dated, non-Beta entry is treated as 'not yet GA', per the brief's instruction not to convert an announcement into a completed change.	direct_record	SUPPORTED	high	C-S08	

D-F01	Across 663 Facebook ad experiments (Nov 2019-Mar 2020, ~7.9B user-experiment observations, 38B impressions), median RCT-measured lift was 29% (upper funnel), 18% (middle funnel) and 5% (lower funnel/purchase), while double/debiased ML (DML) non-experimental estimates showed median relative errors of 83%, 58% and 24% respectively, and stratified propensity-score matching (SPSM) showed 173%, 176% and 64%.	affiliated_measurement	QUALIFIED	high	D-S02	Across 663 US Facebook ad experiments (Nov 2019-Mar 2020), median RCT-measured lift was 29% (upper funnel), 18% and 5% by funnel/purchase, while double/debiased ML (DML) non-experimental estimates showed median relative errors of 83%/58%/24% (DML), and stratified propensity-score matching (SPSM) showed 173%/176%/64% (SPSM).
D-F02	In an earlier, separate study of 15 large-scale US Facebook ad experiments (~500M user-experiment observations, 1.6B impressions), observational methods commonly used in industry practice often failed to recover the same causal effects as the randomized experiments, even after conditioning on extensive demographic and behavioral data.	affiliated_measurement	SUPPORTED	high	D-S01	
D-F03	Across 25 large-scale US digital-ad field RCTs (\$2.8M combined spend, 19 retailers + 6 financial firms), the median 95% confidence interval on measured advertising ROI was over 100 percentage points wide, and individual-level sales volatility (coefficient of variation ~10 relative to per-capita ad cost) means an informative experiment can require more than 10 million person-weeks.	affiliated_measurement	SUPPORTED	high	D-S03	
D-F04	In a 60-day, 210-DMA eBay paid-search field experiment (ads suspended in ~30% of DMAs), naive OLS estimates put paid-search ROI at 4,173% (no controls) or 1,632% (with time/geo controls), while the experimental estimate was -63%; brand-keyword ads showed no measurable short-term benefit.	affiliated_measurement	SUPPORTED	high	D-S04	
D-F05	Using Kantar Vivivix ad-spend, Shopify revenue and SimilarWeb traffic data, a 2025 study finds conversion-optimized Meta ads saw a 37% reduction in click-through rate after Apple's ATT rollout, and e-commerce firms with higher pre-ATT Meta dependence saw firm-wide revenue declines of 8-40% relative to less-exposed firms, concentrated among smaller firms.	independent_measurement	QUALIFIED	medium	D-S05	Using opt-in panels from Facebook and Instagram ad-analytics providers, a 2025 study benchmarked against Shopify and SimilarWeb (Management Science) conversion-optimized ads saw a 37% CTR reduction after ATT, and e-commerce firms more dependent on Meta saw revenue declines of 8-40% relative to less-exposed firms, concentrated among smaller firms.
D-F06	A large-scale Meta-run experiment across hundreds of thousands of advertisers found that removing offsite/off-Meta signal data raised the median cost per incremental customer from \$43.88 to \$60.19 (a 37% increase), with ads using offsite data also generating more long-term customers per dollar measured six months later.	affiliated_measurement	CONTRADICTED	medium	D-S06	In a one-week randomized experiment, Facebook and Instagram employees covering more than 70% of ad delivery (Wernerfelt, Tuchman, Marketing Science 2023) found that removing offsite/off-Meta signal data raised the median cost per incremental customer from \$43.88 to \$60.19 (+37%) under the median baseline was estimated at \$43.88. The study was superseded by a 2022 study.
D-F07	A survey experiment with 11,000 US/UK adults found opt-in rates to Apple's own Personalized Ads consent prompt (~25%) were roughly double opt-in rates to the ATT third-party tracking prompt (~13%), a 12.4 percentage-point gap, with the gap widening to 15.1pp among users who stated a preference for personalized ads.	independent_measurement	QUALIFIED	medium	D-S07	A survey experiment with 11,000 US/UK adults found opt-in rates to Apple's own Personalized Ads consent prompt (~25%) were roughly double opt-in rates to the ATT third-party tracking prompt (~13%), a 12.4 percentage-point gap, with the gap widening to 15.1pp among users who stated a preference for personalized ads.
D-F08	A working paper using web-scraped data on 580,000+ apps in a difference-in-differences design (Apple vs Google Play as comparison) reports that Apple's ATT is associated with a small increase in apps adopting in-app payments after its introduction, strengthening a pre-existing industry trend rather than creating a new one.	independent_measurement	QUALIFIED	low	D-S08	Kesler's working paper (2023; abstract only) reports that Apple's ATT is associated with a small increase in apps adopting in-app payments after its introduction, strengthening a pre-existing industry trend rather than creating a new one.
D-F09	The 'ghost ads' method (Johnson, Lewis & Nubbemeyer 2017) identifies the counterfactual control-group equivalent of ad-exposed users inside a real-time ad-delivery/auction system, which the authors state can reduce experimentation cost and improve measurement precision relative to public-service-announcement (PSA) or intent-to-treat A/B tests, while remaining compatible with platforms that optimize ad delivery dynamically.	affiliated_measurement	SUPPORTED	low	D-S09	
D-F10	AppsFlyer's official documentation states its click-through lookback window defaults to 7 days ("regarded as the industry standard", configurable 1-30 days), view-through lookback defaults to 1 day (configurable 0-24 hours), and engaged-click lookback defaults to 2 days (configurable 1-7 days); deterministic attribution methods (referrer/device-ID matching) are given priority over probabilistic modeling when both occur within the lookback window.	direct_record	SUPPORTED	high	D-S13	
D-F11	AppsFlyer classifies networks like Meta, Google Ads and Snapchat as 'self-reporting networks' (SRNs): AppsFlyer queries these networks' own MMP APIs using the device ID at first app launch (subject to ATT consent on iOS) to determine attribution, a separate mechanism and priority path from non-SRN networks, which rely on install-referrer, device-ID matching or probabilistic modeling.	direct_record	SUPPORTED	high	D-S13	
D-F12	AppsFlyer's reattribution window rules mean a reinstall occurring inside the configured window generates no install postback to any media source and subsequent in-app events are classified organic, whereas a reinstall occurring after the window expires is treated as a standard new install and does generate postbacks.	direct_record	QUALIFIED	high	D-S14	In AppsFlyer, a reinstall occurring inside the configured window generates no install postback to any media source and subsequent in-app events are classified organic, whereas a reinstall occurring after the window expires is treated as a standard new install and does generate postbacks.
D-F13	TikTok discontinued its legacy MMP-postback integration on 31 March 2025, making Self-Attributing Network (SAN) integration mandatory for all apps; TikTok's own documentation states SAN integration provides enhanced conversion recognition and reporting inside TikTok Ads Manager 'without affecting your MMP's final attribution logic,' meaning the two systems can diverge by design rather than by error.	direct_record	SUPPORTED	high	D-S23	

D-F14	TikTok's ad-group attribution settings offer click-through windows of 1, 7 or 28 days; view-through windows of Off, 1 or 7 days; and, for app campaigns specifically, engaged view-through windows (users who watch an ad 6+ seconds without clicking) of 1 or 7 days.	direct_record	QUALIFIED	high	D-S22	TikTok ad-group attribution through windows of 1, view-through Off, 1 or campaigns engaged view watched, or full video days.
D-F15	Singular's documentation of Meta's Advanced Mobile Measurement (AMM, available since 18 June 2025) states that in Meta's aggregate reporting, 'engaged view conversions are reported under conversion type = click-through,' meaning users who merely watched (rather than clicked) an ad can be counted toward click-attributed totals in aggregate reports.	direct_record	QUALIFIED	high	D-S24	Singular's documentation of aggregate reporting of engaged-view conversions over 5 seconds) are reported under conversion type 'click-enabled user-level Ad Measurement (AMM) 2025.
D-F16	Meta's App Events API documentation states that its attribution insights are 'surfaced based on impression or click time, not install or conversion time,' which the documentation notes may cause discrepancies versus Ads Manager reports; the API's stated attribution model is 28-day click-through and 1-day view-through (when view-through is enabled).	direct_record	QUALIFIED	high	D-S16	Meta's legacy App Events API no longer recommended says its attribution insights are surfaced from clicks within 30 (optional) and describes 28-day click/1-day view-through. API docs allow app-in-app optimization on 1-day click-through and 1-day view or engaged view-through offers 7-day and 28-day comparison windows.
D-F17	Per Adjust's documentation of Apple's SKAdNetwork 4, up to three delayed postbacks are sent per install/conversion-value window (0-2 days, 3-7 days, 8-35 days) with 24-144 hour delivery delays; conversion values are 'fine' (6-bit, 0-63) or 'coarse' (none/low/medium/high buckets), and which type is delivered depends on a 4-tier 'crowd anonymity' threshold tied to ad/app/source-identifier volume.	direct_record	SUPPORTED	high	D-S11	
D-F18	Apple's AdServices framework / Apple Ads Attribution API is scoped only to 'app-download campaigns that originate from the App Store on iOS devices' -- i.e. Apple Search Ads campaigns specifically -- and returns an attribution token with a 24-hour time-to-live that a developer or MMP must exchange server-side (POST to api-adservices.apple.com) within that window to fetch the attribution record.	direct_record	QUALIFIED	high	D-S20 D-S21	Apple's AdServices framework Attribution API covers app-download campaigns (search results, search product pages) on iOS with a 24-hour TTL that MMPs must exchange via P
D-F19	Google's official data-driven attribution documentation states the model requires at least 200 conversions and 2,000 ad interactions within a 30-day period across supported networks to function, implying accounts below this threshold default to rules-based (non-data-driven) attribution models.	direct_record	CONTRADICTED	medium	D-S10	Google Ads data-driven attribution default for most conversions is available regardless of Google recommends (require) at least 200 conversions and 2,000 ad interactions for model performance.
D-F20	Google's Meridian (open-source Bayesian MMM) documentation states it estimates channel ROI, response curves and optimal budget allocation using a single Bayesian model with adstock/diminishing-returns transformations, supports injecting business knowledge via custom priors, and integrates geo-experiment (GeoX) results for calibration -- but the documentation makes no independent accuracy claim or benchmark against competing methods.	direct_record	SUPPORTED	medium	D-S17	
D-F21	Meta's Robyn MMM documentation states calibration accepts only point-estimate inputs (start date, end date, incremental response value) from Meta Conversion Lift or Meta GeoLift experiments via a MAPE-based calibration score, and explicitly disclaims that 'we don't guarantee that using Budget Allocator as well as any Robyn's functions will meet business expectations on predicted response's accuracy,' recommending output validation before use.	direct_record	QUALIFIED	high	D-S18	Robyn accepts calibration estimates (start date, end date, incremental response) such as Meta Conversion Lift or other channels' incremental response scored via MAPE(calibration), and guarantees of predicted accuracy.
D-F22	PyMC-Marketing's own documentation states it uses a Bayesian methodology (distinct from Robyn's traditional ML-based optimization), supports lift-test calibration on par with Robyn and Meridian, and claims 'superior results versus Meridian' on its own benchmarking, while acknowledging its added flexibility requires more technical expertise to use than simpler alternatives.	vendor_assertion	QUALIFIED	medium	D-S19	PyMC-Marketing's documentation states it uses a Bayesian (vs Robyn's ML-based) methodology, supports lift-test calibration, and run benchmark claim: 40% lower contribution error than (vendor self-benchmarking)
D-F23	Branch defines an attribution window generically as 'the specific period of time...during which eligible conversion events can be claimed by publishers or media partners,' and states windows are configured per ad network in its dashboard rather than using one universal default.	direct_record	SUPPORTED	medium	D-S12	
EF-F01	Google Ads App campaigns require at least 10 conversions per day (or 300 conversions in 30 days), each with a value greater than 0, to use Target ROAS bidding.	direct_record	QUALIFIED	high	EF-S02	Per Google Ads Help 'bidding', App campaigns require conversions every day each with value greater than 0, Target ROAS, and the events must come from Analytics for Firebase
EF-F02	Google's own best-practices documentation warns that changing an in-flight Google Ads App campaign before it registers its first 100 conversions may disrupt learning, and recommends daily budgets of at least 50x the bid for tCPI/pre-registration campaigns, 10x for tCPA install campaigns, and 15x for tCPA engagement campaigns.	direct_record	SUPPORTED	high	EF-S03	
EF-F03	TikTok's own Smart+ App Campaigns help page documents three optimization goals — mobile app installs (MAI), in-app event optimization (AEO), and value-based optimization (VBO, Android only) — but does not itself state a minimum event threshold or learning-period rule.	direct_record	SUPPORTED	high	EF-S04	

EF-F04	Meta ad sets are widely reported (via secondary sources, since Meta's own help pages could not be rendered by this research's fetch tool) to require roughly 50 optimization events per ad set within a 7-day window to exit the 'learning phase,' with some reporting a reduced ~10-event threshold specifically for install/purchase optimization.	vendor_assertion	QUALIFIED	low	EF-S18 EF-S19	Meta's help centre says learning phase once d which 'usually' happens results in the week after edit (Shops ads: 17 weeks purchases); no lower is documented there.
EF-F05	Apple Search Ads' CPA cap sets a bid ceiling across an ad group's keywords that will never exceed the advertiser's max CPT bid, but the ad group's actual CPA may still exceed the entered cap; Apple's own help page states no fixed minimum-conversion or learning-period rule before adjusting it.	direct_record	SUPPORTED	high	EF-S12	
EF-F06	AppLovin's own MAX in-app bidding FAQ describes a real-time auction where all bidding partners bid simultaneously and the highest bidder wins the impression, but does not document minimum event thresholds, a learning period, or 7-day vs 28-day optimization windows anywhere on that page.	direct_record	SUPPORTED	medium	EF-S14	
EF-F07	Unity LevelPlay's own introduction documentation describes both a legacy CPM-ordered waterfall and a real-time in-app-bidding auction, plus an 'auto-optimization' algorithm comparing network eCPM, but does not state any minimum-impressions or minimum-time threshold for that optimization to activate.	direct_record	SUPPORTED	medium	EF-S15	
EF-F08	Moloco Ads and Liftoff Accelerate both market machine-learning-driven bidding toward CPI, CPA, ROAS and predicted-LTV goals, but this research could not directly open either vendor's help-center documentation (Moloco returned HTTP 403; Liftoff pages were not fetched in full) to verify mechanics or thresholds first-hand.	vendor_assertion	QUALIFIED	low	EF-S14	Moloco Ads and Liftoff ML bidding toward de (Moloco: ROAS and C long-term value; Liftoff ROAS, pLTV'), per the neither page document thresholds or learning
EF-F09	Apple's Product Page Optimization tool uses Bayesian statistical methods and only begins reporting results once at least five first-time downloads are attributed to a test, labeling variants 'Performing Better' or 'Performing Worse' once a treatment reaches 90% confidence versus the baseline.	direct_record	SUPPORTED	high	EF-S01	
EF-F10	Apple's Custom Product Pages are a separate mechanism from Product Page Optimization: each custom page gets a unique URL (with a page ID) usable in external ad channels or a developer's own site, but Apple's documentation for custom pages does not describe an automatic, evenly-split, statistically-scored comparison against a baseline the way Product Page Optimization does.	vendor_assertion	SUPPORTED	medium	EF-S01 EF-S16	
EF-F11	Google Play's Store Listing Experiments let a developer set the minimum detectable effect and the confidence level in advanced settings, and the tool warns that increasing the confidence level 'decreases the likelihood of a false positive'; experiments auto-stop after running six months regardless of outcome.	direct_record	SUPPORTED	high	EF-S05	
EF-F12	App Store and Google Play product-page conversion-rate benchmarks (commonly cited as roughly 25-33%, with Play running about 3-5 percentage points higher than iOS for the same category) come from ASO vendor panels (e.g., Storemaven, AppTweak, SplitMetrics) with differing methodologies and date windows, not from Apple or Google themselves.	commissioned_measurement	CONTRADICTED	low	EF-S16	App-store conversion definition-dependent; 1.6% average conversion product pages (and + product pages), while publish much higher p rates; the 25-33% ran over-iOS gap could not cited source.
EF-F13	The Ninth Circuit (No. 25-2935, decided December 11, 2025) affirmed the district court's contempt finding against Apple and most of the April 30, 2025 injunction's restrictions on non-neutral steering messages and link-styling rules, but held the total ban on charging any commission for external-link purchases was overbroad and remanded for the district court to define a narrower fee tied to costs 'genuinely and reasonably necessary for its coordination of external links.'	direct_record	QUALIFIED	high	EF-S06 EF-S07 EF-S08	The Ninth Circuit (No. 2025) affirmed Apple's most restrictions in the 30, 2025 order, but re commission prohibitio a commission based c and reasonably neces coordinating external compensation for App property.
EF-F14	The U.S. Supreme Court granted certiorari in Apple Inc. v. Epic Games, Inc. (No. 25-1311) on June 30, 2026, limited to 'Question 1 presented by the petition'; as of the September 27, 2026 cutoff, merits briefing was still underway (Apple's opening brief and joint appendix filed September 14, 2026; multiple amicus briefs filed September 16-21, 2026; respondent's brief due November 13, 2026), meaning the scope of any permissible external-link commission in the US remains legally unresolved.	direct_record	QUALIFIED	high	EF-S17	The Supreme Court granted Apple v. Epic (No. 25-2026, limited to whether may rest on violating c 'spirit'; merits briefing the cutoff (respondent 2026) while the district permissible link-out co (stay denied Aug 13, 2
EF-F15	Google's own Play Console Help documentation lists a staged US rollout tied to the Epic v. Google litigation: from October 29, 2025 developers may communicate about and link to external pricing/downloads and offer non-Play billing methods; alternative-billing and external-content-links compliance programs launched December 9, 2025; third-party US Android app stores gain access to developer listings by July 22, 2026 unless a developer opts out; and developers using alternative billing or external links must begin reporting transactions and paying service fees from October 1, 2026 (extended to December 1, 2026 for external-content-link enrollees' download reporting).	direct_record	SUPPORTED	high	EF-S09	
EF-F16	Secondary reporting (not independently verified via a primary court document in this research, after two attempted fetches to law-firm and news sources returned 404) indicates the Ninth Circuit affirmed the jury verdict and injunction against Google on July 31, 2025, and that Epic and Google reached a settlement on redesigned Play Store practices around March 4, 2026, with a revised modified injunction submitted to the district court.	vendor_assertion	QUALIFIED	low	EF-S09	The Ninth Circuit affirmed Google verdict and inj 2025; Epic and Google moved to modify the i 2025; renewed Mar 4, petition dismissed Mar withdrew the renewed 2026, so the original i force under court con

EF-F17	Apple's own StoreKit External Purchase Link Entitlement developer page, as fetched in this research, describes eligibility scoped only to dating apps on the Netherlands App Store storefront (per a Rotterdam district court interim-relief order) with a 3% commission reduction — it does not describe any US external-purchase-link terms, which instead flow from the separate Epic v. Apple injunction/contempt track (EF-F13, EF-F14).	direct_record	QUALIFIED	high	EF-S11	Apple's StoreKit External Purchase Link Entitlement page covers eligibility on the Netherlands storefront with a Dutch ACM order of commission reduction. A Rotterdam court interim-relief order requires storefront apps need include external purchase link terms in App Review Guidelines.
EF-F18	AppsFlyer's own description of deferred deep linking states that when a user without the app clicks a campaign link, the vendor's combined web platform and SDK capture link parameters, route the user through app-store installation, and 'restore that context' on first app open to deliver them to the intended in-app screen.	vendor_assertion	SUPPORTED	medium	EF-S20	
EF-F19	Claims that app landing pages or web-to-app flows produce '2.8x higher conversion rates' or that native shopping apps convert '3x higher' than mobile web are repeated across vendor and aggregator blogs citing unnamed 'Google internal data,' but this research could not locate or open the underlying Google source.	vendor_assertion	WITHDRAWN	low	EF-S09	
EF-F20	OpenAI's ChatGPT Apps SDK (built on the Model Context Protocol) and an associated App Directory are described, in OpenAI's own announcement and help pages, as letting developers build discoverable apps inside ChatGPT — surfaced either when ChatGPT proactively suggests one or when a user calls it by name — but this research could not directly fetch OpenAI's pages (both returned HTTP 403), so the December 2025 App Directory launch date and mechanics rest on unverified search-result synthesis.	vendor_assertion	QUALIFIED	low	EF-S21	OpenAI launched app SDK preview build on October 6, 2025 and encouraged app submissions plus an in-app directory on December 1, 2025. A mention invoked by @mention in a conversation was proactively in-conversation described as experimental.
EF-F21	Google's own Android Developers Blog (May 21, 2026) announced that Google Play discovery is being extended into the Gemini assistant on Android and web, including an 'Ask Play' conversational overlay on store listing pages, with app/game discovery rolling out 'in the coming weeks' — this is a live, documented, platform-official AI discovery surface, distinct from the retired Google Assistant.	direct_record	QUALIFIED	high	EF-S13	On May 19, 2026 Google announced 'Ask Play' (Android Developers Blog) for app discovery in the Gemini assistant on Android and web would be 'in the coming weeks' and included an AI conversational overlay on store listing pages; rollout complete but unverified.
EF-F22	An arXiv preprint (Sept 2025, revised Sept 2026) proposing a 'path signature framework' to detect creative fatigue in ad time-series explicitly relies on synthetic panel data because 'proprietary production data could not be released,' and reports no quantified real-world effect size — illustrating that rigorous, generalizable evidence on creative fatigue's magnitude is not publicly available even in current academic work.	synthesis	QUALIFIED	medium	EF-S10	A 2025-26 arXiv preprint by author at WPP Data Science proposed a path-signature screen fatigue is evaluated on CTR trajectories (94% detected, 31% false alarm) but disclaims causal attribution; no validation, so it gives no effect size.
EF-F23	A frequently repeated claim that an internal Meta study found a '45% CTR drop after the fourth ad repetition' could not be verified in this research: the Medium post commonly cited as its source (published under a 'Analytics at Meta' byline) returned HTTP 403 to direct fetch, and no alternative primary Meta source was located.	vendor_assertion	WITHDRAWN	low	EF-S10	
EF-F24	Evidence AGAINST diversification: a vendor practitioner analysis argues that spreading acquisition budget across multiple channels early can 'dilute learning, slow optimization, and create ad operational overhead,' and that fragmenting spend before any one channel reaches its platform's minimum-signal threshold (e.g., the 100-conversion and budget-ratio rules in EF-F02) can prevent any campaign from ever exiting cold start.	vendor_assertion	QUALIFIED	low	EF-S22	A RevenueCat practitioner analysis argues that adding channels early dilutes learning, slows optimization, creates operational overhead, and that channel gets too little spend for creative testing; the platform thresholds is not verified.
EF-F25	Evidence FOR diversification (against over-concentration): a vendor analysis cites a single example where 10% of a brand's audience was reached on both linear and CTV platforms, 'accumulating nearly twice the ad frequency with diminishing returns,' framed as illustrating that concentrating spend in overlapping channels/placements can waste budget on duplicated exposure.	vendor_assertion	QUALIFIED	low	EF-S23	A vendor blog (Digital) promoting its deduplication algorithm illustrates, with a stylized example, that running linear TV can hit 10% of the audience nearly twice the frequency of a single argument for cross-channel reach; not evidence for or against.
EF-F26	MMM vendor documentation (e.g., Meta's open-source Robyn, and commercial tools like Recast) models ad-spend response with Hill-function/S-curve saturation, under which marginal ROAS (the curve's slope at current spend) is explicitly distinguished from average ROAS, and channels are described as passing through an 'accelerated,' a 'linear,' and a 'plateau' phase as spend increases.	vendor_assertion	QUALIFIED	low	EF-S03	Meta's Robyn MMM model returns with a two-parameter saturation function and curves for budget reallocation; marginal-vs-average ROAS is the three-phase curve attributed to MMM vendor; not verified.
EF-F27	Meta's help center distinguishes a manual 'A/B Testing' (Experiments) tool that randomizes a chosen variable (creative, audience, placement) from 'Dynamic Creative,' which automatically assembles and delivers combinations of uploaded assets based on the delivery algorithm's own predicted performance rather than a controlled random split — per third-party description, since Meta's own pages returned no readable body text to this research's fetch tool.	vendor_assertion	QUALIFIED	low	EF-S18 EF-S19	Meta's A/B Testing tool is separate from version to a separate version; Dynamic Creative is nobody sees both; Dynamic Creative automatically combines and reports aggregate performance and does not recommend split testing. Since Jurisdiction is unavailable for new sales promotion ad sets.
EF-F28	On Apple's own Product Page Optimization page, only Product Page Optimization (not Custom Product Pages) provides a built-in baseline comparison at a stated confidence level; Custom Product Pages are traffic Apple reports per-page analytics for, but does not automatically split evenly or score for statistical significance against a control, per Apple's own documentation.	direct_record	QUALIFIED	high	EF-S01	Apple documents built-in baseline comparison for Product Page Optimization; labels only for Product Page Optimization tests; for custom product pages, per-page impressions and conversion rates in App Store Analytics are a randomized split or control.

EF-F29	Apple's own developer page states that apps can publish up to 70 additional custom product page versions (beyond the default page), routable via unique per-page URLs, Apple Ads placements, in-app StoreKit-rendered ads, and editorial placements, and that 'developers see a 2.5 percentage point increase on average' in conversion rate when referring traffic to a custom product page versus the 1.6% average conversion rate Apple reports for default product pages (a 156% relative increase).	vendor_assertion	SUPPORTED	medium	EF-S25	
G1-F01	Braun & Schwartz (Journal of Marketing, 2025) show that ad platforms' own delivery/targeting algorithms create nonrandom user exposure within A/B test creative arms, confounding a creative's effect with the platform's targeting effect.	synthesis	QUALIFIED	medium	G1-S01	Braun & Schwartz (Journal of Marketing, 2025) document in one Meta A/B test video impressions to 96,150 weeks; web lead-form each ad reached a different audience; they argue formally that relative to targeting confounds a creative's effect with algorithmic user targeting; changing the magnitude of A/B results.
G1-F02	Divergent delivery means a platform's own creative-split test result reflects both the ad content and the platform's targeting algorithm, so the same creative test run on different platforms (or reaching different audiences via delivery optimization) is not directly comparable across platforms.	synthesis	QUALIFIED	medium	G1-S01	Braun & Schwartz advise that A/B tests suit predicting performance on the same settings, but not on as causal evidence effects for use outside channels, strategy; not the same creative test an inference from this.
G1-F03	A 21-month randomized field experiment across ~35 million Pandora listeners (9 ad-load arms) found ad-load sensitivity (effect on listening hours/days and probability of listening) roughly three times larger than a one-month version of the same experiment would show.	affiliated_measurement	QUALIFIED	medium	G1-S02	In a 21-month (June 2024) randomized experiment ad-supported Pandora ad-load arms plus a 10% manipulated in the iOS each extra audio ad per month listening hours: days by 1.91%; the heaviest listened 2.9% less than the lightest (3x) 1.8% more; elasticity was about three times one- or two-month test.
G1-F04	In the same Pandora randomized ad-load experiment, observational (non-experimental, correlational) estimates of the ad-load/engagement relationship were biased and sometimes directionally wrong compared to the randomized estimates.	affiliated_measurement	SUPPORTED	medium	G1-S02	
G1-F05	In the same experiment, higher randomized ad-load assignment was associated with more paid (ad-free) subscription conversions, alongside reduced listening.	affiliated_measurement	QUALIFIED	low	G1-S02	In the Pandora experiment randomized higher ad-load increased ad-free subscription conversions, alongside reduced listening. Additional ad per hour increased probability of being a subscriber by 0.14 percentage points (0.75 cents of monthly revenue per listener) and probability of not listening by 0.34 points; listeners less likely as 13-24s to upgrade; authors were not permitted revenue trade-off.
G1-F06	A micro-randomized trial (Bidargaddi et al., JMIR mHealth 2018) randomized the timing of push notifications in a workplace mobile-health app to isolate their proximal (near-term) effect on engagement, distinct from confounded pre/post comparisons.	independent_measurement	QUALIFIED	low	G1-S03	A micro-randomized trial JMIR mHealth uHealth users of a commercial being app over 89 days; one randomly selected whether to send a tailored notification; sending reduced in-app self-monitoring by 3.9% relative (risk ratio 1.01-1.08), with larger effect on weekends.
G1-F07	In a single freemium mobile RPG (Age of Ishtaria, ~2,505 churned paying users in the evaluated cohort), a CNN model predicting 365-day-forward LTV had a 5.72% mean percent error across all paying users versus 8.96% for the Pareto/NBD benchmark model, on a held-out test set.	vendor_assertion	QUALIFIED	medium	G1-S04	In one freemium mobile game Ishtaria; 2,505 paying users May 2016 to May 2017; CNN had a max-norm of 5.72% vs 8.96% (Pareto/NBD average spend) or 9.0% (SMAPE); on SMAPE the CNN was 73.76% vs 95.65%.
G1-F08	In the same study, prediction error for the top 20% of spenders ('top spenders') was 15.64% for the CNN model versus 33.35% for the Pareto/NBD model — parametric BTYD models systematically underestimate top-spender ('whale') future spending, roughly double the error of the deep-learning models on this subgroup.	vendor_assertion	QUALIFIED	medium	G1-S04	For the 20% of paying users most in the year, max-norm error was 15.64% for the CNN model vs 33.39% for Pareto/NBD; BTYD models systematically underestimate top spenders; generate up to 50% of revenue (single game, ~500-600 users).

G1-F09	Two more recent (2024-2025) arXiv preprints report large relative LTV-prediction-error reductions in industrial mobile-game deployments (a Monte Carlo Dropout uncertainty model improving top-5%-spender MAPE; a since-withdrawn 'SHORE' long-term LTV model claiming a 47.91% relative error reduction), both citing delayed payment behavior, data sparsity, and high-value outliers as the core failure modes of standard LTV models.	vendor_assertion	QUALIFIED	low	G1-S07 G1-S08	Two recent non-peer-preprints claim LTV ga Cao, Xu and Yang (Tel about 3 million new us predict next-month sp of play; Monte Carlo D MAPE from 0.48 to 0.1 0.20 (DCNv2), with G SHORE (2025) claims error reduction and ne payments, sparse earl value outliers, but arX removed it because th the right to agree to th is available.
G1-F10	Retargeting DSP Jampp reports (vendor blog, Aug 2026) that its always-on 'Ghost Bids' holdout methodology measured an 86% increase in new riders and 122% increase in rides for ride-hailing app FREENOW, and a 93% growth in listers and 92% lift in conversions for marketplace app Wallapop, versus its ~6%-of-users control group.	vendor_assertion	QUALIFIED	low	G1-S06	Jampp's blog (Aug 20 period, baseline or un FREENOW saw an '86 riders', a '122% increa '30% campaign lift', a growth in listers' and i conversions', under its Bids lift measurement, recommends a ~6% n group.
G1-F11	Retargeting DSP Rmerge markets three incrementality-measurement methods for app re-engagement campaigns — continuous 'Ghost Bids' holdout, Intent-to-Treat (ITT) holdout, and a distribution-matching validation check — but its published case studies (Miniclip 8 Ball Pool, Socialpoint Dragon City, PhotoSI, Delivery Hero HungerStation) describe qualitative outcomes without disclosing quantified lift percentages on the page reviewed.	vendor_assertion	QUALIFIED	low	G1-S05	Rmerge's incrementa Ghost Bids (always-or and Causal Impact (ID econometric, for insta engagement), with ITT in linked content; its fe (Miniclip retargeting, v Miniclip's own distribu Socialpoint, PhotoSI a Hero/HungerStation u no quantified lift figure
G1-F12	No app-install-specific or app-usage-specific randomized or credible quasi-experimental study of CTV or TV advertising effects could be located in this research; the closest available high-quality causal TV advertising evidence (Shapiro, Hitsch & Tuchman, Econometrica 2021, 288 CPG brands) measures brand sales elasticity/ROI, not app outcomes.	inference	QUALIFIED	high	G1-S09	No randomized or qua study of TV/CTV effec usage was found in th available; treat as an e medium confidence. T TV evidence located, Tuchman (Econometri brands), measures br and ROI, not app outc
G2-F01	AppLovin reported Q2 2026 (three months ended June 30, 2026) revenue of \$1,923.7 million, up 53% year-over-year from \$1,258.8 million in Q2 2025.	direct_record	SUPPORTED	high	G2-S01 G2-S02	
G2-F02	AppLovin's Q2 2026 net income was \$1,266.5 million and adjusted EBITDA was \$1,613.8 million, an 83.9% adjusted EBITDA margin (non-GAAP).	direct_record	SUPPORTED	high	G2-S01 G2-S02	
G2-F03	AppLovin's revenue-recognition accounting policy states the company acts as agent, not principal, for its advertising business and presents revenue net of advertising-inventory costs paid to publishers.	direct_record	SUPPORTED	high	G2-S01	
G2-F04	No primary source accessed in this pass (AppLovin's FY2025 10-K, Q1 2026 10-Q, Q2 2026 10-Q, or the Q1/Q2 2026 earnings releases) confirms that AppLovin completed a 'public opening' of Axon Ads Manager to e-commerce or web advertisers in June 2026; 'Axon Ads Manager' is used only as the brand name for AppLovin's core ad-buying platform, and e-commerce is mentioned solely as a named emerging vertical for future expansion.	inference	QUALIFIED	medium	G2-S01 G2-S02 G2-S03 G2-S04 G2-S05	AppLovin's filings do 2026 'public opening' commerce/web adver since 2024 it has been based e-commerce ad solutions have been r web-based advertiser market expansion', FY self-serve platform (A launched in 2025, and Q, the core product is Ads'.
G2-F05	As of June 30, 2026, AppLovin's 10-Q states it had no material legal-proceedings loss contingencies requiring accrual or disclosure, and does not name any specific SEC, DOJ, FTC investigation or securities/class-action lawsuit in that filing's text.	direct_record	CONTRADICTED	high	G2-S01	AppLovin's Q2 2026 1 (Commitments and Co Proceedings), states: and December 31, 20; no material loss contie legal proceedings for disclosure was require nonetheless discloses securities class action N.D. Cal.; motion to di of February 2026) and shareholder derivative pending that motion; if government investigat
G2-F06	Unity reported Q2 2026 total revenue of \$546.5 million and Grow Solutions revenue of \$388.9 million (up 35.5% year-over-year from \$287.2 million), which the company attributes to growth in the Unity Ads Network 'driven by Unity Vector', partially offset by declines in the ironSource Ads Network.	direct_record	QUALIFIED	high	G2-S06 G2-S07	Unity reported Q2 20 \$546.5M and Grow Sc \$388.9M, up 35% (iss unrounded) from \$28; Unity Ads Network gr Vector', partly offset b Network.
G2-F07	Unity CEO Matt Bromberg described 'the ongoing success of Unity Vector AI' as helping drive value for creators, players and shareholders in the Q2 2026 earnings release, without disclosing a standalone revenue or performance metric for Vector.	vendor_assertion	SUPPORTED	medium	G2-S07	

G2-F08	Liftoff Mobile filed its first 10-Q as a public company for the quarter ended June 30, 2026 (following an IPO completed June 5, 2026 at \$23.00/share, raising approximately \$472.4 million in net proceeds from 21.9 million shares, with \$409.2 million applied to debt repayment).	direct_record	SUPPORTED	high	G2-S09	
G2-F09	Liftoff reported Q2 2026 revenue of \$219.5 million, up 35.4% year-over-year from \$162.1 million in Q2 2025.	direct_record	SUPPORTED	high	G2-S09	
G2-F10	Liftoff's Q2 2026 10-Q discloses that one customer accounted for approximately 10% of revenue for both the three and six months ended June 30, 2026, with no customer exceeding 10% concentration in the prior-year periods; the filing does not disclose a gaming vs non-gaming revenue split.	direct_record	SUPPORTED	high	G2-S09	
G2-F11	Liftoff's revenue-recognition policy states it acts as agent (not principal) because it does not control ad inventory, and reports revenue net of 'consideration payable to publishers or third parties.'	direct_record	SUPPORTED	high	G2-S09	
G2-F12	Digital Turbine reported fiscal-year 2026 (ended March 31, 2026) total revenue of \$565.3 million (+15% YoY), with App Growth Platform segment revenue of \$185.7 million (+21% YoY) and On Device Solutions revenue of \$382.4 million (+12% YoY); adjusted EBITDA was \$122.5 million (+69% YoY).	direct_record	SUPPORTED	high	G2-S10 G2-S13	
G2-F13	Digital Turbine's revenue-recognition policy is split by business line: it recognizes App Growth Platform Marketplace (exchange) revenue net, stating it 'is an agent in transactions on its Marketplace platforms' and acts as 'an intermediary between DSPs and publishers,' while recognizing Brand and Performance offerings gross, stating it 'is a principal' and pays publishers' shares 'as a cost of revenue.'	direct_record	SUPPORTED	high	G2-S11 G2-S12	
G2-F14	Unity's Grow Solutions revenue is described in its Q2 2026 10-Q as 'the amount we retain from the transaction we are facilitating through our auction and mediation platform,' language consistent with a net/agent presentation, but the formal principal-vs-agent accounting-policy paragraph (as found for AppLovin, Liftoff and Digital Turbine) was not located within the portions of Unity's 10-Q or FY2025 10-K retrieved in this pass.	direct_record	QUALIFIED	medium	G2-S06 G2-S08	Unity's FY2025 10-K s Grow Solutions advert net basis where it fac between advertisers a without controlling the publisher being its cus gross basis where it is Q2 2026 10-Q describ 'the amount we retain transaction'.
G2-F15	Mobvista reported FY2025 (calendar year) total revenue of approximately \$2.05 billion, up 35.7% year-over-year, driven by Mintegral; the company's own press release does not state whether this revenue is recognized gross or net of publisher/traffic-acquisition payouts, and the formal accounting-policy passage from Mobvista's HKEX annual report could not be retrieved in this pass.	direct_record	QUALIFIED	medium	G2-S14	Mobvista reported FY about \$2.05 billion (+ gross ('generally the g Ad-tech'); its non-IFRS revenue after traffic-p was about \$518.6 milli \$245.1M + H2 \$273.5M headline is not compa reporting peers such a Grow or Liftoff.
G2-F16	Stripe's published US online card-processing rate is 2.9% + \$0.30 per successful domestic-card transaction, with an additional 1.5% for international cards and 1% for currency conversion.	direct_record	SUPPORTED	high	G2-S15	
G2-F17	RevenueCat's State of Subscription Apps 2026 report finds that web-revenue adoption scales sharply with app performance tier: 41% of its top-performing ('Tier 5') apps generate web revenue versus 1.3% of 'Tier 1' hobby apps, and web revenue represents 3.2% of total revenue globally but 4.9% in North America and 0.8% in India/SEA, based on a panel of over 115,000 apps and more than \$16 billion in tracked revenue on RevenueCat's own platform.	vendor_panel	SUPPORTED	medium	G2-S16	
G2-F18	No maintained, non-paywalled public registry or count of live SKAdNetwork/AdAttributionKit-registered ad-network identifiers was accessible in this pass: Apple's ad-network-list page (developer.apple.com/app-store/ad-network-list/) returned HTTP 404, its SKAdNetwork documentation landing page returned only a page title (JavaScript-rendered body not retrievable via WebFetch), and an attempted MMP partner-directory page (AppsFlyer's SKAdNetwork marketplace) also returned HTTP 404.	inference	SUPPORTED	high	G2-S17	
G2-F19	No public, non-paywalled, quantified estimate of US or global in-app advertising revenue (money quantity Q2) covering a Q2-2026-relevant period was found in this pass; EMARKETER's dedicated in-app mobile ad spending page returned HTTP 403 (paywalled), consistent with the AB stream's independent finding (AB-F27) that this series requires a subscription.	inference	SUPPORTED	high	G2-S18	
G3-F01	In Johnson, Lewis & Nubbemeyer's (2017, JMR) 'ghost ads' field experiment with an online retailer, retargeting display ads lifted website visits by 17.2% and purchases by 10.5% relative to the ghost-ad counterfactual.	affiliated_measurement	QUALIFIED	medium	G3-S01	In Johnson, Lewis & N 2017) 'ghost ads' exp Google's display ad sy unnamed online appar retailer's two-week de campaign on the Goo retargeting ads lifted v 17.2% and purchases predicted-ghost-ad co (published abstract). A Google employee whe done, so this is platfor independent of the ad
G3-F02	Sahni, Narayanan & Kalyanam (2019, JMR) find that switching on experimental retargeting causes 14.6% more users to return to an online home-improvement retailer's website within four weeks, with the effect decaying over time (33% of week 1's effect occurring on day 1) and higher week-2 effects when week-1 advertising was nonzero.	independent_measurement	SUPPORTED	medium	G3-S02	
G3-F03	Simonov, Nosko & Rao (2018, Marketing Science) find brand-keyword search ads have a statistically significant positive effect of 1%-4% on traffic when no competitor is present, but when a brand with contested keywords stops bidding, competitors in paid positions 2-4 'steal' 18%-42% of the focal brand's clicks (versus 1%-5% stolen when the brand's own ad is still present).	affiliated_measurement	QUALIFIED	medium	G3-S03	Simonov, Nosko & Ra 2018; randomized exp fraction of US-located nine days in January 2 while two authors wer Research) find brand- focal-brand clicks by competitor is present; is present, competitor take 1%-5% of its clic that face competition advertise on their bran take 18%-42% of clic authors note may refle

G3-F04	A 2026 quasi-experimental (regression-discontinuity) study using push-notification-provider data finds that delivering a notification immediately upon a user's screen-on event reduces same-day app logins by approximately 16% versus a short 1-6 minute delay, with peak engagement at that modest delay and a similar pattern in app uninstalls.	independent_measurement	QUALIFIED	medium	G3-S04	A 2026 SSRN working (Zhong) using data from notification provider finds discontinuity, that delivered immediately at screen day app logins by about instrumental-variable engagement peaks at with suggestive consistency app uninstalls and a similar app.
G3-F05	On February 26, 2025, Culper Research published a report titled 'AppLovin Corporation (NASDAQ: APP): Force-Feeding Users with Silent Backdoor Installs and Copying Meta's Homework. Straight to the Principal's Office, Please,' alleging (among other things) a 'Backdoor Installation Scheme' in which AppLovin's MAX SDK updates allegedly smuggled a permission into partner apps enabling silent, one-click direct app installs outside app stores.	vendor_assertion	SUPPORTED	medium	G3-S07	
G3-F06	Also on February 26, 2025, Fuzzy Panda Research published a report titled 'AppLovin (APP) - Formers Allege Ad Fraud; Is DTC Hype Actually 'Stealing' Meta's Data; Illegal Tracking of Children & Serving Sex Ads to Kids,' alleging AppLovin used its Meta-mediation SDK access to 'peek' at Meta ad performance data and that the company was 'caught using fake clicks [and] other dirty tricks to game installs.'	vendor_assertion	SUPPORTED	medium	G3-S07	
G3-F07	On February 20, 2025, Edwin Dorsey (author of The Bear Cave, who the underlying complaint itself notes 'is not a short seller' and 'does not take positions against companies profiled') published a report titled 'Problems at AppLovin (APP)' raising concerns about ad-quality, dependency on required Meta ad spend from advertisers, and lack of transparency.	vendor_assertion	SUPPORTED	medium	G3-S07	
G3-F08	Muddy Waters Research's March 27, 2025 report claims, based on 'conversion log-level files provided by a leading independent demand-side platform' covering 37 million+ users across five advertisers, that over 50% of AppLovin's e-commerce conversions were attributable to retargeting and only 25%-35% were incremental, alongside an estimated ~23% Q1 2025 advertiser churn rate.	vendor_assertion	QUALIFIED	medium	G3-S05 G3-S07	Muddy Waters (disclosed 27 Mar 2025) estimates AppLovin's e-commerce retargeting and that 'it ~25%-35%', citing (per complaint) conversion an independent DSP covering across five advertisers; ~23% Q1 2025 churn; e-commerce advertisers of AppLovin's pixel from are a short seller's estimate undisclosed method, and measurements.
G3-F09	Muddy Waters Research's May 7, 2025 follow-up report alleges that AppLovin CEO Adam Foroughi's March 31, 2025 statement denying that AppLovin 'creates or uses persistent identifiers' was false, and that CTO Basil Shikin's blog-post explanation was 'a lie of omission,' asserting the company continued using persistent identifiers that violate partner-platform terms of service.	vendor_assertion	SUPPORTED	medium	G3-S06	
G3-F10	AppLovin publicly responded to the short-seller reports: CEO Adam Foroughi published a blog post titled 'A Note from Our CEO: Discussing Web Advertising Opportunity and Unpacking Pixels' within hours of the March 27, 2025 Muddy Waters report, and CTO Basil Shikin published a further blog post on March 31, 2025 defending the company's e-commerce data practices; both posts incorporated AI-generated analysis from xAI's Grok 3 chatbot.	direct_record	QUALIFIED	medium	G3-S07	AppLovin publicly reb 26 Feb 2025 'Note from complaint); CEO Adam from Our CEO: Discuss Advertising Opportunity Pixels', issued hours after Muddy Waters report, which includes a section Grok3, an AI model by Shikin's 'Examination practices' (31 Mar 2025 complaint), which says its drafting; and Foroughi Advertising: How we can handle data', which M 31 Mar 2025.
G3-F11	AppLovin's Class A stock fell 8.9% on February 20, 2025 (Bear Cave Report), fell approximately 12.2% on February 26, 2025 (Culper/Fuzzy Panda Reports, from \$377.06 to \$331.00), and plunged 20.1% on March 27, 2025 (Muddy Waters Report, from \$327.62 to \$261.70), described in the complaint as the stock's 'steepest drop on record.'	direct_record	SUPPORTED	high	G3-S07	
G3-F12	Five related AppLovin lawsuits were filed in the N.D. California federal court in 2025: three putative securities class actions (Quiero, 4:25-cv-02294, filed 2025-03-05, voluntarily dismissed May 2025; Brownback, 4:25-cv-02772, filed 2025-03-24, the surviving lead action; Wayne County Employees' Retirement System, 4:25-cv-03438, filed 2025-04-17, voluntarily dismissed May 2025) and two shareholder derivative suits (Patel v. Foroughi, 4:25-cv-02780, filed 2025-03-25; Smith v. Foroughi, 4:25-cv-04261, filed 2025-05-19), the latter two consolidated and stayed pending the securities case.	direct_record	QUALIFIED	high	G3-S11 G3-S10	In 2025 five related securities Cal.: securities class action cv-02294, filed 5 Mar 2025, dismissed, closed 27 Mar 2025; Brownback (4:25-cv-02772, 2025; the lead action); ERS (4:25-cv-03438, voluntarily dismissed, 2025), and derivative Foroughi (4:25-cv-02780, 2025 per the docket); Foroughi (4:25-cv-04261, 2025), consolidated as of 5 Jun 2025. A separate action, Talbot v. AppLovin, filed 16 Sep 2026; class action (Aug 2026), concerns products, not the 2025 allegations.

G3-F13	As of AppLovin's Q2 2026 10-Q (filed 2026-08-05), the lead securities class action (Brownback, consolidated with an Amended Complaint filed 2025-09-12 naming Basil Shikin as an additional defendant and alleging a class period of November 7, 2024 to March 27, 2025) had a motion to dismiss that was 'fully briefed as of February 2026,' with no ruling reported in that filing; the related shareholder derivative suits remain consolidated and stayed pending that motion's resolution.	direct_record	QUALIFIED	high	G3-S10	As of AppLovin's Q2 2026 (filed Aug 2026), the motion to dismiss (Brownback Amended Complaint, Sep 2025; class period Mar 2025; Shikin added as a defendant) has been fully briefed since Feb 2026. The docket shows the submission on 14 Apr 2026, together with a motion to supplement the SEC and state derivative suits. No ruling appears on the docket (last refreshed Aug 2026). AppLovin's filings through the derivative suits remain consolidated and stayed.
G3-F14	None of AppLovin's three most recent SEC filings accessible at cutoff (10-Q filed 2025-11-05 for Q3 2025, 10-K filed 2026-02-19 for FY2025, and 10-Q filed 2026-08-05 for Q2 2026 — the last filed roughly ten months after Bloomberg's reported October 2025 SEC-inquiry story) discloses, names, or otherwise confirms any SEC subpoena, inquiry, or investigation into AppLovin; only generic risk-factor boilerplate about 'litigation or regulatory action ... or investigations by regulators' appears.	direct_record	QUALIFIED	medium	G3-S08 G3-S09 G3-S10	None of AppLovin's filings from Nov 2025 to Aug 2026, FY2025 10-K, Q1 2026, or any SEC, DOJ, attorney-general investigation or SEC proceedings sections, that the company is civil or criminal legal proceedings and governmental and other investigations and probes (Bloomberg reported on Oct 2025 via Reuters) that the SEC subpoenaed AppLovin's data-collection. AppLovin said it does not have any potential regulatory matters. AppLovin plaintiffs in Brownback v. AppLovin reported SEC and state derivative suits to their complaint. The suits are not confirmed by Bloomberg.
G3-F15	Article 6(8) of the EU Digital Markets Act requires a designated gatekeeper to give advertisers, publishers, and their authorized third parties free access, on request, to the gatekeeper's own performance-measuring tools and to the aggregated and non-aggregated data necessary to independently verify advertisement inventory and measure performance.	direct_record	SUPPORTED	high	G3-S12	
G3-F16	Article 6(10) of the EU Digital Markets Act requires a designated gatekeeper to give business users and their authorized third parties, on request and free of charge, 'effective, high-quality, continuous and real-time access to, and use of' aggregated and non-aggregated data (including personal data under opt-in consent conditions) generated in the context of business users' and end users' use of the gatekeeper's core platform services.	direct_record	SUPPORTED	high	G3-S12	
G3-F17	Apple's official developer policy states developers 'may not derive data from a device for the purpose of uniquely identifying it' (covering browser/device configuration, location, or network-connection signals), and that apps or referenced SDKs (including ad networks, attribution services, and analytics SDKs) found engaging in fingerprinting 'may be rejected from the App Store.'	direct_record	SUPPORTED	high	G3-S13	
G3-F18	AppsFlyer's own attribution-methods documentation lists 'Probabilistic modeling' as one of its supported attribution techniques on iOS (alongside Android, and CTV/PC/console, but not UWP), explicitly labeled as using no device ID and a 'Probabilistic' (not deterministic) technique, and states that in its attribution-priority waterfall it 'prioritizes clicks over impressions, and deterministic over probabilistic methods.'	vendor_assertion	QUALIFIED	high	G3-S14	AppsFlyer's attribution documentation (edited 12 Aug 2026) lists 'Probabilistic modeling' (no device ID) as support and CTV/PC/console as support and that it 'prioritizes clicks and deterministic over probabilistic methods', and describes probabilistic modeling as a statistical method that produces aggregate data reports rather than identifying individual devices or users (vendor assertion).
G3-F19	Adjust's official documentation states that SKAdNetwork attribution 'is always deterministic' and 'attributes regardless of the user's ATT status,' whereas ATT-framework (non-SKAN) attribution 'can happen with probabilistic modeling when allowed by the ATT policy' and 'is limited for users who have not consented per the ATT policy.'	vendor_assertion	SUPPORTED	high	G3-S15	
G3-F20	AdCP's documented governance protocol enforces a three-party separation of duties (an 'orchestrator' that proposes/executes buys but cannot set its own spending limits or approve its own plans; a 'governance agent' that validates plans against policy and tracks budgets but cannot execute buys; and a 'seller' that fulfills buys but cannot override governance decisions), with a check_governance API call returning structured findings at 'must,' 'should,' or 'may' severity, and transactions exceeding a human-set per-transaction authority threshold are held for asynchronous human approval rather than executed automatically.	direct_record	QUALIFIED	medium	G3-S16	AdCP's campaign governance protocol documented at v3.1.24 (released at cutoff; marked as 'Request for Comments') and 'Request for Comments' duties among an orchestrator (its own spending limits and approvals), a governance agent (execute buys) and a seller (override governance decisions). The check_governance API call returns must/should/may severity findings above a human-set authority threshold. Escalated asynchronous human approval. Governance agent must call check_governance API. Buyer has registered campaign on the account, and must call all buyers.

G3-F21	IAB Tech Lab's AAMP 2.0 (announced 2026-04-23) provides 'transaction-ready' Buyer and Seller Agent SDKs with documented 'human-in-the-loop approval gates and guardrails' and 'configurable approval gates' on the Seller Agent side, described as ensuring 'automation remains controlled and reliable, preventing errors and maintaining alignment with business objectives'; as of the article's last-modified date (2026-08-28), AAMP 3.0 was described as 'well underway,' extending the framework into more transaction and media types.	direct_record	QUALIFIED	high	G3-S17	IAB Tech Lab's AAMP (Apr 2026) released 'transaction-ready' Buyer and Seller Agent SDKs. The Seller Agent is described with 'human-in-the-loop approval gates' and 'configurable approval gates'. On 22 Aug 2026, IAB Tech Lab announced AAMP 3.0 in an OpenProposal specific to the EU, currently in proposal stage, which comment until 22 Oct 2026.
G3-F22	On July 8, 2025, the Eighth Circuit (Custom Communications, Inc. v. FTC, Nos. 24-3137, 24-3388, 24-3415, 24-3442, consolidated, per curiam) granted the petitions for review and vacated the FTC's negative-option ('click-to-cancel') Rule in its entirety, holding the Commission 'failed to follow procedural requirements under Section 22 of the FTC Act' by not conducting a required preliminary regulatory analysis, without reaching petitioners' other substantive challenges.	direct_record	SUPPORTED	high	G3-S18	
G3-F23	Following the Eighth Circuit's vacatur, the FTC's own 'Negative Option Rule' page shows the agency published a Federal Register notice on 'Revision of the Negative Option Rule' (2026-02-12) and an Advance Notice of Proposed Rulemaking seeking public comment on potential amendments (published 2026-03-13; comment period opened via a 2026-03-11 notice), i.e., at the 2026-09-27 cutoff the original 2024 click-to-cancel Rule remains vacated and the FTC is in an early-stage re-rulemaking process rather than having a rule currently in effect.	direct_record	QUALIFIED	high	G3-S19 G3-S18	After the Eighth Circuit's decision, the FTC updated its 'Negative Option Rule' page on 12 Feb 2026 re: 'Revision of the Negative Option Rule'. That narrower 'Use of Prenotification Plans' (16 CFR 425) remains in effect. On 13 Mar 2026 the FTC published a notice seeking comment on a proposed rule (comments closed 13 Sep 2026) no proposed rule and no federal click-to-cancel effect.
G3-F24	Google's Ads API v25.1 (released 2026-08-19) added a new read-only LiftMeasurementConfig resource, 24 new Conversion Lift metrics and winner-score statistical metrics, and an Experiment.Lift_measurement_config field, i.e., Conversion Lift study results became newly accessible via the Ads API rather than only through the UI; the same release also added comparable Brand Lift measurement API resources.	direct_record	QUALIFIED	high	G3-S20	Google Ads API v25.1 read-only reporting resources (LiftMeasurementConfig, LiftMeasurementFlight, LiftMeasurementWinnerScore, optional Experiment.Lift_measurement_config) as well as Brand Lift data resources, so lift-study results retrieved through the API, not describe creating any change to eligibility.
GH-F01	Apple's standard App Store commission on auto-renewable subscriptions is 70% net revenue to the developer during a subscriber's first year of paid service and 85% after one year of accumulated paid service (free trials excluded from the one-year count).	direct_record	SUPPORTED	high	GH-S02	
GH-F02	Apple's App Store Small Business Program gives a flat 15% commission (vs. the standard 30%/70-85% tiered structure) to developers with up to \$1,000,000 USD in prior-year proceeds, with re-qualification the year after falling below the threshold.	direct_record	QUALIFIED	high	GH-S01	Apple's App Store Small Business Program charges 15% to enroll in the program. Proceeds are calculated as Developer Accounts, not net revenue, and so far in the current year developers can also qualify for the program in the current year regardless of future sales. Developers can fall back below the threshold the following year.
GH-F03	Under Apple's Alternative Terms for Apps in the EU (effective 2026-10-01), commission tiers are 26%/15% (Apple in-app purchase), 20%/10% (alternative payment processing in-app), and 15%/10% (out-of-app link-outs within a 7-day attribution window), with a new 5% Core Technology Commission per transaction replacing the prior per-install Core Technology Fee for apps distributed via alternative marketplaces or the web.	direct_record	QUALIFIED	high	GH-S03	Apple's unified EU business model was updated 18 Aug 2026; the 27 Sep 2026 cutoff of 26%/15% for Apple in-app alternative payment processing (sales with tap), and replace the prior Core Technology Fee with a new Core Technology Commission for apps distributed outside the EU.
GH-F04	Google Play's baseline service fee (in markets not yet migrated to the 2026 restructuring) remains 15% on the first \$1M USD of a developer's annual earnings and 30% above that threshold for non-subscription digital goods, while auto-renewing subscriptions are charged a flat 15% regardless of annual revenue tier.	direct_record	QUALIFIED	high	GH-S04 GH-S05	In markets not yet migrated to the 2026 fee model, developers pay 15% service-fee tier plus 30% on earnings above US\$1M earned each year. Auto-renewing subscriptions are charged 15% regardless of earnings. The EEA moved on 30 Jun 2026, Korea on 31 Dec 2026, and Japan on 31 Dec 2026, with global completion of the world on 30 Sep 2026.
GH-F05	Google Play is replacing its flat commission with a split 'service fee' + 5% 'billing fee' structure following the Epic v. Google settlement: new-install recurring transactions and subscriptions are charged 10% service fee + 5% billing fee; non-recurring new-install transactions are 20% (or 10-15% under program enrollment) + 5% billing; existing-install non-recurring transactions run 25% (or 15-20%) + 5% billing. Rollout is 2026-06-30 for US/UK/EEA, then AU/Japan, then South Korea, with global completion by 2027-09-30.	direct_record	QUALIFIED	high	GH-S06	Google Play's new fee structure was announced 18 Aug 2026. US, UK and EEA on 30 Sep 2026, Australia and Japan on 31 Dec 2026, Korea on 31 Dec 2026, and the rest of the world on 30 Sep 2026. New transactions pay 10% service fee + 5% billing. Existing transactions pay 20% on new installs, or 25% on existing installs, or 20% on web link. A developer pays 10% + 5% billing. Google Play Billing transactions in the US, UK and EEA. Program 20% existing) are published programs open only on Google Play.

GH-F06	As of the Ninth Circuit's December 11, 2025 ruling, Apple is barred from charging any commission on US App Store external purchase links until a district court approves a specific 'reasonable' fee reflecting only coordination costs and some IP compensation (explicitly excluding security/privacy costs); the original injunction's blanket lifetime ban on commissions was found to be an abuse of discretion, and the matter was remanded.	synthesis	QUALIFIED	low	GH-S18	The Ninth Circuit (11 D Apple's civil contempt court abused its discretion by banning all link-out commissions, remanded, recommending charge none until the court based on necessary compensation plus some IP compensation (security/privacy). As a result, the fee had been approved in prior proceedings were on remand. Supreme Court had granted certiorari (30 Jun 2026).
GH-F07	RevenueCat's State of Subscription Apps 2026 report is built from a panel of 115,000+ apps using RevenueCat's subscription-billing infrastructure, representing \$16 billion+ in tracked revenue and 1 billion+ transactions, primarily reflecting 2025 activity.	vendor_assertion	SUPPORTED	medium	GH-S07	
GH-F08	In RevenueCat's 2026 panel, apps using a hard paywall (no free access) show a 10.7% median Day-35 download-to-paid conversion rate vs. 2.1% median for freemium apps, roughly a 5x difference.	vendor_assertion	SUPPORTED	medium	GH-S08	
GH-F09	In RevenueCat's 2026 panel, longer free trials convert better: trials of 17-32 days show a 42.5% median trial-to-paid conversion rate vs. 25.5% for trials under 4 days (~70% relative difference).	vendor_assertion	SUPPORTED	medium	GH-S08	
GH-F10	RevenueCat's 2026 panel shows Year-1 annual-subscription cancellation reaching approximately 72% (worsening from ~56% in the 2025 report), with 35% of annual cancellations occurring in Month 1 and 55.4% of 3-day-trial cancellations occurring on Day 0 (up from 51% in 2025).	vendor_assertion	QUALIFIED	medium	GH-S08	RevenueCat's 2026 report shows Year-1 retention of 28% (2024 cohort; 31% for 2025 cohort). The report restates this as ~30% retention, which is consistent with subscribers cancelling in the prior report, compared to the report's different basis for 35% of annual cancellations occurring on 3-day-trial cancellations.
GH-F11	RevenueCat's 2026 panel attributes 31% of Google Play subscription cancellations to involuntary churn (billing/payment failures), up from 28.2% in 2025, vs. 14% on the App Store -- a persistent platform gap in payment-failure-driven churn.	vendor_assertion	QUALIFIED	medium	GH-S08	RevenueCat's 2026 report shows 32.2% of Google Play Store subscription cancellations attributed to payment errors (its summary states 31%).
GH-F12	RevenueCat's 2026 panel reports apps categorized as 'AI apps' show 41% higher median Year-1 realized LTV (\$30.16) than non-AI apps (\$21.37), while AI-app monthly-plan 12-month retention is 36% worse than traditional (non-AI) apps.	vendor_assertion	SUPPORTED	medium	GH-S08	
GH-F13	Adaptly's 2026 panel (16,000+ apps, \$3B+ subscription revenue, 10,000+ paywalls compared) reports a global average install-to-trial conversion of 10.9% and trial-to-paid conversion of 25.6%, with Health & Fitness apps highest among named categories at 35.0% trial-to-paid and Entertainment lowest at 19.1%.	vendor_assertion	QUALIFIED	medium	GH-S09	Adaptly's 2026 report shows global average install-to-trial conversion of 10.9% and trial-to-paid conversion of 25.6% (Health & Fitness 35.0% and Entertainment 19.1%).
GH-F14	AppsFlyer's State of App Monetization 2026 Edition (Jan 2025-Mar 2026, verified revenue observed via AppsFlyer's platform) reports \$900M in verified IAP revenue, \$800M in verified in-app-subscription (IAS) revenue, and \$7.2B in verified in-app-advertising (IAA) revenue across its measured base; by Day 7, in-app-advertising revenue reaches 89% of its eventual Day-60 total.	vendor_assertion	QUALIFIED	medium	GH-S10	AppsFlyer's State of App Monetization 2026 (Jan 2025-Mar 2026) reports \$900M verified IAP revenue, \$800M verified IAS revenue, and \$7.2B verified IAA revenue. IAA reaches 89% of its eventual Day-60 total.
GH-F15	AppsFlyer's 2026 monetization panel reports non-gaming one-time-buyer conversion of 9.84% and repeat-buyer (2+ purchases) conversion of 4.64% within 30 days of install.	vendor_assertion	SUPPORTED	medium	GH-S10	
GH-F16	GameAnalytics' 2026 benchmark report (16,000+ mobile games, min. 1,000 MAU, 9 regions, iOS+Android, Jan-Dec 2025) shows median mobile D1 retention ~22%, D7 just under 4%, and D30 ~0.68-0.79%, with extreme right-skew: P99 games hold 64-68% D1 and 13-15% D30.	vendor_assertion	QUALIFIED	medium	GH-S11	GameAnalytics' 2026 benchmark report (16,000+ games using MAU, iOS+Android) shows median D1 ~22%, D7 just under 4%, and D30 ~13-15%.
GH-F17	GameAnalytics' 2026 report shows iOS mobile games retain better at the top end than Android: the top 25% of iOS titles show 31-33% D1 retention vs. 25-27% for the top 25% of Android titles.	vendor_assertion	WITHDRAWN	medium	GH-S11	
GH-F18	Google's ZILN (zero-inflated lognormal) loss function for LTV prediction, published by Wang, Liu and Miao (2019), was validated by its authors' open-source companion repository only on two non-mobile public datasets: the Kaggle 'Acquire Valued Shoppers' dataset (~350M transaction rows, 300,000+ shoppers) and the KDD Cup 98 charity-donor dataset (~200,000 lapsed donors); no mobile-app install-cohort validation is documented in the paper or its repository.	direct_record	SUPPORTED	high	GH-S12 GH-S13	
GH-F19	Industry commentary on backtesting mobile LTV models (Seufert, 2015) argues that historical backtests systematically miss changes in the competitive/market environment (e.g. rising UA-inventory competition over a 12-24 month LTV-projection horizon) and cohort-size effects (network effects or saturation), making backtest accuracy a poor guarantee of forward predictive accuracy; the piece recommends continuous real-time recalibration over static backtested models.	inference	SUPPORTED	low	GH-S14	
GH-F20	Singular (an MMP) describes SKAN-era predicted-LTV practice as modeling missing conversion data from a single 6-bit SKAN postback per install, encoding cohort/revenue/behavior signals, and states (via an internal executive quote, not an independent audit) that sophisticated implementations can reach performance 'comparable to' IDFA-era measurement; the vendor's own article discloses no accuracy percentages, sample sizes, or validation methodology.	vendor_assertion	QUALIFIED	low	GH-S19	A Dec 2021 Singular blog post describes SKAN-era: one 6-bit postback per install, encoding cohort/revenue/behavior signals, and states Singular's conversion value to be comparable to IDFA and quotes Singular's customer claims of reliability at 'level of scale' as IDFA accuracy data or validation methodology.

GH-F21	A Unity-authored analysis of 8 high-DAU apps (iOS and Android) found users who engaged with rewarded video ads were 4.5x more likely to make an in-app purchase than non-engaged users, showed an average weighted 326% increase in user spend (7 days before vs. after first rewarded-ad engagement, range ~200%-500% across apps), a 34% average weighted increase in daily sessions, and single-rewarded-ad-completion 30-day retention of at least 50% vs. an unspecified 13% baseline (rising to 53-68% for higher engagement, 3.5-5x the stated benchmark).	vendor_assertion	SUPPORTED	low	GH-S15	
GH-F22	OneSignal's 2026 best-practices guide cites named-customer case studies (Beach Bum Games: push CTR from <1% to 12% and a 250% engagement increase from quest-based personalization; Gedi Digital: 80% YoY growth in iOS subscription activations; Whisker: 20% increase in mobile-driven sales) with no disclosed sample size, test duration, control condition, or independent verification.	vendor_assertion	SUPPORTED	low	GH-S17	
GH-F23	Sensor Tower's State of Mobile 2026 report states global app-store consumer in-app-purchase spend reached \$167 billion in calendar-year 2025 (+10% YoY), with non-gaming IAP revenue surpassing gaming IAP revenue for the first time, growing 21% YoY and attributed partly to generative-AI services.	vendor_assertion	SUPPORTED	medium	GH-S16	
IJ-F01	MRC's June 2, 2017 interim guidance requires measurement vendors to evaluate mobile in-app SIVT/GIVT detection discretely per-app rather than uniformly, supplementing the October 2015 IVT Detection and Filtration Guidelines Addendum.	synthesis	QUALIFIED	medium	IJ-S01	Per a Pixelate summary of the 2017 interim guidance, the Oct 2015 IVT Addendum requires vendors to treat mobile in-app as a distinct environment with different detection parameters where it is a material source of traffic.
IJ-F02	comScore announced on December 4, 2017 that it received MRC accreditation for Sophisticated Invalid Traffic (SIVT) detection and filtration specifically covering mobile in-app traffic, completing its accreditation across desktop, mobile web, and in-app.	vendor_assertion	QUALIFIED	medium	IJ-S16	comScore announced its vCE product had received MRC accreditation for mobile in-app traffic; comScore does not appear to have a digital accreditation list as of September 2026.
IJ-F03	TAG's Certified Against Fraud program requires that 100% of a certified participant's monetizable transactions (impressions, clicks, conversions) be filtered for invalid traffic compliant with an MRC-recognized standard, and its scope explicitly includes in-app inventory alongside display, video, and mobile web.	vendor_assertion	QUALIFIED	medium	IJ-S02	TAG's Certified Against Fraud program (v11.0, July 2026) requires participant's monetizable transactions (impressions, clicks, conversions) be IVT-filtered under a TAG-recognized standard (MRC IVT Guidelines or China's T/CAAD 001-2017) and in-app obligations such as filtering and app-ads.txt.
IJ-F04	A 2019 TAG-commissioned benchmark study reported a 1.41% invalid-traffic (IVT) rate across more than 201 billion display, video, and mobile impressions measured within TAG-certified distribution channels, described as an 88% reduction versus a broader industry average.	commissioned_measurement	QUALIFIED	low	IJ-S03	TAG reports that its 2019 Invalid Traffic Study (conducted by TAG) found a 1.41% IVT rate across more than 201 billion display, video and mobile impressions measured within TAG Certified channels, below the broader industry average of 1.68% over 75B impressions.
IJ-F05	Apple's App Store Review Guideline 3.2.2(x) prohibits apps from requiring users to rate, review, download other apps, or take similar store-related actions to access functionality, content, or compensation, while explicitly permitting apps to incentivize users for in-app actions such as watching an ad or completing a level.	direct_record	QUALIFIED	high	IJ-S13	App Store Review Guideline 3.2.2(x) (June 2026) bars apps from requiring users to rate or review the app, download other apps, or take other store-related actions to access functionality, content, or compensation, while allowing in-app actions such as completing a level or watching an ad.
IJ-F06	Google's June 5, 2017 Play Store policy update prohibits manipulating app rankings via fraudulent or incentivized installs, reviews, or ratings, while stating that incentivized installs can still serve as a legitimate user-acquisition channel when not designed to inflate rankings.	direct_record	SUPPORTED	high	IJ-S14	
IJ-F07	A peer-reviewed 2020 study (ACM Internet Measurement Conference) documents incentivized-install platforms that pay users to install apps and complete in-app tasks, inflating daily-active-user counts, session length, and revenue metrics, and finds lax Google Play enforcement against this practice.	independent_measurement	SUPPORTED	medium	IJ-S15	
IJ-F08	AppsFlyer's Protect360 documents detecting click injection through anomaly analysis of click-to-install-time (CTIT) timestamps, and click spamming/flooding through statistical (maximum-likelihood-estimation) analysis of the CTIT distribution, applying both real-time pre-attribution blocking and post-attribution reconciliation.	vendor_assertion	QUALIFIED	medium	IJ-S17	AppsFlyer markets real-time attribution blocking plus post-attribution reconciliation (Protect360); Protect360 glossary describes identifying and blocking flooding sources by looking at device distributions, low click conversion and high non-contributor rates; specific methods (e.g. MLE) are used to check the sources.
IJ-F09	Adjust documents an SDK-signing mechanism ('SDK Signature') using custom encryption keys to detect and reject SDK-spoofed (replayed) install/event requests, and describes device/install farms as partly detected via anonymous-IP and proxy-traffic heuristics.	vendor_assertion	QUALIFIED	low	IJ-S18	Adjust documents SDI (proprietary signing with anti-emulation checks) to detect invalidly signed requests against SDK spoofing, IP filter (VPNs, Tor, data centers), and MaxMind that routes traffic to 'Untrusted Devices' but is aimed at device farms.

IJ-F10	Uber's November 13, 2017 cross-complaint against Fetch Media and Phunware alleged fraudulent, non-viewable, or otherwise ineligible ad placements tied to over \$82.5 million Uber paid for a mobile ad campaign run 2014 through Q1 2017; the complaint was amended in 2019 to add civil RICO claims against named Phunware executives.	direct_record	QUALIFIED	medium	IJ-S07 IJ-S06	Uber sued Fetch Media in Sept 2017, alleging it paid \$82.5M in 2016-Q1 2017 for fraudulently attributed ad impressions. Separately, in Phunware Court suit, Uber cross-sued Phunware and Fetch Media, adding civil RICO claims against Phunware-linked individuals in 2019.
IJ-F11	Uber's ad-fraud dispute resolved in two settlements: Uber and Fetch Media settled in March 2019 on undisclosed terms, and separately Uber, Phunware, and four named individuals executed a \$6,000,000 settlement on October 9, 2020 (\$1.5M from insurers, \$4.5M from Phunware in three \$1.5M installments through September 30, 2021), with all parties denying wrongdoing.	direct_record	QUALIFIED	high	IJ-S04 IJ-S06	Uber, Phunware and four named individuals settled with insurers; \$4.5M from Phunware in three installments through 3/31/21 denying wrongdoing, and Uber struck Phunware's plea in Aug 2020; the reported settlement is not confirmed source here.
IJ-F12	IAB Tech Lab's Open Measurement SDK (OM SDK), which lets third-party viewability/verification vendors access in-app measurement signals through a common API (OMD) instead of separate per-vendor SDKs, saw in-app coverage in one viewability vendor's (IAS) own measured impressions rise from roughly 1% to roughly 60% within about a year, while DSP-side use of the resulting signals lagged supply-side adoption.	vendor_assertion	QUALIFIED	low	IJ-S19 IJ-S20	IAB Tech Lab released 2018 so verification vendors measure in-app ads via IAS (which originally did not report OM coverage) and app impressions rising from 1% to ~60% (Sept 2018) to ~60% (Sept 2020); DSPs used OM signals.
IJ-F13	app-ads.txt adoption reached an estimated 66% among the top 1,000 Google Play apps versus roughly 24% overall adoption across all measured apps, with about 25% of app bid requests still originating from apps lacking a known app-ads.txt file.	vendor_assertion	QUALIFIED	low	IJ-S21	HUMAN Security (verified) reported that 25% of app bid requests observed were for apps lacking app-ads.txt file (vs 3% lacking ads.txt).
IJ-F14	sellers.json coverage gaps were reported at approximately 8.2% of web bid requests and 9.1% of app bid requests coming from sellers without a known sellers.json record.	vendor_assertion	SUPPORTED	low	IJ-S21	
IJ-F15	AppsFlyer's CTV-to-mobile attribution uses IP address and user-agent-based cross-platform link matching (OneLink) to attribute an app install to a prior CTV ad exposure, and explicitly allows multiple installs within the same household (across iOS, Android, and other platforms) to be credited to a single CTV ad view - a household-level, probabilistic bridge, distinct from AppsFlyer's separately-sold incrementality product for causal measurement.	vendor_assertion	QUALIFIED	medium	IJ-S09 IJ-S25	AppsFlyer's CTV-to-mobile attribution uses cross-platform OneLink-based probabilistic matching to attribute app installs by multiple users to the same CTV impression; measurement is a separately-sold incrementality product.
IJ-F16	Kochava's Device Link/IdentityLink graph performs deterministic device matching only when device ID, device user agent, and IP address are all passed together; absent all three it falls back to probabilistic matching, and Kochava's Samba TV partnership layers this device graph on top of Samba TV's self-reported opted-in smart-TV panel (approximately 50 million TVs across all 210 US DMAs plus nine other countries), using 'synthetic control groups' - a modeled comparison, not a randomized holdout - for lift measurement.	vendor_assertion	QUALIFIED	low	IJ-S10	Kochava's Samba TV blog, Dec 2025) links to Samba TV's ~50M opted-in smart-TV panel; operational in 210 US household/device identifier measures lift with post-exposure control groups, a modeled comparison, not a randomized holdout.
IJ-F17	MNTN's 'Verified Visits' technology claims household-level deterministic matching, via a proprietary identity graph said to cover 99% of available US households, crediting a conversion when the same household is linked between a full/non-skipable CTV ad view and a later website or app visit.	vendor_assertion	QUALIFIED	low	IJ-S12	MNTN's Verified Visits technology credits a website visit defined window when identity graph, claimed to cover 99% of available US households, links a household that watched a skipable CTV ad; the details are not disclosed.
IJ-F18	Amazon DSP can attribute a mobile app install to a CTV ad exposure (e.g., on Prime Video/Fire TV) when the same user is logged into an Amazon account on both surfaces - a deterministic, individual-account-level identity basis, distinct from the IP/household-graph matching used by other CTV-to-app vendors examined in this stream.	vendor_assertion	WITHDRAWN	medium	IJ-S26	
IJ-F19	Moloco's Performance CTV product (announced April 2026) reports vendor-cited early results of up to 1.5x higher ROI on CTV than mobile in cross-channel campaigns, with approximately two-thirds of resulting app installs occurring within six hours of ad exposure, while its own press release does not disclose the identity-matching method (deterministic vs. probabilistic, household vs. individual) or any causal/incrementality test design, instead crediting installs via 'the buyer's preferred MMP.'	vendor_assertion	SUPPORTED	low	IJ-S08	
IJ-F20	Roku's Ads Manager integrates with AppsFlyer, Adjust, and Branch via pixel/Conversions API for app-install measurement and offers 'Action Ads' letting viewers text themselves a download link or go to an app store, but Roku's own public documentation does not disclose the underlying identity-matching mechanism (device ID vs. household IP) used for its app-install attribution.	vendor_assertion	QUALIFIED	medium	IJ-S11	Roku Ads Manager does not disclose the pixel and a server-to-server API to record outcome and to optimize likely-converting households; disclosing the identity-matching mechanism and MMP integrations and verified on this page.
IJ-F21	Vendors in this space describe 'incrementality'/causal testing as a distinct, separately-designed layer from exposure-based identity attribution: tvScientific markets deterministic 1:1 ID-based exposure attribution alongside separately-branded 'always-on incrementality testing,' and iSpot.tv's case-study figures (e.g., a reported 23% increase in leads for one brand) are presented without a disclosed randomized control group, reading as exposed-population before/after or exposed-only metrics rather than audited causal lift.	synthesis	UNRESOLVED	low		
IJ-F22	No independently designed, published, peer-reviewed, or regulator-audited randomized/geo-holdout study isolating CTV advertising's causal effect on mobile app install or in-app event outcomes specifically (as distinct from general TV-advertising-to-sales causal literature) was found as of the cutoff; every CTV-to-app 'incrementality' or 'lift' claim located in this research was vendor-designed and vendor-reported.	inference	UNRESOLVED	medium		

IJ-F23	AppLovin's e-commerce advertising business (Axon Ads Manager) launched as a referral-only self-serve platform on October 1, 2025 with an initial eligibility threshold of at least \$10 million in gross merchandise value (GMV) for participating advertisers, expanding from a few hundred initial advertisers, with a scheduled fully public opening in June 2026.	direct_record	QUALIFIED	medium	IJ-S23	AppLovin launched Axon Ads Manager as a referral-only self-serve platform on 1 Oct 2025 (advertisers live), after limited to brands with 100,000+ app installs. On 22 Jun 2026 it opened to all advertisers and renamed AppLovin Ad Manager. The platform launched without a referral code requirement through Q2 2026 due to low e-commerce revenue.
IJ-F24	AppLovin's total company revenue for Q1 2026 was \$1.84 billion (up 59% year-over-year), a company-wide figure not broken out by e-commerce (Axon Ads Manager) versus mobile-app-install advertising in available public reporting, while CEO Adam Foroughi's stated target of roughly \$7 billion in first-year ad spend from 100,000 new e-commerce customers is a forward-looking management projection, not a realized or reported figure.	forecast	QUALIFIED	medium	IJ-S22	AppLovin reported Q1 2026 revenue of \$1.842B (+59% YoY; +10% QoQ). E-commerce segment, no e-commerce revenue reported. The call the CEO offered during the Q1 2026 earnings exercise that 100,000 new e-commerce customers in the first year of launch would imply ~\$7B in ad spend, which is not a realized result.
IJ-F25	Moloco Commerce Media (a retail-media/onsite-advertising product, distinct from app-install advertising) reports Costco selected it to power onsite advertising on Costco.com around April 2026, and separately claims that a majority of its retail-media customers 'at least doubled' ad revenue as a percentage of GMV within their first four weeks on the platform.	vendor_assertion	QUALIFIED	low	IJ-S24	Moloco claims (internal documents, 2026) that most Moloco retail-media customers (retailers at least double share of GMV between app and retail-media) within weeks and their latest platform; Costco is shown as a customer.
IJ-F26	Google Play's policy stance (tolerating incentivized installs as a legitimate acquisition channel absent ranking manipulation) and Apple's stance (banning any requirement to install/rate/review another app to unlock functionality, while allowing in-app-only incentives like rewarded ads) draw the line between legitimate rewarded engagement and abusive incentivized-install manipulation differently by platform.	synthesis	QUALIFIED	medium	IJ-S13 IJ-S14	Apple bars forcing users to rate or review other apps, or to rate or review the app to unlock functionality, while allowing in-app-only incentives such as rewarded ads. Google Play's current policy bans incentivized installs but placement through illegitimate channels is allowed. These include incentivized installs, incentivized ratings, and incentivized app reviews. Apple's 2017 post called incentivized installs a legitimate channel for app discovery, but not used to game rank or review. Apple's policy text does not require
IJ-F27	MMP fraud-remediation documentation (AppsFlyer) describes post-attribution reconciliation in general terms as letting advertisers 'reclaim wasted spend and recover partner payouts,' but does not publicly specify a standardized, contractual refund/makegood process (e.g., a stated SLA, dispute window, or automatic chargeback mechanism) with ad-network partners for confirmed fraud.	vendor_assertion	QUALIFIED	medium	IJ-S17	AppsFlyer's fraud-prevention documentation describes post-attribution reconciliation in general terms as letting advertisers 'reclaim wasted spend and recover partner payouts,' but does not publicly specify a standardized, contractual refund/makegood process (e.g., a stated SLA, dispute window, or automatic chargeback mechanism) with ad-network partners for confirmed fraud.
IJ-F28	MRC's current digital accreditation listing shows Sophisticated Invalid Traffic (SIVT) accreditation covering mobile-app (MA) inventory held by DoubleVerify, Google Ads, Google DV360, HUMAN, Integral Ad Science, Pixelate, and Protected Media (each also covering desktop, mobile web, and CTV) - but no mobile measurement partner that app marketers use for install attribution (AppsFlyer, Adjust, Kochava, Singular, Branch) appears on this list.	direct_record	QUALIFIED	high	IJ-S27	MRC's digital accreditation listing shows SIVT accreditation covering mobile-app (MA) inventory held by DoubleVerify, Google Ads, Google DV360, HUMAN, Integral Ad Science, Pixelate, and Protected Media (each also covering desktop, mobile web, and CTV) - but no mobile measurement partner that app marketers use for install attribution (AppsFlyer, Adjust, Kochava, Singular, Branch) appears on this list.
KL-F01	82% of ANA member companies operated an in-house agency as of a 2023 ANA survey of 162 respondents, up from 78% in 2018 and 58% in 2013.	commissioned_measurement	SUPPORTED	medium	KL-S16	
KL-F02	In ANA's 2026 in-house survey of 404 award jurors, only 9% cited cost-saving as the primary benefit of in-house agencies, down from 30% in the 2023 survey, while 92% (2023 figure) still used external agencies for specialized capacity.	commissioned_measurement	QUALIFIED	medium	KL-S15 KL-S16	In ANA's 2026 State of Advertising survey of 404 award jurors (representing 100% of ANA's 2026 member survey), the survey found that 9% named cost savings as the primary benefit of in-house agencies, citing cost efficiencies as the primary benefit. The survey wording both changed and remained the same for-like trend.
KL-F03	Google's official Google Ads API MCP server (released ~Apr 28, 2026) is documented as strictly read-only: it can list accessible accounts, run GAQL queries, and fetch resource metadata, but cannot modify bids, pause campaigns, or create assets.	direct_record	QUALIFIED	high	KL-S06	Google's open-source MCP server (announced Sep 25, 2026) is documented as strictly read-only: it lists accessible accounts, runs GAQL queries, and fetches resource metadata, but cannot modify bids, pause campaigns, or create assets.
KL-F04	Meta's officially hosted Ads MCP server (opened in beta ~Apr 29, 2026 at mcp.facebook.com/ads) exposes write-capable tools across categories including ad/catalog creation and management and A/B test/conversion-lift-study operations, alongside read-only reporting and activity-log tools.	direct_record	QUALIFIED	medium	KL-S07	Meta's hosted Ads MCP server (opened in beta ~Apr 29, 2026 at mcp.facebook.com/ads) is documented as strictly read-only: it lists accessible accounts, runs GAQL queries, and fetches resource metadata, but cannot modify bids, pause campaigns, or create assets.
KL-F05	TikTok's official Agentic Hub / MCP Server exposes approximately 400 advertising API functions including campaign creation, budget/bid changes, and audience/creative operations, and its own help documentation states agents can execute complete workflows 'without human intervention at every step,' with no approval-gate mechanism described on that page.	direct_record	SUPPORTED	high	KL-S08	

KL-F06	AdCP (Ad Context Protocol) documents a live agent-to-agent media buy using real inventory and money on October 16, 2025, and its current specification is versioned 3.2.0-rc.3 as of the site content reviewed.	direct_record	QUALIFIED	medium	KL-S01	AgenticAdvertising.org agent-to-agent media Ads inventory) ran on (self-reported); at the stable release is 3.1.24 candidate (3.2.0-rc.6)
KL-F07	AdCP's documented governance model routes campaign actions above a human-set authority threshold (e.g., a budget cap) to human approval before execution, and provides a get_plan_audit_logs function giving a full decision trail of who proposed, approved, and what ran.	direct_record	QUALIFIED	medium	KL-S03	AdCP's campaign-gov specification (v3.1.24) opt-in surface. A plan reallocation threshold governance agent esc review. Every check_g produce an audit log € get_plan_audit_logs. B check_governance on governance agent is c plan. Governance age mode, approving ever attached. Approval-flc tests apply only to sel governance_aware.
KL-F08	AdCP is governed by AgenticAdvertising.org through four equally-weighted voting classes (brands, agencies, publishers, technology providers), with an interim board through May 6, 2026 and a target elected board of 40 seats (10 per class) plus CEO.	direct_record	SUPPORTED	high	KL-S02	
KL-F09	IAB Tech Lab's Agentic RTB Framework (ARTF) v1.0 is a finalized specification operating at the OpenRTB/data-center infrastructure layer (segment activation, deal management, bid shading, metrics), distinct in scope from AdCP's advertiser-facing planning/buying/creative/measurement tasks; its own README does not reference AdCP or an advertiser-facing approval mechanism.	direct_record	SUPPORTED	high	KL-S04	
KL-F10	AAMP ('Agentic Advertising Management Protocols') is a real, currently-maintained IAB Tech Lab umbrella initiative (version 3.0, last updated Sept 22, 2026) for agentic AI in advertising, distinct from — and not cross-referenced with, on either primary page reviewed — both IAB Tech Lab's own separately-documented ARTF (KL-S04) and AgenticAdvertising.org's AdCP (KL-S01).	direct_record	CONTRADICTED	high	KL-S35	AAMP (Agentic Adver Protocols) is IAB Tech initiative for agentic A (AAMP 3.0; page last 2026); its GitHub repo among its component page nor repository m
KL-F10b	AAMP 3.0 (Sept 2026) covers end-to-end ad decisioning/delivery workflows including OpenDirect direct transactions, Programmatic Guaranteed, Preferred Deals, PMPs, open programmatic marketplace, and Linear GRPs, and newly adds 'OpenProposal,' described by IAB Tech Lab as automating the discovery-and-planning phase of a campaign — a scope that overlaps with AdCP's stated discovery/planning/buying/creative/measurement scope (KL-S01), without either primary source acknowledging the other.	synthesis	SUPPORTED	medium	KL-S35 KL-S01	
KL-F11	AppLovin's Axon AI legal disclosure page states the model uses device/network engagement signals, MAX auction win/loss notifications, and advertiser-shared data, explicitly excludes sensitive personal information and other bidders' data, but contains no language addressing whether advertisers retain rights to models or learnings trained on their own data.	direct_record	SUPPORTED	high	KL-S12	
KL-F12	Automation-of-bidding features (Google AI Max for Search's automatic campaign upgrades in Sept 2026, Moloco's per-impression ML bidding, AppLovin Axon AI's predictive impression-value scoring) are documented as live vendor features, and should be treated as distinct from agentic planning/execution (AdCP agent-to-agent buys, write-capable MCP servers) per the brief's requested separation.	synthesis	SUPPORTED	medium	KL-S11 KL-S13 KL-S12	
KL-F13	Unity's 'agentic Vector' generative-creative capabilities (GenAI Creative, Creative Optimization, Creative Insights) were described in Unity's own June 2026 Ads newsletter as 'coming with' agentic Vector, indicating announced/beta status rather than confirmed general availability at the cutoff.	vendor_assertion	QUALIFIED	low	KL-S14	A June 2026 Unity Ad on github.io, not a unit GenAI Creative, Creat Creative Insights as c Vector' on its H2 2026 announced, not beta c at the cutoff.
KL-F14	The FTC's 2025 COPPA Rule amendments took effect June 23, 2025, with the general compliance deadline April 22, 2026 (one year after Federal Register publication), except for certain safe-harbor-related subsections that had earlier compliance deadlines.	direct_record	SUPPORTED	medium	KL-S17	
KL-F15	The FTC finalized a stipulated order against Kochava Inc. and its subsidiary Collective Data Solutions on May 4, 2026, filed in the U.S. District Court for the District of Idaho, banning the sale, licensing, transfer, or sharing of sensitive location data without consumer consent and requiring a sensitive-location identification program, supplier consent assessments, and a data-deletion schedule.	direct_record	QUALIFIED	high	KL-S19	The FTC filed a propo against Kochava Inc. & Solutions in the U.S. D District of Idaho on 4 entered it (Dkt. 138) o bars the defendants fr transferring, sharing c location data tied to s identified under a req one exception needs i relationship, affirmativ and a service the con requested. It also req assessments and data defendants neither ad allegations. The order
KL-F16	The FTC previously finalized similar sensitive-location-data orders against X-Mode/Outlogic (April 2024) and InMarket (May 1, 2024), establishing a consistent enforcement pattern that the Kochava order (2026) extends.	direct_record	SUPPORTED	medium	KL-S20 KL-S21	
KL-F17	Utah's App Store Accountability Act was amended by HB498 (signed March 18, 2026) to delay its effective date to May 6, 2027 and remove the state Attorney General's enforcement authority, leaving only a private right of action; CCIA voluntarily dismissed its constitutional challenge to the original law on April 21, 2026 as a result.	direct_record	SUPPORTED	medium	KL-S22 KL-S23	

KL-F18	Texas's App Store Accountability Act (SB2420) was preliminarily enjoined by a federal district court on December 23, 2025, but the Fifth Circuit stayed that injunction on June 4, 2026, finding the district court likely erred in applying strict-scrutiny review; the law is currently in effect and enforceable while the underlying appeal continues.	direct_record	SUPPORTED	medium	KL-S24	
KL-F19	Louisiana's app-store age-verification law was recast by HB977, signed by the Governor as Act No. 185; three sources give three different effective dates for the resulting requirement — the official Louisiana Legislature bill-status page's own status line reads 'Act 185 (effective 05/15/2026)' (the signature date), while secondary legal commentary separately reports a July 1, 2026 effective date and, elsewhere, a delay to 2027 — this three-way conflict is unresolved.	direct_record	CONTRADICTED	low	KL-S25 KL-S25b	Louisiana Act 185 of 2026 (May 15, 2026) repeals the 2021 app-store law (Act 481) but is not effective on signature; the replacement app-store law requires a category requirement by July 1, 2027.
KL-F20	The Ninth Circuit's March 12, 2026 ruling in NetChoice v. Bonta kept five of six challenged provisions of California's Age-Appropriate Design Code Act enjoined (including the data-use restrictions and the dark-patterns prohibition, found likely unconstitutionally vague), while remanding the age-estimation/coverage-definition provision to the district court for further factual development.	direct_record	QUALIFIED	medium	KL-S26	In NetChoice v. Bonta (9th Cir. 2026) the court rejected the whole-statute challenge to the AADC, finding the injunction on the AADC provision, while the District Court's 90-day notice-and-comment rulemaking for the dark-patterns provision's severability was remanded to the district court.
KL-F21	Google's Play Age Signals API (in beta) returns age-range signals to developers and began live rollout in Brazil on March 17, 2026 (tied to Brazil's Digital ECA) and in Texas on May 28, 2026 (tied to SB2420), with Google's terms explicitly prohibiting use of the signal for advertising, marketing, user profiling, or analytics.	direct_record	SUPPORTED	high	KL-S29	
KL-F22	Apple's Declared Age Range framework/API exists for iOS apps to request a person's age range, and at least one Apple developer-forum thread ties its implementation guidance to Texas SB2420 compliance, but in this research could not retrieve the full official documentation body (only page titles/headers rendered).	vendor_assertion	QUALIFIED	low	KL-S27	Apple's Declared Age Range framework/API exists for iOS apps to request a person's age range, and at least one Apple developer-forum thread ties its implementation guidance to Texas SB2420 compliance, but in this research could not retrieve the full official documentation body (only page titles/headers rendered).
KL-F23	Apple's App Store Review Guideline 5.1.2(i) requires apps to obtain permission via the App Tracking Transparency API before tracking user activity, and separately bans surreptitious user profiling (5.1.2(iii)) and use of HomeKit/HealthKit/ARKit/facial-mapping data for marketing or advertising (5.1.2(vi)); Guideline 5.1.4 restricts third-party analytics/advertising in apps primarily directed at children.	direct_record	SUPPORTED	high	KL-S28	
KL-F24	Google Play's Families policy requires child-directed ad serving to use only Families Self-Certified Ads SDKs, bans interest-based advertising and remarketing directed at children, and prohibits several ad content categories and non-dismissible/full-screen-at-launch ad formats.	direct_record	SUPPORTED	high	KL-S30	
KL-F25	Apple updated its EU business terms on August 18, 2026, replacing a layered fee structure (Core Technology Fee, Initial Acquisition Fee, Store Services Fee) with a single set of business terms per payment path, following ongoing engagement with the European Commission under the DMA; the Commission separately found Apple's app-sideload compliance mechanism still does not satisfy DMA requirements.	direct_record	QUALIFIED	medium	KL-S32 KL-S33	Apple updated its EU business terms on August 18, 2026, replacing a layered fee structure (Core Technology Fee, Initial Acquisition Fee, Store Services Fee) with a single set of business terms per payment path, following ongoing engagement with the European Commission under the DMA; the Commission separately found Apple's app-sideload compliance mechanism still does not satisfy DMA requirements.
KL-F26	The three major ad platforms' officially hosted MCP servers show materially different control postures at launch: Google chose read-only by design (Apr 28, 2026), while Meta (Apr 29, 2026) and TikTok (May 13, 2026) shipped write-capable servers, illustrating that 'agentic' access to ad platforms is not a single uniform capability but a spectrum from reporting-only to full campaign-execution control.	synthesis	QUALIFIED	high	KL-S06 KL-S07 KL-S08	The three platforms' control postures differ in control posture: Google chose read-only by design (Apr 28, 2026), while Meta (Apr 29, 2026) and TikTok (May 13, 2026) shipped write-capable servers, illustrating that 'agentic' access to ad platforms is not a single uniform capability but a spectrum from reporting-only to full campaign-execution control.
MAIN-F01	AppsFlyer estimates global app marketing spend at \$109B in 2025: \$78B on user acquisition (up 13% year over year) and \$31.3B on remarketing (up 37%), with remarketing's share of the total rising from 25% in 2024 to 29% in 2025.	vendor_assertion	QUALIFIED	medium	MAIN-S01	AppsFlyer estimates, for 2025, that global app marketing spend is \$109B (32B paid installs; 45K installs), global app marketing spend is \$109B for 2025 - \$78B on user acquisition (+13% share 25% to 29% - in 2025) and \$31.3B on remarketing (+37% share 25% to 29% - in 2025) - in 2025, published 10 December 2025, neither the measurement nor the panel spend is scaled.
MAIN-F02	In AppsFlyer's 2025 panel estimate, UA spend growth came entirely from iOS (+35%) while Android was flat (-1%); non-gaming UA rose 18% to \$53B and gaming rose 3% to \$25B.	vendor_assertion	SUPPORTED	medium	MAIN-S01	
MAIN-F03	The US commanded 42% of global UA spend in AppsFlyer's 2025 panel, which implies roughly \$33B of US app user-acquisition spend (0.42 x \$78B), before remarketing.	synthesis	QUALIFIED	low	MAIN-S01	AppsFlyer states the US commanded 42% of global spend in its 2025 panel (basis not explicitly stated, but applied to UA spend, not total spend); \$33B of US app UA spend (0.42 x \$78B) - in 2025, published 10 December 2025, neither the measurement nor the panel spend is scaled.
MAIN-F04	AppsFlyer reports iOS paid installs rose 31% in the US in 2025 while Android paid installs rose 8%.	vendor_assertion	SUPPORTED	medium	MAIN-S01	
MAIN-F05	Google Ads Help states App campaigns need at least 10 conversions every day (or 300 in 30 days), each with value above zero and coming from the Firebase SDK, to use Target ROAS.	direct_record	QUALIFIED	high	MAIN-S03	Google Ads Help states App campaigns need at least 10 conversions every day (or 300 in 30 days), each with value above zero and coming from the Firebase SDK, to use Target ROAS, that the bid on (and sending value from the Google Analytics SDK, and recommend first to establish a baseline.

MAIN-F06	Apple states a 1.6% average conversion rate on default product pages and says developers see a 2.5 percentage point increase on average when referring people to a custom product page, with no period, definition or test design given.	vendor_assertion	SUPPORTED	medium	MAIN-S04	
MAIN-F07	Meta says ad sets exit the learning phase once delivery is stable, which usually happens after about 50 results in the week after the last significant edit, and advises combining similar ad sets because many ads and ad sets reduce what the system learns about each.	direct_record	SUPPORTED	high	MAIN-S05	
MAIN-F08	In February 2017 VIZIO agreed to pay \$2.2 million to settle FTC and New Jersey charges that it collected second-by-second viewing data from 11 million smart TVs without consumers' knowledge or consent; the order requires affirmative express consent for such collection and sharing.	direct_record	SUPPORTED	high	MAIN-S06	
MAIN-F09	Meta's help page for Advantage+ app campaigns limits targeting to operating system, country and language, applies Advantage+ placements automatically, and lists placement among supported reporting breakdowns; it describes no breakdown by publisher app.	direct_record	SUPPORTED	medium	FX1-S23	
V1-F01	AppLovin's FY2025 10-K attributes its revenue step-change primarily to its own Axon Ads Manager product and AXON model upgrades, per company language, rather than to an external, independently-audited efficiency measure.	vendor_assertion	QUALIFIED	low	V1-S01	AppLovin's FY2025 10-K revenue increase (to \$ of publisher costs) prior to Axon Ads Manager pe installations up 3% an installation up 72%, an grew rapidly 'in partic of Axon AI'; it cites no efficiency or incremen
V1-F02	Mobvista's H1 2026 interim report discloses that smart-bidding products (Target ROAS variants and Target CPE) generate over 90% of Mintegral's revenue, and that Mintegral's non-gaming revenue grew 15.5% YoY to reach 24.4% of Mintegral revenue in H1 2026 (six months ended 30 June 2026), versus 75.6% for gaming (+27.5% YoY).	direct_record	SUPPORTED	high	V1-S37	
V1-F03	Moloco documents a 'built-in' ghost-bidding incrementality-testing method for both its Performance CTV product and its Re-engagement product, presented as a standard, non-custom feature of the platform rather than a bespoke or third-party add-on.	vendor_assertion	QUALIFIED	medium	V1-S13 V1-S14	Moloco's Performance documents built-in gh incrementality tests; it page says only that M incrementality testing/ method or calling it bu
V1-F04	Among the six V1 deep-profiled vendors, only AppLovin, Moloco and Unity have a documented campaign-management API (not just reporting) confirmed via a directly-opened technical source in this research; Liftoff's and Mintegral's self-serve/API depth could not be confirmed, and Digital Turbine's documented self-serve control (Micro Bidding) is UI-based rather than API-based in the pages found.	synthesis	SUPPORTED	medium	V1-S06 V1-S17 V1-S34 V1-S37 V1-S46	
V2-F01	Verve Group SE (ticker VER) was renamed to Verve Group Media SE at its 2026-06-05 AGM in Stockholm, which also approved relocating the registered office from Sweden to Ireland.	direct_record	QUALIFIED	medium	V2-S12	Verve Group SE now t Group Media SE (com Sep 2026). Its 5 June moving the registered Stockholm to Dublin; t taken effect at the 27 (expected 2 Oct 2026 IE0002VY0237).
V2-F02	Smadex operates as a fully managed-service DSP, not self-serve: Entravision's FY2025 10-K states Smadex teams 'configure bidding parameters, manage fraud prevention protocols' on the client's behalf.	direct_record	SUPPORTED	high	V2-S08	
V2-F03	Aarki rebranded to RZR on 2026-03-17, positioning as a 'connected, cross-screen performance platform' spanning UA, retargeting (Encore), and CTV.	vendor_assertion	QUALIFIED	medium	V2-S18	Aarki has rebranded a Inc.; site notice live by RZR positions itself as screen system spann retargeting and CTV, r bidding engine; curre disclosed.
V2-F04	Chartboost (DSP + SSP/mediation + SDK) was acquired by LoopMe from Zynga, closing 2024-12-10; deal value undisclosed and post-acquisition DSP product roadmap unspecified.	direct_record	QUALIFIED	high	V2-S27	LoopMe announced o that it had acquired C (programmatic advert mediation platform) fr terms and closing dat
V2-F05	Jampp documents an always-on incrementality-testing product ('Ghost Bids,' modeled on Ghost Ads) offered to all campaigns at no additional cost, one of the more specifically-named incrementality claims found across V2 vendors.	vendor_assertion	QUALIFIED	low	V2-S22	Jampp's product page on incrementality ('liff UA and retargeting ca 'Ghost Bids', available additional cost; the m documented.
V2-F06	Verve Datasat's own product materials claim CTV spend is 'validated via controlled incrementality testing' but do not name a specific measurement design (geo/PSA/ghost/ITT) or identity method for CTV-to-mobile linkage.	vendor_assertion	SUPPORTED	low	V2-S13	
V2-F07	Affle 3i Limited's app-marketing portfolio comprises at least five separately branded DSP/ad-network products discovered in this research — Jampp, YouAppi, RevX, mediasmart, and Appnext — each still operating under its own brand and site.	direct_record	QUALIFIED	medium	V2-S21 V2-S22 V2-S23 V2-S16 V2-S30 V2-S31 V2-S32	Affle 3i owns at least t branded app-marketir (Jampp, YouAppi, Rev Appnext's Affle link w from sources opened, agreed to acquire the platform and brand fr Digital Turbine).
V2-F08	Applier's AIQUA (a frequently-cited Applier product) is a CRM/lifecycle-personalization tool, not a user-acquisition DSP; Applier's separate 'Ad Cloud' suite (including Retargeting and AIBID) is the actual media-buying/retargeting product relevant to the V2 app-DSP category.	vendor_assertion	SUPPORTED	medium	V2-S29	
V2-F09	Persona.ly, a candidate omission/closure to check, remains an active, independently operated mobile DSP as of the research access date — no acquisition or shutdown found.	direct_record	QUALIFIED	medium	V2-S28	Persona.ly's website i markets a programma UA and re-engagemen shutdown notice was; ownership is not state

V3-F01	Meta's own Business Help Center states Advantage+ App Campaign advertisers 'won't have the option to see reporting insights for specific placements,' with targeting limited to OS, country, and language.	direct_record	CONTRADICTED	medium	V3-S05	Meta's help page says campaigns limit target system, countries and automatically use Adv and a 90-day exclusio support reporting brei placement (SKAN can in delivery data).
V3-F02	Google's own App campaigns API documentation does not specify which placements, networks, or apps receive ads, and provides no placement-level reporting for App campaigns.	direct_record	CONTRADICTED	medium	V3-S11	Google's Ads API pag campaigns document group- and asset-leve placement view. Goog separately documents placement report: app website URL, video tit Google-owned invent as the only metric and placements grouped i
V3-F03	TikTok's official Smart+ App Campaigns help page states manual placement selection 'is not supported by Smart+ Campaigns at this time,' defaulting to automatic placement selection.	direct_record	SUPPORTED	high	V3-S01	
V3-F04	Amazon DSP's app-campaign events manager, announced Feb 13, 2024, is documented as available only for Android and Fire TV/tablet devices in the US, Canada, Mexico, and Brazil, with no iOS support stated.	direct_record	QUALIFIED	medium	V3-S10	Amazon's 13 Feb 2024 described DSP events support as available in Mexico and Brazil for TV/tablet, with expans September 2026 App closed Amazon DSP (K SKAdNetwork (third-p
V3-F05	Meta, Google, and TikTok each document a named conversion-lift/incrementality product with a randomized-controlled-trial design, but all three restrict access via account-team gating rather than self-serve availability.	direct_record	SUPPORTED	high	V3-S06 V3-S07 V3-S09	
V3-F06	Amazon Marketing Cloud is documented to provide advertisers with event-level and user-level data, including ad-attributed impressions, clicks, and conversions, across 'hundreds of fields,' but access for Amazon DSP advertisers requires a request to an Amazon Ad Tech Account Executive rather than being self-serve by default.	direct_record	QUALIFIED	medium	V3-S13	Amazon Marketing Cl query hundreds of ev attributed impressions conversions) inside a outputs are aggregate only. Amazon DSP ad request access throug Account Executive; sp advertisers self-serve
V3-F07	TikTok's official SKAN documentation (last updated February 2025) confirms full SKAN 4.0 support: fine-grained (6-bit) and coarse-grained conversion value schemas, three postback windows (0-2, 3-7, and 8-35 days), crowd anonymity tiers, and a four-digit source_id; no AdAttributionKit support is mentioned on that page.	direct_record	SUPPORTED	high	V3-S08	
V3-F08	Apple's own Ads Help documentation states Apple Ads registered with AdAttributionKit on April 10, 2025 (initially covering SKAdNetwork versions 1-3), operating alongside the pre-existing AdServices attribution API, which retains data for up to 21 days.	direct_record	SUPPORTED	high	V3-S02	
V3-F09	Adjust, a mobile measurement partner, has been a subsidiary of AppLovin Corporation since an acquisition announced in February 2021; AppLovin separately owns the MAX mediation platform and the AppDiscovery/Axon ad network, making Adjust an MMP owned by the same company that also sells mediation and demand-side ad products.	synthesis	QUALIFIED	medium	V3-S14 V3-S15	Adjust has been a wh subsidiary since 20 A cash plus \$352.0M in securities plus up to \$ AppLovin also sells M AppLovin Ads UA plat states Adjust data is n unless a customer dir
V3-F10	Statsig, an experimentation/feature-flagging platform, was acquired by OpenAI (reported by TechCrunch as occurring around September 2025) and then had its product/business acquired by Amplitude, Inc. on May 5, 2026, per Statsig's own blog post; the original Statsig founding team moved to OpenAI, and a new Amplitude-appointed leadership team now runs Statsig.	direct_record	CONTRADICTED	high	V3-S19 V3-S20	OpenAI acquired Stats 2025 (TechCrunch). U agreement, Amplitude customer contracts, tr perpetual non-exclusi technology (an asset i company); the Statsig under Amplitude, and at OpenAI.
V3-F11	data.ai (formerly App Annie) is now branded 'a Sensor Tower company' and its own domain states it 'Been Acquired by Sensor Tower'; the exact acquisition date could not be confirmed via a primary source opened in this research.	direct_record	QUALIFIED	high	V3-S16	data.ai (formerly App by Sensor Tower, ann 2024 (price undisclos operates as 'a Sensor
V3-F12	Kochava's marketing pages name specific incrementality-adjacent products (Marketing Mix Modeling described as 'Always-On Incremental Measurement,' 'Incrementality Testing,' and 'Media Lift Studies'), while Singular's and Branch's homepages/about pages do not name any comparable incrementality product.	vendor_assertion	QUALIFIED	low	V3-S23 V3-S22 V3-S24	Kochava names incre (MMM, Incrementality Studies). Singular doc Incrementality feature exposed test group w control group (random and Adjust documents synthetic control grou apps. Branch's about
V3-F13	TikTok's Value-Based Optimization (VBO) for Smart+ App Campaigns is documented as Android-only, meaning iOS advertisers on TikTok cannot access this value/ROAS-tier optimization goal through Smart+ as documented.	direct_record	QUALIFIED	high	V3-S01	TikTok's Smart+ App (value-based optimizat (iOS Smart+ bidding: I TikTok's standard app value-based optimizat and iOS.

V3-F14	All five benchmark channels profiled (Meta, Google, TikTok, Apple Ads, Amazon DSP) document a campaign-management API allowing programmatic creation/editing of campaigns and budgets/bids, not merely a reporting API — self-serve/API-level buyer control is documented as universal at this tier, even where automation (Advantage+, Smart+, AI Max) reduces the number of manual levers exposed.	synthesis	QUALIFIED	medium	V3-S11 V3-S12	All five benchmark channels document a campaign-management reporting API (excluding App Campaigns, Google Ads, Apple Ads, Meta (POS /act_id>/campaigns, objectives), TikTok (campaign_id), and DSP (self-serve campaign and budget management app-specific)).
--------	--	-----------	-----------	--------	------------------	---

D. Metric dictionary

Metric dictionary				
Metric	Unit	Definition	Can be added to	Watch out
Install	device event	A first launch after download that an attribution system records. Includes reinstalls and redownloads unless filtered.	Other installs of the same app and window only	Not a person and not a customer; reinstalls inside a reattribution window are often marked organic.
CPI (cost per install)	USD per install	Media spend divided by attributed installs.	n/a (a ratio)	Depends on attribution rules; cheap installs can churn faster.
CPA (cost per action)	USD per event	Media spend divided by an in-app event (registration, trial, purchase).	n/a	The event is a proxy; confirm it predicts value.
Media-only CAC	USD per acquired customer	Media spend divided by newly acquired customers (not installs, not reactivated users).	n/a	Excludes fees, creative, measurement and staff.
Fully loaded CAC	USD per acquired customer	Media plus platform and agency fees, creative production, measurement tools and allocated staff cost, divided by newly acquired customers.	n/a	The denominator must exclude reactivated and organic users.
Activation	share of installs	Share of installs reaching a defined first-value event within a stated window.	n/a	Define the event and the window; compare only like with like.
Retention (exact day)	share of cohort	Share of an install cohort active on exactly day N.	n/a	Lower than rolling retention; compare only at matching cohort age.
Retention (rolling or unbounded)	share of cohort	Share of a cohort active on day N or any later day.	n/a	Rises as later data arrives; incomplete cohorts are censored.
Churn	share of payers or subscribers per period	Share of paying users who cancel or lapse in a period; split voluntary from involuntary (payment failure).	n/a	Involuntary churn is a billing problem, not a media problem.
Gross bookings	USD	What users pay before store fees, refunds and taxes.	Other gross bookings	Not revenue to the developer.
Net revenue	USD	Gross bookings minus store or payment fees, refunds, chargebacks and sales taxes; plus ad revenue earned.	Other net revenue	State whether ad revenue is included.
Contribution	USD	Net revenue minus variable costs (cost of goods, payment costs, promotions, servers where material).	Other contribution	The basis for payback; revenue is not profit.
Observed LTV	USD per install or per customer	Cumulative contribution (or net revenue, stated) actually realized by a cohort up to its current age.	n/a	Always state cohort age.
Predicted LTV	USD per install or per customer	A model's forecast of cumulative value to a horizon.	n/a	A forecast; publish backtest error on mature holdout cohorts and drift by segment.
ROAS	ratio	Revenue attributed to a campaign divided by its cost. State revenue basis (gross, net, ad revenue, contribution) and cost basis (media only or fully loaded).	n/a	Attributed, not causal; the same campaign can show very different ROAS on different bases.
Incremental ROAS (IROAS)	ratio	Incremental revenue (treatment minus control, scaled) divided by incremental spend.	n/a	Undefined only when incremental spend is zero; zero or negative incremental revenue gives an IROAS of zero or below, which must be reported with its interval, not dropped.
Incremental CAC	USD per incremental customer	Spend divided by customers that would not have been acquired without it, estimated against a control.	n/a	Do not compute with zero or negative incremental customers.
Payback period	months	Cohort age at which cumulative contribution per customer first equals fully loaded CAC.	n/a	If payback falls in the predicted part of the curve, it depends on the forecast.
Postback (SKAN / AdAttributionKit)	one anonymous report per install, detail set by privacy thresholds	Apple's privacy-preserving attribution message: up to three per install, delayed, with detail limited by crowd-anonymity tiers.	Never to platform-reported conversions	Not a user; one winner per install.
Household (CTV)	household	A set of devices sharing an IP address or a graph-linked identity.	n/a	Exposure of a household is not attention by a person, and not the installer.

E. Policy and product status register

Policy and product status register
Official status of measurement frameworks, store rules and regulation at the cutoff. Existence of a standard is kept separate from live vendor support.

ITEM	OS	VERSION	REGION	ANNOUNCED	EFFECTIVE	STATUS	FUNCTIONS	IMPLEMENTATION EVIDENCE	SOURCE
App Tracking Transparency (standard alert)	iOS	iOS 14.0+	global	2020-06 (WWDC20)	2021-04-26	live	requestTrackingAuthorization(completionHandler); trackingAuthorizationStatus	developer.apple.com framework reference, GA (non-Beta)	C-S07
App Tracking Transparency (Beta expanded EU interface)	iOS	iOS/iPadOS/Mac Catalyst 27.2+	EU (full-page sheet mandatory in France, Germany, Italy, Poland, Romania; optional elsewhere in EU; no effect outside EU)	documented as Beta, exact announcement date not separately confirmed (found live on developer docs as of 2026-09-27)	not yet GA at cutoff	beta	full-page sheet, Markdown-formatted usage description (NSUserDefaultsTrackingMarkdownUsageDescription), optional 'Additional Information' button/callback	developer.apple.com API reference, marked Beta	C-S08
SKAdNetwork 4	iOS	iOS 16.1+	global	2022 (WWDC22)	iOS 16.1 release	live	3 conversion windows, coarse/fine conversion values, crowd anonymity tiers 0-3, hierarchical source identifier (2-4 digits), lock windows	developer.apple.com technical docs; AppsFlyer/Adjust/Singular all document live postback-forwarding support	C-S01 C-S02 C-S03 C-S19 C-S20 C-S21
AdAttributionKit	iOS	iOS 17.4+ (base); re-engagement iOS 18 cycle (June 2024); country-code/attribution-rules/cooldown iOS 26 cycle (June 2025)	global (App Store + alternative marketplaces)	2024-03 (iOS 17.4 release); expanded 2024-06 and 2025-06	iOS 17.4+ release; feature-specific dates per changelog	live	install + re-engagement conversions, postback bridging with SKAdNetwork, view-through impressions (SKAdImpression-equivalent), configurable attribution rules and cooldown windows, country-code in postbacks	developer.apple.com framework + changelog; AppsFlyer/Adjust/Singular document live support incl. Singular's dedicated re-engagement opt-in key	C-S04 C-S05 C-S06 C-S19 C-S20 C-S21
Apple Ads Attribution API (AdServices)	iOS	iOS 14.3+	global	2020-12	iOS 14.3 release; view-through added 2025-03-27; pre-order attribution added 2025-10	live	click/tap-through attribution (30-day window), view-through attribution (added March 2025), pre-order campaign attribution, supplyPlacement field (Jan 2026), attribution=false for age/gender-targeted campaigns (Sept 2026)	developer.apple.com framework + dated changelog through September 2026	C-S09 C-S10
Privacy Sandbox on Android (Attribution Reporting, Topics, Protected Audience, SDK Runtime, On-Device Personalization, Protected App Signals)	Android	n/a (platform-level API, not tied to one Android version in reviewed docs)	global	2025-10-17 (retirement announcement); original APIs announced 2022	not completed at cutoff -- status page shows 'Scheduled for phaseout' with no stated removal date	announced	n/a (being wound down)	privacysandbox.google.com official status page and blog announcement	C-S14 C-S15
Google Play Install Referrer API	Android	Google Play app 8.3.73+	global	n/a (long-standing API)	current	live	referrer URL, click/install timestamps (client+server), first-install app version, instant-experience interaction (7-day)	developer.android.com official docs	C-S17
Android Advertising ID (AAID) / AD_ID permission	Android	Android 12+ (opt-out/reset); Android 13+ targeting requires AD_ID manifest permission	global	2021 (late 2021 rollout on Android 12)	2022-04-01 (all Play-supported devices)	live	device-level identifier with user reset (returns all-zero string on reset); com.google.android.gms.permission.AD_ID manifest declaration required for Android 13+ targets	support.google.com policy page; developer.android.com AAID reference	C-S18 C-S16
France ATT competition decision (25-D-02)	iOS	n/a (regulatory)	France	2025-03-31	2025-03-31 (fine imposed); no Apple product-change deadline stated in the decision text reviewed	unconfirmed (appeal status not found)	n/a	autoritedelaconurrence.fr official press release	C-S11
Italy ATT competition decision (A561)	iOS	n/a (regulatory)	Italy	2025-12-22	2025-12-22 (fine imposed)	unconfirmed (Apple reportedly intends to appeal per Reuters; outcome not in official record reviewed)	n/a	en.agcm.it official press release	C-S13
Germany ATT competition decision (Bundeskartellamt, Sec. 19a GWB)	iOS	n/a (regulatory)	Germany	proceeding opened 2022-06; preliminary assessment 2025-02; decision 2026-08-17	commitments legally binding as of 2026-08-17; Apple has 4 months from service to implement (not yet reached at cutoff)	proposed (binding commitment, implementation pending)	aligned consent-prompt wording/design between first- and third-party apps; removal of discouraging symbols/wording; more room to combine ATT consent with GDPR consent	bundeskartellamt.de official press release	C-S12

FTC Negative Option ('Click-to-Cancel') Rule	n/a	16 C.F.R. Part 425	US	2024-10-16 (final rule)	vacated 2025-07-08	retired	n/a	Eighth Circuit vacated the rule in full (Custom Communications, Inc. v. FTC); FTC subsequently opened an Advance Notice of Proposed Rulemaking (Mar 2026) to redo it; no click-to-cancel rule was in force under this rulemaking as of 2026-09-27	G3-S18 G3-S19
EU DMA Article 6(8) — advertiser/publisher performance-measurement data access	n/a	Regulation (EU) 2022/1925	EU	2022-09-14 (adopted); published in the Official Journal 2022-10-12	in force 2022-11-01; applies from 2023-05-02 (Art. 54); each gatekeeper must comply within six months of designation (Art. 3(10))	live	Free access to gatekeeper's own performance-measuring tools and aggregated/non-aggregated verification data, on request	Statutory text confirmed directly from EUR-Lex; applies only to core platform services for which a company is actually designated a gatekeeper	G3-S12
EU DMA Article 6(10) — business-user data access/portability	n/a	Regulation (EU) 2022/1925	EU	2022-09-14 (adopted); published in the Official Journal 2022-10-12	in force 2022-11-01; applies from 2023-05-02 (Art. 54); each gatekeeper must comply within six months of designation (Art. 3(10))	live	Continuous, real-time access to aggregated/non-aggregated (incl. conditionally personal) data generated by business users' and end users' activity, free of charge, on request	Statutory text confirmed directly from EUR-Lex	G3-S12
Apple device-fingerprinting prohibition	iOS	Apple Developer Program License Agreement	global		current policy as of access date	live	n/a (prohibition, not a feature)	Verbatim official FAQ text re-verified by direct fetch	G3-S13
AdCP governance protocol (check_governance / sync_governance, three-party model)	n/a	AdCP docs v3.1.24 (root site tagged release v3.2.0-rc.3 at cutoff)	global		documented as current at access date	live (documented spec; independent verification of production-scale enforcement not established in this pass)	Plan registration, pre-spend policy checks with must/should/may severity, async human-approval escalation above authority thresholds, audit logging	Primary technical documentation, quote-verified	G3-S16
IAB Tech Lab AAMP governance (human-in-the-loop approval gates)	n/a	AAMP 2.0 (AAMP 3.0 'well underway' per Aug 2026 update)	global	2026-04-23	2026-04-23 (2.0); 3.0 not yet released per this source	live (2.0 SDKs available for download; 3.0 in progress)	Configurable human approval gates, notifications, audit logging for buyer/seller agent transactions	Standards body's own announcement, quote-verified directly	G3-S17
COPPA Rule 2025 amendments	n/a	n/a	US	2025-01	2025-06-23	live	Separate opt-in consent for third-party disclosure/targeted advertising (per secondary summaries; not independently confirmed against primary text), data retention schedule, written information security program, updated safe harbor requirements	FTC landing page confirms rule and a Feb 25, 2026 companion policy statement on age-verification technologies	KL-S17 KL-S18
FTC order: X-Mode/Outlogic	n/a	n/a	US	2024-01	2024-04	live	Ban on selling/sharing sensitive location data; deletion of prior data/derived products	Official FTC press release (not directly opened in this research)	KL-S20
FTC order: InMarket	n/a	n/a	US	2024-01	2024-05-01	live	Ban on selling/licensing precise or sensitive-location-categorized consumer data	Official FTC press release (not directly opened in this research)	KL-S21
FTC order: Kochava / Collective Data Solutions	n/a	n/a	US	2026-05-04	2026-06-25 (stipulated order entered, Dkt. 138; FTC posted 2026-06-26)	live: stipulated order entered 2026-06-25	Bars sale, licensing, transfer or disclosure of sensitive location data; narrow exception needs a direct consumer relationship, express consent and a requested service; defendants neither admit nor deny	Entered court order (OA2-S01); FTC timeline item (OA2-S02)	KL-S19 OA2-S01 OA2-S02
Utah App Store Accountability Act	iOS/Android	n/a	US-Utah	2025	2027-05-06 (HB 498, signed 2026-03-18; state enforcement removed)	enacted; delayed (only harmed minors or parents may sue)	Age verification and parental consent obligations on app stores; AG enforcement removed by HB498, private right of action retained	Law-firm alerts corroborating HB498 amendment and CCIA dismissal (bill text itself not retrievable via WebFetch in this research)	KL-S22 KL-S23
Texas App Store Accountability Act (SB2420)	iOS/Android	n/a	US-Texas	2025-05-27 (signed)	2026-01-01 (statutory; confirmed via official Texas Legislature record)	live: in force after the 5th Cir. stayed the district court's injunctions (28 May 2026 administratively; 4 Jun 2026 pending appeal); argued 4 Aug 2026; no merits ruling at cutoff	Age category verification via app store; parental account linkage and consent for minors before app download or in-app purchase	Official Texas Legislature bill-history page (KL-S36) confirms signed/effective dates; Fifth Circuit's June 4, 2026 stay of the district court's preliminary injunction (which had blocked the Jan 1, 2026 effective date) reported via law-firm alert (KL-S24), not independently confirmed against the court order itself	KL-S36 KL-S24

Louisiana app store age-verification law (HB977, Act No. 185)	iOS/Android	n/a	US-Louisiana	2026 (signed by Governor)	Act 185 (signed 15 May 2026) repealed the 2025 law on signature; new app-store duties from 2027-07-01	enacted; not yet in force	Developer must verify minor status via app-store data-sharing method and obtain verifiable parental consent before download/purchase	Official Act text (FX2-S18)	KL-S25 KL-S25b
California Age-Appropriate Design Code Act (AB 2273 / AADC)	n/a	n/a	US-California	2022-09-15 (chaptered as Chapter 320)	2024-07-01 statutory operative date (confirmed via official CA Legislature record); substantively enjoined in part per 9th Cir.	unconfirmed (statute operative but core provisions enjoined)	Statute requires Data Protection Impact Assessments and other design-code obligations from July 1, 2024; data-use restrictions and dark-patterns prohibition remain enjoined per 9th Circuit; age-estimation/coverage-definition provision remanded for further proceedings	Official CA Legislature bill-status page (KL-S37) confirms enactment/operative dates but does not reference the litigation; two independent law-firm alerts describe the March 12, 2026 Ninth Circuit ruling	KL-S37 KL-S26
Apple Declared Age Range API	iOS	iOS/iPadOS/macOS 26.0+ (age-check method from 26.2; regulatory features from 26.4)	global/US	2025-2026 (WWDC)	unconfirmed	live	Requests a person's declared age range for age-appropriate app experiences; referenced by developers for Texas SB2420 compliance (forum-thread title only, not independently verified)	Apple documentation JSON (FX2-S19)	KL-S27
Google Play Age Signals API	Android	beta	Brazil, US-Texas (expanding globally by end of 2026 per secondary reporting)	2026	2026-03-17 (Brazil); 2026-05-28 (Texas)	beta	Returns age-range signal (default bands 0-12, 13-15, 16-17, 18+); explicitly barred from ad/marketing/profiling/analytics use	Directly fetched official Android developer documentation	KL-S29
Apple App Store Review Guidelines 5.1.1/5.1.2/5.1.4	iOS	current	global	ongoing	current as of access date	live	Consent/disclosure for data collection; ATT requirement for tracking; profiling ban; sensitive-API marketing ban; kids-category third-party ad/analytics restriction	Directly fetched and quoted official guideline text	KL-S28
Google Play Data safety & Families policy (ads)	Android	current	global	ongoing	current as of access date	live	Families Self-Certified Ads SDK requirement; ban on interest-based/remarketing ads to children; ad-format and content restrictions	Directly fetched and quoted official policy text	KL-S30
Apple Privacy Manifest / Required Reason APIs	iOS	current	global	2023	2024-05 (new apps); 2024-08 (updates) — medium confidence	live	Mandatory PrivacyInfo.xcprivacy declarations for SDKs and Required Reason API usage; non-compliant apps blocked at submission	One fetch returned detailed content; not independently cross-checked against a second source in this research (see reliability note on KL-S31)	KL-S31
EU DMA — Apple EU business terms restructure	iOS	n/a	EU	2026-08-18	2026-10-01 (after the cutoff)	announced; not in force at cutoff	App Store: 26% via Apple IAP, 20% other in-app payment, 15% link-out within seven days; reduced rates for small businesses and year-two subscriptions; 5% core technology commission only for apps distributed outside the App Store. At the cutoff, the opt-in Alternative Terms Addendum (17%/10% + 3% payment fee, EUR 0.50 per first annual install above 1m) applied	Apple EU support page (GH-S03); Alternative Terms Addendum (OA2-S03)	KL-S32 GH-S03 OA2-S03
EU DMA — Apple app-distribution/sideloading compliance	iOS	n/a	EU	2024	unresolved	proposed	European Commission finding that Apple's sideloading-compliance mechanism does not satisfy DMA obligations; investigation ongoing at cutoff per secondary reporting	Not directly opened; press-release title and secondary reporting	KL-S33

F. Benchmark and study table

Benchmark and study table												
Design, sponsor, sample and limits for the causal studies used. App-specific vendor holdouts are in the evidence ledger (G1).												
STUDY	AUTHORS	YEAR	SPONSOR	DESIGN	UNIT	SAMPLE	DURATION	CATEGORY	GEO	OS	OUTCOME	COMMENTS
A Comparison of Approaches to Advertising Measurement: Evidence from Big Field Experiments at Facebook	Gordon, Zettelmeyer, Bhargava, Chapsky	2019	Facebook (two co-authors were Facebook employees; academic leads at Northwestern/NBER)	RCT vs. multiple observational models	user	15 US ad experiments; ~500M user-experiment observations; 1.6B ad impressions	not stated in sections read	mixed (E-commerce, Retail, Travel, Entertainment/Media per the related 2023 paper's description of comparable experiments)	US	n/a (Facebook feed ads, not mobile app installs)	Whether observational methods recover RCT-measured causal ad effects	Randomized control

Close Enough? A Large-Scale Exploration of Non-Experimental Approaches to Advertising Measurement	Gordon, Moakler, Zettelmeyer	2023	Meta (one co-author is Meta Ads Research staff; academic leads at Northwestern/NBER)	RCT compared against double/debiased machine learning (DML) and stratified propensity-score matching (SPSM)	user	563 US Facebook experiments (Nov 2019 to Mar 2020) giving 663 test-and-control pairs; ~38 billion impressions	November 2019-March 2020	mixed (E-commerce, Retail, Travel, Entertainment/Media)	US	n/a	Accuracy of DML/SPSM in recovering RCT-measured causal lift across purchase-funnel stages	Randomized control within 663 e
The Unfavorable Economics of Measuring the Returns to Advertising	Lewis, Rao	2015	Yahoo! Inc. (authors' employer at time of study)	RCT (randomized holdout of targeted users from display ad exposure)	user	25 field experiments; 19 retailers + 6 financial-service firms; \$2.8M combined ad spend	not stated precisely in sections read	non-gaming (retail, financial services)	US	n/a	Precision of experimentally-measured advertising ROI	Randomized control
Consumer Heterogeneity and Paid Search Effectiveness: A Large-Scale Field Experiment	Blake, Nosko, Tadelis	2015	eBay (all three authors affiliated with eBay Research Labs)	Geo-based RCT (DMA-level ad suspension)	Nielsen Designated Market Area (DMA)	68 test DMAs (ads suspended) vs. 142 control DMAs, out of 210 US DMAs total	60 days	non-gaming (e-commerce marketplace)	US	n/a (desktop/paid search, not mobile app)	Causal ROI of paid search advertising	DMA-level on
Evaluating the Impact of Privacy Regulation on E-Commerce Firms: Evidence from Apple's App Tracking Transparency	Aridor, Che, Hollenbeck, Kaiser, McCarthy	2025	Academic funding (MSI, UCLA, LEC); revenue panel from Grips Intelligence, a firm a co-author works with	Event study + difference-in-differences	firm (e-commerce merchant)	Opt-in panels from an anonymous ad-analytics provider and Grips Intelligence; benchmarked against Kantar Vivivix, Shopify and SimilarWeb	Pre-ATT baseline April 2020-April 2021; panels span 2019-2022	commerce	US	iOS	Meta ad click-through rate and firm-wide e-commerce revenue after ATT	Firms' baseline spend dependent
Estimating the Value of Offsite Data to Advertisers on Meta	Wernerfelt, Tuchman, Shapiro, Moakler	2025 (Marketing Science; 2022 draft superseded)	Meta: two of four co-authors were Meta staff at the time; Meta could review for proprietary data but had no right to restrict publication on the results; experiment run on Meta's platform	Large-scale in-platform experiment (offsite-data signal removed for treated advertisers/campaigns)	advertiser/campaign	More than 70,000 advertisers; one-week randomized experiment, fall 2021	Experiment period plus 6-month post-experiment purchase tracking	all (noted especially CPG, Retail, E-commerce)	not specified in sections read	both	Cost per incremental customer, with vs. without offsite targeting data	Business offsite target simulation offsite
ATT vs. Personalized Ads: User's Data Sharing Choices Under Apple's Divergent Consent Strategies	Baviskar, Chowdhury, Li, Sokol	2024	Not funded by Meta per author disclosure, but two of five co-authors are Meta employees with a financial interest in Meta	Randomized survey experiment (prompt-framing manipulation)	individual survey respondent	11,000 US and UK online adults	not stated precisely in sections read	all (tested across social media, news, delivery, CPG app contexts)	US, UK	iOS	Stated data-sharing opt-in rate under ATT-style prompt vs. Apple's own Personalized Ads prompt	Within subject; subject comparison two p
The Impact of Apple's App Tracking Transparency on App Monetization	Kesler	2022/2023	not stated in abstract-level material read	Difference-in-differences (Apple vs. Google Play as comparison)	app	580,000+ apps (per search-engine abstract summary)	before/after April 2021 ATT introduction; exact window not confirmed	all	not confirmed	both	In-app payment / paid-app adoption	Google (not s ATT)
Ghost Ads: Improving the Economics of Measuring Online Ad Effectiveness	Johnson, Lewis, Nubbemeyer	2017	not stated in abstract-level material read	Method paper proposing 'ghost ads' as an alternative to PSA/intent-to-treat A/B testing within real-time ad-delivery/auction systems	user (within the ad-delivery system)	not stated in abstract	not stated in abstract	all	not stated in abstract	n/a	Cost and precision of ad-effectiveness measurement	PSA (service announcement intent test

G. Author corpus map

Author corpus map

The author's own prior work reviewed for this paper. Dispositions: retain, update, test, qualify or reject. Prior claims are hypotheses, never corroboration.

ID	TITLE	TYPE	PUBLISHED	REVISED	RELEVANT POSITIONS	RELEVANCE	DISPOSITION	REASON
CP-01	Mobile App Growth Playbook	playbook	2026-06-05	2026-06-19	Mobile DSP is true but too small; 'Apps DSP' (install+engagement+retention+commerce+CTV+LTV decision layer) is the better buyer-facing category frame; AppLovin sold its mobile-gaming studios to Tripdot Studios (announced May 2025, closed June 2025) and is now a single-segment ad-tech public company (Nasdaq: APP).; On 1 Oct 2025 AppLovin rebranded its self-serve ad platform as 'Axon' (Axon Ads M	Q1,Q2,Q5,Q6,Q9,Q10,Q11,Q12; chapters 1 (Apps DSP framing), 9 (CTV/commerce), 11 (buyer control)	TEST	Marketing/GTM collate authored by the paper's author selling advisory this exact buyer set; 'validated mid-2026' is months stale by the Sep 2026 cutoff and the Chrome-vs-Android Privacy Sandbox claim needs primary-source verification
CP-02	App DSP Landscape Matrix — data file	data	2026-06-06	2026-06-08	File's own governing rule: do not fabricate — capability fields default to 'validate'/'unknown' rather than asserting an unverified capability as 'strong'; AppLovin: Tripdot divestiture (~\$400M cash + ~20% Tripdot equity, ~\$800M total transaction value, not all-cash); post-divestiture reports as single advertising segment; Axon Ads Manager self-serve launched Oct 2025.; Adjust is listed repea	Q5,Q6,Q11; V1/V2 vendor-universe rows for Stream V1/V2 profiles	UPDATE	Structurally disciplined (evidence-graded, non-fabricating) but stale: AppLovin's Aug 2026 C print (revenue \$1.92B, EBITDA margin, per the author's own later essay postdates this file's Jun 2026 validation and is reflected.
CP-03	Video & Mobile Ad Delivery Standards	standards	2026-06-11	2026-09-08	IDFA is gated by ATT (available iOS/iPadOS 14.0+); never assume availability — it is a permission outcome, not a device property.; SKAdNetwork current documented version is 4: up to three conversion windows from iOS 16.1, up to three postbacks, a winning postback plus up to five runner-ups; two older SKAN methods are deprecated but the framework remains documented.; AdAttributionKit (from iOS 17.4	Q2,Q3,Q7; Stream C (measurement matrix)	RETAIN	Internally consistent, di (last validated Sept 202 and appropriately hedge treat as a hypothesis pointing to Apple/Google primary docs rather than the primary source itself brief's source hierarchy
CP-04	Privacy & Consent Standards: GPP, TCF, SKAN & Platform APIs	standards	2026-06-11	2026-09-08	Frames the Oct 2025 Privacy Sandbox retirement (Topics, Protected Audience, Attribution Reporting) primarily as a CHROME/browser event, following April 2025 Chrome cookie U-turn and Nov 2025 Intent-to-Deprecate threads (Chrome 144 deprecation, Chrome 150/153 removal milestones) — does not equate this with an Android OS-level retirement.; ATT (iOS/iPadOS 14.0+) has 'no deprecation notes in public d	Q3,Q7; Stream C policy_status register — directly bears on CP-01's Android Privacy Sandbox claim	TEST	This page's own Chrome-centric framing appears conflict with CP-01's br 'Chrome and Android' retirement claim — veri against Google's official 2025 'update on plans Privacy Sandbox technologies' and developer.android.com before either claim is u
CP-05	Measurement, Verification & Media Quality: MRC, IVT, OM SDK	standards	2026-06-11	2026-09-08	MRC's Invalid Traffic Detection and Filtration Guidelines (issued Oct 2015, addendum Jun 2020, interim updates Apr 2024) split invalid traffic into GIVT (list/pattern-detectable) and SIVT (requires advanced analytics/human review, e.g. app-ID spoofing, domain laundering).; OM SDK is a data-collection standard, not a fraud-detection product; it exposes data for third-party verification vendors but	Q3,Q7; Stream IJ (fraud/quality registers)	RETAIN	Standards-doc summary with appropriate hedge (explicitly disclaims accreditation verification as a pointer to MRC Tech Lab primaries rather than as the evidence its
CP-06	CTV, Streaming & Live Event Advertising Standards	standards	2026-06-12	2026-09-08	Person-level reach/frequency/co-viewing on CTV is typically modeled from household-level signals, not observed — treat as methodology-bearing estimates, not raw measurement.; IAB Tech Lab CTV Ad Format Portfolio (6 formats) finalized July 22, 2026 after a second public-comment round.; Live Event Ad Playbook's Concurrent Streams API v1.0 is the one shipped LEAP component as of mid-2026; Forecasting	Q9; Stream IJ ctv_bridge register — CTV-to-app measurement-bridge chapter	RETAIN	Directly supports the p CTV-to-app bridge cha household-vs-person distinction; the household modeling caveat should applied verbatim to any CTV-to-app case evidence including Samba TV (author's current empl — apply stricter scrutiny brief).
CP-07	IAB Incrementality Guidelines Decoded	standards	2026-07-12	2026-09-08	Decodes the joint IAB/IAB Europe 'Guidelines for Incremental Measurement in Commerce Media' (final Nov 3, 2025): four method families graded by causal strength — experiment-based (strongest), model-based counterfactual, econometric (MMM), hybrid proxies (weakest).; Platform-reported incrementality is explicitly classified by the guidelines themselves as a weak 'hybrid proxy' relying on platform-sp	Q1,Q3; central methodology framework for the whole paper's incrementality-vs-attribution chapter	RETAIN	Strong primary-adjace decode with an explicit sponsor/interest caveat already applied to its headline statistic; the 6.5x/83% figure is retail media-specific (not mobile-app-native) and must be scoped as an analogy, direct mobile-app evidence when reused.
CP-08	Retail & Commerce Media Measurement	standards	2026-06-12	2026-09-08	GPP consent metadata (gpp_string, gpp_sid) and an mmt_only flag mark events for measurement-only use, excluded from ads-delivery optimization — schema does not itself constitute a legal compliance opinion.; Closed-loop attribution ties credit for sales to a retailer's own transaction data; MRC's viewable-impression requirement for attributed outcomes is explicitly noted as something 'most retail m	Q9,Q10; commerce-app expansion chapter (marketplace/retail/QSR apps in CP-01)	RETAIN	Useful for the commerce-app-growth sections b retail-media-network-centric, not app-install-centric — scope claims app commerce/market use cases only, not general mobile UA.
CP-09	Gaming Playbook	playbook	2026-06-05	2026-09-17	Mobile gaming remains the strongest performance-ad heritage channel, but measurement now runs through SKAN/AdAttributionKit, MMP, incrementality and MMM thinking rather than last-click installs.; Cites (vendor-attributed) that ~3.4B+ people play games globally and ~59% of the US population are gamers (eMarketer, 2026), yet gaming draws under 5% of worldwide media investment.; Frames gaming measure	Q2,Q9; gaming-vertical segment required by the brief's 'gaming vs non-gaming' segmentation rule	UPDATE	The 3.4B/59%/5%-of-s figures are asserted with an in-page source local must be traced to the eMarketer 2026 public and geography/definition confirmed before use a market-size data point.

CP#	Title	Type	Published	Reviewed	Summary	Key Findings	Status	Notes
CP-10	Performance Playbook (native/recommendation/commerce)	playbook	2026-06-05	2026-06-19	Explicit weak-vs-better-vs-strong evidence ladder per claim type: 'drives conversions' (last-click platform reporting, weak) → MMP/site-analytics reconciliation (better) → incrementality test + BI view (strongest); Frames CTV/cross-screen native's core measurement risk as identity, dedupe, and measurement-bridge quality rather than creative or reach; Positions retail/commerce media's different	Q3,Q9; supports the paper's evidence-hierarchy framework for grading vendor performance claims generally (not mobile-app-specific)	QUALIFY	Adjacent GTM playbook (native/content-discovery not app-install DSPs); useful only for its generic evidence-ladder framework not as mobile-app-growth market evidence.
CP-11	Outcome Underwriting Playbook	playbook	2026-09-12	2026-09-12	States plainly it is 'not an incrementality methodology' — every counterfactual-estimation method is itself a negotiated contract term, not a neutral fact, so the method belongs in the contract, not an appendix; Frames a 'measurement' contract clause as 'who counts it, and with what access' — the current state described as 'partly built' (attribution, holdouts, incrementality methods each indepen	Q1,Q3,Q10; underpins the paper's framing that 'a lower CPI/higher ROAS is not proof of profitable growth' by making measurement-choice a contractual/negotiated variable	QUALIFY	Conceptual/contractual framework (guaranteed outcome deals), not an growth-specific data set; useful as a governance for the paper's Q1/Q3 framing only.
CP-12	Measurement Governance Playbook	playbook	2026-09-12	2026-09-12	Distinguishes what 'earns' trust (a holdout/incrementality test that survived a counterfactual) from what merely 'defends' a number (correlation wearing a causal label) — ties directly to what justifies moving budget between lines.	Q1,Q3,Q7; governance framework for reconciling attribution/MMP/experiments/finance (Stream D)	QUALIFY	Generic governance framework, not app-growth specific; use only as a conceptual scaffold for paper's reconciliation chapter.
CP-13	DSP / Agentic Buying — Ecosystem Surface Deep Dive	playbook	2026-06-05	2026-09-13	DSPs are evolving from bid-optimizing consoles into 'decision engines' that interpret goals, evaluate supply paths, use curated inventory, and increasingly automate setup/bidding/optimization/measurement; Notes market consolidation: Microsoft Invest (ex-Xandr) stopped supporting media buying Feb 28 2026; MediaMath now trades as Infillion MediaMath post-bankruptcy — independent DSP field is consol	Q11,Q12; Apps DSP category-definition chapter and AI/agentic-workflow chapter	QUALIFY	Open-web/CTV DSP surface, not an in-app/mobile DSP deep dive; useful only for the general agentic-DSP evolution framework the mobile paper's 'Apps DSP' argument leans on by analogy.
CP-14	BI / MMM / Decision Intelligence — Ecosystem Surface Deep Dive	playbook	2026-06-05	2026-09-13	Core thesis: 'BI describes the past; MMM, experiments, and decision intelligence turn signals into budget decisions' — the constraint is comparable metric definitions and calibration, not more dashboards; MMM 'revived' specifically because privacy changes weakened user-level tracking; positions MMM as causal, aggregate, and durable to identifier loss — directly analogous to the mobile paper's 'MM	Q3,Q4,Q12; MMM-for-apps and LTV-prediction chapters	QUALIFY	General cross-industry BI/MMM framework, not app-specific; use only to support the paper's MMM-for-apps definitional section not as app-growth market evidence.
CP-15	iROAS Is Not a Number, It's a Negotiation	essay	2026-07-13	2026-07-14	Reports the Ovative/Albertsons/Kellogg finding (42 campaigns, 54 methodology combinations) that within-campaign iROAS varied 6.5x on average (median 2.5x), with 83% of campaigns flipping sign purely on methodology choice, no change to the underlying campaign; Frames iROAS's numerator (what counts as incremental) and denominator scope (whose spend against which audience) as the two loci of methodo	Q1,Q3; same underlying study as CP-07, retail-media context reused as methodology cautionary tale	QUALIFY	Same primary study as CP-07 (IAB Incrementality Guidelines page); retail display context, not mobile app-install context — clarifies the underlying study or not twice, and scope explicitly as an analogy app-side iROAS/incrementality c
CP-16	One Event, Three Machines: Orchestrating Conversions Across PMax, Advantage+, and OpenAI Ads	essay	2026-07-23	2026-08-09	Describes a within-platform deduplication pattern (single event_id minted on form submit, carried to both server-side and browser-side legs) so each platform independently sees one deduplicated conversion rather than double-counting browser+server events; Distinguishes within-platform identity (shared event_id) from cross-platform identity (each walled garden's own click ID captured into a record	Q3,Q7; directly transferable to mobile MMP/SKAN dedup logic (browser-based lead-gen example, not app-install, but same event-identity architecture problem)	QUALIFY	Web lead-gen (form submission architecture, not mobile install/SKAN); the deduplication identity framework is transferable by analogy; MMP/SKAN reconciliation but is not itself mobile-evidence.
CP-17	The Loop Closed Inside the Wall	essay	2026-07-21	2026-08-16	Analyzes Snowflake's July 21 2026 Meta-integration blueprint (Conversions API skill + Meta ads MCP) as a governed agentic loop where an agent prepares actions and a human approves before money moves; Argues the loop is 'identity-only' (hashed-PIL match against one walled garden's identity spine) with no contextual lane, and that feeding richer conversion signals into a platform's own attributed o	Q7,Q12; governance/agentic-loop chapter, directly relevant to whether AI/agentic workflows in app-growth DSPs (Axon Ads Manager, Advantage+ app campaigns) are causal-evidence-producing or just cleaner attribution	QUALIFY	Case study is Snowflake's Meta web/CRM conversion; not a mobile app-install use only as an analogy evaluating whether app-growth agentic feature (AppLovin Axon, Google campaigns AI Max) run holdouts or merely improve attribution hygiene.
CP-18	Signal Containerization: The Next Abstraction Layer for Agentic Advertising	essay	2026-06-07	2026-08-06	Proposes 'signal container' as a portable, governed decision object (intent, semantic meaning, source data, methodology, privacy rules, activation path, identity/contextual logic, evaluation method, allowed outputs, audit trail) replacing a plain audience segment; Names two IAB Tech Lab specifications (ARTF — Agentic Real-Time Framework — and an embeddings-based Agentic Audiences concept) as the	Q11,Q12; the framework underlying the paper's treatment of what 'genuinely agentic' app-growth buying would need beyond automated bidding	QUALIFY	General open-web/exchange-level framework (segments, bidstream), not mobile-DSP-specific; author could lead AdCP's Signals & Measurement WG — please brief, treat any AdCP-adjacent claims here neutrally, no advocacy.
CP-19	Nobody Sells an Outcome	essay	2026-08-21	2026-08-28	Reports an informal reader poll (73 ballots, LinkedIn 'Friday Thought Experiment'): 57% picked 'Outcomes' as what autonomous buying agents will ultimately buy, vs 27% Impressions, 8% Audiences, 6% Attention — explicitly labeled as an opinion poll, not market data; Argues 'outcome optimization' already exists inside walled gardens but 'outcome insurance' (a warranted, priced, capital-backed guaran	Q1,Q10,Q11; conceptual scaffold for evaluating whether app-growth 'outcome-based' pricing claims (CPA/CPI/ROAS optimization) are actually warranted outcomes	QUALIFY	Evidence class is an informal LinkedIn poll (illustrative_scenario/or not measurement); do not cite the 57/27/8/6 split market data — use only underlying conceptual distinction (optimization warranted outcome).

CP#	Title	Type	Published	Updated	Summary	Key Findings	Status	Notes
CP-20	The Risk You Can Price	essay	2026-08-23	2026-08-26	Cites Criteo's own Dec 2017 SEC filing estimating Apple's Safari ITP-style change cost Criteo about 22% of the following year's revenue ex-TAC — and that no insurance claim existed because the platform-dependency risk had never been priced or collateralized.; Argues the first real 'performance insurance' for open-web/app advertising would protect the sell side (publisher receivable, intermediary s	Q5,Q11; relevant precedent for evaluating platform-dependency risk in app-growth vendors reliant on Apple/Google measurement changes (ATT 2021, SKAN/AAK, Privacy Sandbox)	QUALIFY	Open-web display/Criteo case study, not a mobile DSP; useful as a historical precedent for platform-dependency risk (directly analogous to ATT's 2021 impact on app-install advertisers) but not its app-growth market evidence.
CP-21	The Open Web Isn't Dead. It's Uninsured.	essay	2026-08-22	2026-08-24	Cites AppLovin's Q2 2026 SEC filing/press release: revenue \$1,923.7M (up 53% YoY), 84% adjusted-EBITDA margin, but landed about \$1M under its own adjusted-EBITDA guidance floor and fell 19.66% the next trading day.; States AppLovin's 10-K discloses no fixed price per action, and that Meta's own documentation states its cost cap is not guaranteed — used to argue that even the most 'closed-loop' app	Q5,Q11; AppLovin financial-scale and buyer-control evidence directly supersedes CP-02's June-2026-dated figures	UPDATE	Highest-value item in the corpus for Stream AB's AppLovin filings require — pulls Q2 2026 10-Q-adjacent figures newer app-dsp-landscape.ts's June 2026 validation; v the cited SEC exhibit U directly rather than rely on the essay's transcrip
CP-22	The CMO Owns the Action Space (Post-Agentic Marketing, Part 4)	essay	2026-07-05	2026-08-16	Argues autonomy is already real in decisioning (not spending) — i.e., automated bidding/targeting decisions run without per-decision human sign-off, but budget-level actions still require an approval gate.; Frames the post-agentic CMO's core deliverable as 'guardrail architecture' (a brand constitution + procurement test) rather than a channel plan — measurement's job shifts from explaining the pa	Q12; governance/procurement checklist directly reusable for the paper's AI-features register (Stream KL)	QUALIFY	Cross-channel marketing governance framework app-growth/mobile-DS specific; the procurement test structure is reusable template but its content not itself app-growth evidence.
CP-23	The Mandate Finished Last	essay	2026-08-14	2026-09-15	Reports an informal reader poll (42 ballots): 78% picked 'proprietary data' as the durable edge once every advertiser has a capable buying agent, vs 9% faster execution, 7% privileged access, 4% clearer mandate — again explicitly an opinion poll, not market data.; Author's own reflection concedes he could name, for any past client, which data assets they held that competitors lacked, but could not	Q5,Q12; compounding-advantage chapter (which advantages compound for scaled platforms — event data vs. workflow integration vs. governance)	QUALIFY	Informal opinion poll (illustrative_scenario), measured market evidence usable only for its proprietary-data-vs-verifiability argument, not a data point on what compounds for app-growth platforms.
CP-24	From Meridian to NNN: How Transformers Are Redefining Marketing Mix Modeling	essay	2025-04-21	2026-08-30	Describes Google Meridian (launched 2024) as a Bayesian MCMC MMM using scalar spend/impression inputs with Adstock (lag) and Hill (saturation) parametric functions.; Describes Google's NNN (2025) as a Transformer-based successor using high-dimensional embeddings blending quantitative spend with qualitative creative/query attributes, claimed (early benchmarks) to outperform prior MMM by ~22% in pre	Q3,Q4; MMM-for-apps chapter (saturation, seasonality, organic-lift reconciliation)	TEST	Pre-dates the research window by over a year (originally an April 2021 LinkedIn post) and the accuracy-lift figure is unsecured in the review text — trace to Google's NNN documentation/p before using the numb
CP-25	Debunking Cross-Device Myth	essay	2015-08-13	2026-06-09	Explains deterministic (login-based) vs. probabilistic (statistical-pattern) device-graph matching as the two foundational cross-device identity methods.	Q9 (identity-bridge background only, e.g. CTV household/device matching precedent)	REJECT	Written in 2015, over a decade before ATT, SKAN, GDPR/CCPA and the AI identifier changes that define this exact problem; the deterministic/probabilistic framing is directionally valid but every specific is superseded — do not for current CTV-to-app identity-bridge evidence
CP-26	Is Apple Harvesting Adtech Data?	essay	2023-12-15	2026-07-06	Reacts to a 2023 Digiday video framing Apple's privacy posture (App Tracking Transparency, privacy manifests) as itself a data-collection/competitive-intelligence strategy, listing the four components of Apple's privacy 'nutrition label' manifest.	Q7 (background context on ATT's strategic framing only)	REJECT	Thin, largely reactive commentary on a third-party video rather than original analysis; pre-dates SKAN 4/AdAttributionKit and current ATT status — superseded by CP-03/04's current standards treatment.
CP-27	Measurement, on the Browser's Terms	essay	2026-07-13	2026-08-30	Reads the W3C Attribution API (Private Advertising Technology Group; editors from Apple, Google, Meta, Mozilla) as a governance question — quoting Mozilla's Martin Thomson that the test of any privacy technology is 'who decides', not which data flows are technically permitted.; Argues the browser (via a W3C-standardized measurement API) is positioning itself as the counterparty that grants measure	Q7,Q12; governance-of-measurement chapter, browser-side analog to the app-side SKAN/AAK governance question	QUALIFY	Browser (W3C Attribution API) governance debate an OS-level mobile-app attribution mechanism (SKAN/AAK/GAID); relevant only as a structural analogy for who controls app-measurement standards
CP-28	Glossary — Mobile App Growth term block	glossary	2026-06-03	2026-09-18	Apps DSP: defined as a reframe of the mobile DSP — an app-growth decision layer spanning install, engagement, retention, LTV, commerce and CTV extension, describing the buyer problem rather than the channel.; SKAdNetwork (SKAN): current version is SKAN 4 — explicitly notes 'no SKAN 5 exists' as of authoring.; AdAttributionKit: introduced WWDC 2024, adds re-engagement attribution and alternative-ma	Q3,Q4,Q9,Q11; definitional backbone for the paper's own terminology, directly overlapping the brief's shared-definitions section	RETAIN	Well-hedged definition: ('validate current documentation' on the time-sensitive claim); uses the paper's working vocabulary but confirms 'no forced SKAN cutoff claim is still true at the 27 2026 cutoff.

CP-29	Glossary — cross-cutting incrementality/attribution terms (IROAS, Incrementality, Conversion Lift, Attribution window, Prebid Mobile)	glossary	2026-06-03	2026-09-18	IROAS defined with the same Ovative/Albertsons 6.5x-variance finding as CP-07/CP-15, framed as reason to ask which methodology family produced any given IROAS number.; Incrementality (defined twice, decision-intelligence and gaming contexts) is explicitly distinguished from attribution: 'the causal contribution of an activity...distinct from attribution, which assigns credit to conversions that ma	Q3; cross-cutting definitional support for the whole causal-vs-attribution argument of the paper	RETAIN	Consistent, well-scope definitions; note the Ovative/Albertsons iRO figure recurs at three si locations (CP-07, CP-15 29) — cite the underlying study once in the paper each site restatement.
CP-30	AI Can Interpret Data. It Can't Vouch For It.	byline	2026-08-06	2026-08-06	Argues AI can interpret data and bridge systems but cannot itself establish trust; as agents move from assisting to deciding, standards must cover provenance, permission, delegation and auditability as a real-time execution layer, not an after-the-fact audit.	Q12; general agentic-trust argument applicable to app-growth DSP AI features	QUALIFY	General agentic-advert column, not mobile-app-specific; background c for the AI/agentic chap only — fetch and read t full AdExchanger piece before quoting.
CP-31	Why Agentic Measurement Will Reprice The Ad Market	byline	2026-05	2026-05	Argues agent-mediated measurement collapses the gap between exposure and outcome, with pricing implications for the ad market broadly.	Q3,Q12; agentic-measurement-repricing argument relevant to the paper's Q12 chapter	QUALIFY	General cross-format argument, not mobile-a-specific; fetch and read full piece before citing specific claim.
CP-32	The Future of Marketing Measurement: From Reports to Real-Time Feedback	byline	2026-06	2026-06	Argues that as agents make decisions in real time, measurement must shift from delayed post-campaign reports to live, machine-readable feedback functioning as a pricing signal.	Q3,Q7,Q12; real-time-feedback argument relevant to the paper's measurement-operating-model chapter	QUALIFY	General DMEXCO color (author is an official DMEXCO columnist —, paid/affiliated commen channel, disclosed on t About page), not mobile-app-specific; fetch full before citing.
CP-33	How Has Your Data Strategy Changed With Agentic AI at Your Doorstep?	byline	2026-09-09	2026-09-09	Argues the governance question for agentic marketing is decision rights and action rights, not raw data access — an agent needs decision-ready intelligence (what data means, what it may be used for, what action is permitted), not just accessible data; the audit record is what makes accountability real rather than asserted.	Q12; decision-rights/action-rights framing directly reusable for the paper's AI-governance chapter	QUALIFY	General agentic-data-strategy column, not m app-specific; most rec byline in the corpus (Se 2026) — fetch full text before citing given pro to the research cutoff.
CP-34	About page — author roles and disclosures	data	2026-06-01	2026-09-23	Author currently holds 'Global Head of Enterprise' at Samba TV (May 2025-present), owning growth for Data License/Measurement/Identity — direct commercial stake in CTV-measurement claims.; Author was EVP, GM International Growth at Verve Group (Sep 2022-Aug 2024), with P&L responsibility across Verve DSP/Match2One/Moments.AI — Verve Dataseat is scored as a vendor in the mobile-app-growth vendor un	Disclosures for all chapters touching CTV (Samba TV), mobile DSPs (Verve/Dataseat), and agentic standards (AdCP)	RETAIN	Primary source for the disclosures list; page it notes financial specific (ACV/EBITDA/deal size were deliberately abstr at the author's own req so treat quantitative cla about these employers qualitative only.

H. Scoring anchors

These anchors were fixed in the research brief before any product was scored. A 0 means documented absence on any dimension; D5 and D6 also define a specific 0. "n/e" means not enough public evidence to score and is never counted as zero. "n/a" means the dimension does not apply. Scores describe documented capability, not performance.

DIMENSION	1	2	3
D1 Optimization objectives	Install or cost-per-install only	Post-install event or cost-per-action optimization documented	Value, return or retention-based optimization documented as live
D2 Re-engagement and suppression	Retargeting claimed without mechanism	Re-engagement campaigns using advertiser or attribution-company audiences documented	Plus suppression controls and a documented holdout option for re-engagement
D3 iOS privacy measurement	Generic "SKAN-ready" claim	SKAdNetwork 4 support documented (postbacks, conversion values or optimization)	SKAdNetwork 4 and AdAttributionKit support documented, with postback-level or aggregated reporting for the advertiser
D4 Creative workflow	Standard assets only	Creative services or testing tools documented	Creative testing with a stated design, plus playable or interactive formats
D5 App and placement visibility	"Transparency" claimed without detail (0: blind, no publisher reporting)	App or bundle-level reporting documented	App and placement or format-level reporting, plus block and allow lists
D6 Fee and cost transparency	Pricing model stated only (0: undisclosed bundled margin, no breakdown)	Media cost separated from platform fee, or a published fee model	Bid or impression-level cost data, or a published fee schedule with media-cost reporting
D7 Experiment and incrementality	Incrementality claimed without method	Holdout or lift method available on request or managed	Self-serve or standard experiment product with a stated method and results reporting
D8 Reporting export	Dashboard only	Reporting API or scheduled aggregated export	Log-level or impression-level export to the buyer
D9 Buyer control	Managed only; buyer cannot change campaigns	Self-serve interface for campaigns, budgets and bids	Campaign-management API that creates and edits campaigns, budgets and bids
D10 CTV-to-app	Announced or claimed	CTV buying with app-outcome reporting through an attribution company or household matching	CTV-to-app product with a stated identity method and a causal measurement option
D11 Non-gaming evidence	Non-gaming focus claimed	Two or more named non-gaming case studies (vendor-reported)	Plus third-party or filing-level evidence of material non-gaming business

I. Data and reproduction

The package contains the canonical evidence files, every table above as CSV, and the figure data. It also holds the scripts that build the figures and the page, and a README that lists inputs, assumptions and how to rebuild. The scripts use the Python standard library only. Observed data, modeled estimates and synthetic teaching values are kept in separate files and labelled in every figure.

Data downloads

Buttons generate CSV files from the data embedded in this page, so they work offline. Links point to the files shipped in the package.

Vendor profiles (Figure 14 detail)

Profiles for the eleven execution products scored in chapter 9, printed apart from the matrix so the matrix can be read as one comparison. Scores and sources match the matrix.

AppLovin · Axon Ads Manager AppLovin Corporation · roles: DSP, ad network · profile group V1		
OWNERSHIP Public (ticker APP). One reportable segment since the sale of its game studios to Triplicated closed on 30 June 2025, for \$430.6m in cash after adjustments plus about 20% of Triplicated's fully diluted equity, valued at \$285.0m (\$715.6m in total) [AB-S01; AB-S03; FX1-S01]. AppLovin also owns its ad platform (AppLovin Ads, formerly Axon Ads Manager), the MAX mediation auction, the Adjust attribution company (bought 2021) and Wurl (bought 2022) [AB-S01; G2-S01].	FINANCIAL SCALE From filings: FY2025 revenue \$5,481m, up 70%, booked net of publisher payouts because AppLovin acts as agent; adjusted EBITDA \$4,512m, 82% of revenue [AB-S01; AB-S02]. These are continuing operations, after the Triplicated sale. Q2 2026: revenue \$1,924m, up 53%; adjusted EBITDA margin 83.9% [G2-S01; G2-S02]. The filings do not break out e-commerce revenue [G2-S01].	WHAT IS SOLD A self-serve, ML-driven (Axon/AXON engine) auction connecting app and web/e-commerce advertiser demand to AppLovin's owned and MAX-mediated publisher supply, billed dynamically (dCPI, cDPM, CPM, or performance-goal billing) rather than fixed per-impression/per-action pricing.
BUYER AND APP FIT App advertisers and, since 2024, web e-commerce advertisers buying scaled, model-driven acquisition [G2-S01; OA4-S06]. AppLovin recommends a daily budget that buys at least 15 to 20 conversions a day; ad-revenue and blended ROAS goals also require the app to use MAX [V1-S07]. Reports show the source app and placement type, but no advertiser block or allow list is documented [V1-S05; RS-S06].	ECONOMICS pricing_model: Dynamic; dCPI (install-optimized), cDPM (ROAS/CPA-optimized), CPM (paced/creative-testing spend), CPI (proven-winner fast spend); web/e-commerce campaigns additionally billed on CPA/ROAS outcomes.; minimums: not disclosed	DATA REQUIREMENTS MMP or server-to-server postbacks for event and ROAS optimization. On iOS, the only advertiser-side SKAdNetwork step documented is an attribution-partner setting to share the SKAN transaction ID. AppLovin's SKAdNetwork page covers DSPs bidding into MAX, not advertiser reporting, and no SKAN 4 or AdAttributionKit handling is documented [RS-S02; V1-S02].
SUPPLY DEPENDENCIES Owned/aggregated publisher SDK supply via MAX mediation plus AppLovin Exchange; CTV supply via owned Wurl (acquired 2022).	MEASURABLE CONTROLS Reporting API: aggregated JSON/CSV with cost, ROAS, retention and billing-method fields, source app and placement type, and a 45-day request window; no log-level export [V1-S05; RS-S05; RS-S06]. Campaign Management API for app and web campaigns: create and update campaigns, goals, budgets (global or by country), targeting and creative sets [V1-S06; RS-S01]. No advertiser block or allow list is documented.	IMPLEMENTATION BURDEN Low to moderate. Self-serve; AppLovin opened the platform to all advertisers, without a referral code, on 22 June 2026 under the name AppLovin Ads [OA4-S06]. Event and ROAS optimization need MMP or server-to-server events, and ad-revenue and blended ROAS goals need MAX [V1-S07].
PROOF QUALITY - Safe Sleeve / Sweet Flexx: AppLovin training material (Ads Playbook); narrative walkthrough with no comparator [V1-S03] - Kikoff (CTV): AdExchanger report quoting AppLovin and Kikoff; before-and-after results with no holdout described [V1-S09] - Geo-holdout tests: AppLovin's own blog describes running geographic holdouts for clients (20% to 50% of the footprint held out, at least 90% statistical power, campaigns out of learning), for web and e-commerce brands; a method description with no published results [V1-S04]	SWITCHING LIMITS Not documented in the sources we read. AppLovin's disclosure on its Axon model says nothing about advertisers' rights over the use of their campaign data [KL-S12].	RISKS AND CONFLICTS - Ownership overlap: AppLovin sells media, runs the MAX mediation auction and owns the Adjust attribution company. It says Adjust data is not shared with it unless the customer directs; that assurance has not been independently checked. This is a potential conflict of interest, not evidence of misconduct [AB-S01]. - Model concentration: the FY2025 10-K attributes the 70% revenue rise mainly to improved Axon Ads Manager performance (installations up 3%, net revenue per installation up 72%) and cites no independent efficiency or incrementality measure [V1-S01]. - Unproven allegations: 2025 short-seller reports alleged improper data practices and retargeting-heavy e-commerce results, and AppLovin denied them. A securities class action had no ruling on the motion to dismiss by the cutoff, and AppLovin's filings name no specific investigation (chapter 2) [G3-S05; G3-S07; G3-S10].
POOR FIT		

Buyers who need an advertiser block or allow list (source apps are reported, but no block control is documented), a split of media cost from fees, a self-serve lift test (geo holdouts are a managed service described for web brands), re-engagement campaigns, or documented SKAN 4 and AdAttributionKit handling [V1-S05; RS-S06; RS-S07; V1-S04; RS-S02].

DIMENSION	FINAL	AVAILABILITY	GRADE	RATIONALE	SOURCES	SECOND SCORER	RECONCILIATIO
D1 Optimization objectives	3	live	A	Reporting API exposes ROAS/ad_roas and retention (1-28d) metrics plus bidding_and_billing_method; Campaign Management API supports automated bidding to a cost/ROAS target.	V1-S05 V1-S06	3	R1 agree
D2 Re-engagement & suppression	n/e	unconfirmed	C	No official Axon Ads Manager documentation describing a retargeting/re-engagement campaign product or mechanism was found or independently opened in this research.		n/e	R1 agree
D3 iOS privacy measurement (SKAN / AAK)	1	live	A	Only SKAN evidence for advertisers: AppsFlyer setup toggle 'Share SKAN transaction ID'. No SKAN 4, conversion-value, AdAttributionKit or SKAN reporting documentation; the SKAN page is for MAX demand partners.	RS-S02 V1-S02	1	R2 conservative: lower score
D4 Creative workflow	2	live	B	Official guidance documents iterative testing across video/playable/static/DPA formats and a playable-ad spec; no disclosed randomized or allocation-based testing methodology found.	V1-S03 V1-S07	2	R1 agree
D5 App & placement visibility	2	live	A	Advertiser report exposes hashed source-app ID ('site ID'), source application name, placement_type and encrypted placement ID; {APP_ID} macro to MMP. No documented advertiser block/allow-list control found.	V1-S05 RS-S06	2	R3 evidence found by one scorer
D6 Fee & cost transparency	1	live	A	Billing documents prepay daily charges and bidding strategies (target goal with CPI billing; auto-bidding with CPM billing). Cost is a single advertiser-spend figure; no media-cost vs fee separation or bid-level cost data.	RS-S07 RS-S01	1	R2 conservative: lower score
D7 Experiment / incrementality support	2	managed	B	AppLovin's own blog describes running geo-holdout incrementality experiments for clients with a stated method (>=90% power, 20-50% holdout) -- available on request/managed, not a self-serve standardized product.	V1-S04	n/e	R3 evidence found by one scorer
D8 Reporting export	2	live	A	Reporting API returns aggregated JSON/CSV data across campaign/creative/platform/country/device/placement dimensions; no impression-level export documented.	V1-S05	2	R1 agree
D9 Buyer control (API & self-serve)	3	live	A	Axon Campaign Management API documents POST endpoints to create/update campaigns and creative sets and to set global or country-level budgets.	V1-S06	3	R1 agree
D10 CTV-to-app	2	live	B	Wurl/AppDiscovery CTV inventory is bought on a CPI basis with MMP attribution to installs; no stated causal/incrementality measurement method specific to CTV was documented.	V1-S08 V1-S09	2	R1 agree
D11 Non-gaming vertical evidence	2	live	B	FY2025 10-K states e-commerce/web-advertiser expansion as a stated objective; named non-gaming case studies include Safe Sleeve, Sweet Flexx and Kikoff (fintech); no filing-level quantification of non-gaming revenue share was found.	V1-S01 V1-S03 V1-S09	2	R1 agree

Moloco Ads

Moloco, Inc. · roles: DSP, CTV DSP · profile group V1

OWNERSHIP status: private; note: Most recent disclosed valuation is a >US\$2B 2023 secondary share sale (per press); reported (Bloomberg, Jan 2026, via Investing.com syndication) to be in early-stage IPO-adviser discussions -- unconfirmed, still private at cutoff. Backers reported: Fidelity Management & Research, Tiger Global Management.	FINANCIAL SCALE Not disclosed (private company; no filings)	WHAT IS SOLD An ML/'compound AI' performance-advertising DSP for app install and post-install value optimization, extended to Performance CTV (household/IP-matched, MMP-attributed) and offered separately for retail/commerce media (Moloco Commerce Media, out of V1 scope).
BUYER AND APP FIT App marketers (gaming and non-gaming) needing ROAS/LTV-oriented bidding with SKAN support and	ECONOMICS oCPM is the only documented pricing model. Moloco asserts a 'non-variable DSP margin' without	DATA REQUIREMENTS MMP postback integration (AppsFlyer, Adjust, Singular, Branch, Kochava) for event/ROAS

a documented campaign-management + reporting API; also fits marketers wanting to extend the same measurement stack to CTV.

publishing the rate; log fields for media cost and fee percentage exist but are marked deprecated, so no current fee split is documented [RS-S18; RS-S19]. Minimum spend: not stated in the sources we read.

optimization and re-engagement; SKAdNetwork conversion-value configuration read/written via the DSP API.

SUPPLY DEPENDENCIES

Programmatic exchange buying; reports break results out by app or site, sub-publisher and exchange [V1-S15; RS-S09]. The sources we read document no owned publisher supply. Performance CTV bids at household (IP) level for CTV app-install campaigns [V1-S13; RS-S24].

MEASURABLE CONTROLS

Campaign/AdGroup CRUD API; Report API (aggregated); Log API (impression/click/conversion event-level, disabled by default, available on request); CreativeGroups (A/B-test structure) API.

IMPLEMENTATION BURDEN

Moderate: requires MMP postback configuration and, for full log-level visibility, a specific request to enable the Log API (not on by default).

PROOF QUALITY

- Nexon (Performance CTV): Moloco case study describing a ghost-bid test with an even test/control split, designed and analysed with Moloco's data scientists; vendor-run, with no independent check [V1-S11; RS-S20]
- Wayfair, Bevmo, Costco, Benjamin, Musinsa, Experian, FreeNow, ManoMano: Moloco case studies; narrative walkthroughs with no comparator on the index page [V1-S11]
- Performance CTV launch: Moloco reports up to 1.5 times higher ROI on CTV than on mobile and about two-thirds of installs within six hours of exposure; no test design described [IJ-S08]

SWITCHING LIMITS

Not documented in the sources we read. Event-level logs can be exported on request, with 90-day access [V1-S16; RS-S21].

RISKS AND CONFLICTS

- Incrementality evidence is self-reported. The Performance CTV page describes built-in ghost-bid tests designed and analysed with Moloco's data scientists; the re-engagement page says only that Moloco offers incrementality testing, without naming a method. We found no independent audit of results [V1-S13; V1-S14; RS-S20].
- Private company: financial scale, headcount and customer concentration are not disclosed [V1-S20].

POOR FIT

Buyers needing a fully self-serve, always-on log-level data feed should note the Log API is disabled by default and requires a specific request.

DIMENSION	FINAL	AVAILABILITY	GRADE	RATIONALE	SOURCES	SECOND SCORER	RECONCILIATION
D1 Optimization objectives	3	live	A	Developer API documents reading/writing SKAN conversion-value configuration and campaign goals spanning install/event/ROAS; case studies report ROAS-based outcomes across verticals.	V1-S17 V1-S11	3	R1 agree
D2 Re-engagement & suppression	3	live	A	APP_REENGAGEMENT campaign type; CustomerSet IDFA/ADID lists and app/event audiences can be included or excluded; RE page states incrementality testing is offered for re-engagement (method not detailed there).	RS-S09 RS-S11 RS-S12 V1-S14	3	R2 higher score has grade-A evidence
D3 iOS privacy measurement (SKAN / AAK)	2	live	A	Developer API documents reading and updating SKAdNetwork conversion-value configuration and SKAN analytics; no AdAttributionKit mention was found.	V1-S17	2	R1 agree
D4 Creative workflow	2	live	A	Developer API documents CreativeGroups supporting an A/B-test structure (a disclosed-allocation design); playable/interactive format coverage was referenced in secondary (help-center) sources not independently opened in this research, so format-coverage is not counted toward a score of 3.	V1-S17	3	R2 conservative: lower score
D5 App & placement visibility	3	live	B	Brand Safety Policy documents a global block list plus customer-configurable inclusion/exclusion lists, and states advertisers receive 'transparent impression level reporting by publisher and exchange'.	V1-S12	3	R1 agree
D6 Fee & cost transparency	1	live	A	oCPM is the only pricing model; 'non-variable DSP margin' asserted without published rate. Log fields for media cost and fee_percent exist but are marked Deprecated, so no current fee separation is documented.	RS-S18 RS-S19	1	R3 evidence found by one scorer
D7 Experiment / incrementality support	2	limited	B	Ghost-bidding RCT (even test/control split, bids without serving) designed and analysed with Moloco data science; described for CTV and offered for RE. No self-serve incrementality product or results UI documented.	RS-S20 V1-S13	2	R2 conservative: lower score
D8 Reporting export	3	live	A	Report API provides aggregated dimensional data; Log API exports impression/click/conversion event-level records.	V1-S15 V1-S16	3	R1 agree
D9 Buyer control (API & self-serve)	3	live	A	Developer API documents full campaign CRUD and AdGroup create/update/list operations, not only reporting.	V1-S17	3	R1 agree

DIMENSION	FINAL	AVAILABILITY	GRADE	RATIONALE	SOURCES	SECOND SCORER	RECONCILIATIO
D10 CTV-to-app	3	live	B	Performance CTV documents a stated household/IP-matching identity method plus a built-in ghost-bidding causal measurement option, with attribution running through the advertiser's own MMP.	V1-S13	3	R1 agree
D11 Non-gaming vertical evidence	2	live	B	Case-studies page names >=8 non-gaming advertisers (Wayfair, Bevmo, Costco, Benjamin, Musinsa, Experian, FreeNow, ManoMano); Moloco is private and does not file, so no filing-level revenue quantification exists.	V1-S11	2	R1 agree

Liftoff Accelerate

Liftoff Mobile, Inc. · roles: DSP · profile group V1

OWNERSHIP

Public (Nasdaq: LFTO). Priced its IPO at \$23.00 on 3 June 2026 and began trading on 4 June; the offering closed on 5 June with the over-allotment exercised, raising about \$472.4m net, of which \$409.2m repaid debt [AB-S11; FX1-S05; FX1-S06]. A first attempt, filed in January 2026, was withdrawn on 17 February 2026 [FX1-S08]. The 10-Q refers to a 'primary private-equity sponsor' without naming it; an affiliate of the sponsor underwrote 2.7m shares [V1-S21; G2-S09]. General Atlantic was allocated about 1.3m shares in the offering [V1-S28].

BUYER AND APP FIT

App marketers (gaming and non-gaming, per named case studies) wanting a single DSP spanning acquisition and re-engagement; Direct specifically fits buyers wanting curated premium publisher access over open-exchange buying.

SUPPLY DEPENDENCIES

Direct: curated direct access to 150k+ publisher apps (per vendor claim); Accelerate: broader UA/re-engagement supply; SDK monetization / exchange supply sits in the separate Liftoff Monetize (Vungle Exchange) product, not scored here.

PROOF QUALITY

- Delivery Hero: Liftoff case study citing a 20% install uplift in Latin America from incrementality testing; test design not disclosed [V1-S26; RS-S36]
- Binance: Liftoff case study ('beats D7 ROAS goals by 30% in key markets'); narrative, no comparator [V1-S26]
- PEP: Liftoff's internal, managed experimentation programme; no holdout method documented [RS-S37]

POOR FIT

Buyers who need a generally available self-serve campaign API (Liftoff's is a closed beta), log-level export, an advertiser block list, a documented lift-test method or a CTV-to-app product. None is

FINANCIAL SCALE

From filings: FY2025 revenue \$685.7m, up 32%, booked net as agent; adjusted EBITDA \$374.4m, 55% of revenue; net loss \$23.1m [AB-S10; FX1-S05; FX4-S06]. Q2 2026 revenue \$219.5m, up 35%, and H1 2026 \$425.1m; one customer made up about 10% of revenue [V1-S21; G2-S09]. Revenue is split by geography only, with no product-line split [V1-S21].

ECONOMICS

Priced per advertising unit, such as installs or impressions. Advertiser reporting shows only spend, with no media-cost versus fee split or published advertiser fee [V1-S21; RS-S31]. Minimums: not disclosed.

MEASURABLE CONTROLS

Reporting API and dashboard CSV export, with publisher app ID, publisher name and ad format; aggregated, with no log-level export [RS-S31; RS-S35]. A Campaign Management API, in closed beta for select customers, creates user-acquisition campaigns and sets daily spend, goals, country targeting and creatives; spend changes are limited to three a day, within 50% up or down [RS-S29]. No advertiser block or allow list is documented [RS-S35].

SWITCHING LIMITS

Not documented in the sources we read.

WHAT IS SOLD

A DSP powered by Liftoff's Cortex models for app user acquisition and re-engagement. Accelerate offers broader access; Direct offers curated publisher access with account-team optimization. Campaigns can be optimized toward CPC, CPI, CPA, ROAS or predicted-LTV goals [V1-S22; V1-S23; RS-S29].

DATA REQUIREMENTS

The Reporting API shows SKAdNetwork installs with and without conversion values, and Accelerate names SKAdNetwork as a supported use case; 'SKAdNetwork Cortex models' were announced for late 2024. No SKAN 4 postback handling or AdAttributionKit support is documented [RS-S31; V1-S22; V1-S27]. An audience-ingestion API lets partners send audiences for targeting [RS-S30].

IMPLEMENTATION BURDEN

Mostly managed. Direct relies on Liftoff's account, creative, product and engineering teams [V1-S23]; the campaign API is a closed beta for select customers [RS-S29].

RISKS AND CONFLICTS

- Liftoff reports revenue net, as an agent. That is an accounting presentation, not an advertiser-facing cost-transparency feature [V1-S21].
- Customer concentration: one customer made up about 10% of Q2 2026 revenue and 17% of receivables at 30 June 2026 [G2-S09].
- Control after the listing: the 10-Q refers to an unnamed primary private-equity sponsor; its remaining stake is not stated in the sources we read [V1-S21].

documented [RS-S29; RS-S31; RS-S35; RS-S37; RS-S38].

DIMENSION	FINAL	AVAILABILITY	GRADE	RATIONALE	SOURCES	SECOND SCORER	RECONCILIATIO
D1 Optimization objectives	3	live	B	Official Accelerate page states the Cortex-powered product suite lets advertisers 'directly optimize for any type of KPI (CPC, CPI, CPA, ROAS, pLTV)'.	V1-S22	3	R1 agree
D2 Re-engagement & suppression	2	live	B	Official re-engagement page documents a data-driven approach intended to avoid cannibalizing organic conversions, with claimed lift/cost figures; no suppression/exclusion mechanism or holdout/incrementality methodology is documented on the page itself.	V1-S25	2	R1 agree
D3 iOS privacy measurement (SKAN / AAK)	1	live	B	SKAdNetwork is named as a supported use case and 'SKAdNetwork Cortex models' were announced for late 2024, but no conversion-value/postback mechanics or AdAttributionKit support were documented in pages opened in this research.	V1-S22 V1-S27	1	R1 agree
D4 Creative workflow	2	live	A	Creative Lab upload/QA for playable and interactive HTML, LXA and triple-page formats; creative services and 'testing technologies' marketed. No randomized or disclosed-allocation creative test design documented.	RS-S33 RS-S34 V1-S22	2	R3 evidence found by one scorer
D5 App & placement visibility	2	live	A	Reporting API returns publisher_app_store_id, publisher_name and ad_format; dashboard breaks down by source app. No advertiser block/allow-list control found in docs.	RS-S31 RS-S35	2	R3 evidence found by one scorer
D6 Fee & cost transparency	1	live	A	10-Q; price per advertising unit (installs, impressions), revenue reported net as agent. Reporting shows only 'spend'; no media-cost versus fee breakdown or published fee for advertisers.	V1-S21 RS-S31	1	R3 evidence found by one scorer
D7 Experiment / incrementality support	1	limited	B	Delivery Hero case reports Liftoff-run incrementality tests with uplift figures but no stated design; PEPr is an internal/managed experimentation programme without a documented holdout method.	RS-S36 RS-S37	1	R2 conservative: lower score
D8 Reporting export	2	live	A	Reporting API (aggregated, cohort windows, CSV/JSON) and dashboard CSV export. No impression- or log-level export documented.	RS-S31 RS-S35	2	R3 evidence found by one scorer
D9 Buyer control (API & self-serve)	3	beta	A	Campaign Management API (closed beta, select customers): create UA campaigns, set/schedule daily spend, goals, country targeting, assets and creatives. Only UA expansion campaigns; spend changes limited to 3x/day within +/-50%.	RS-S29	3	R3 evidence found by one scorer
D10 CTV-to-app	n/e	unconfirmed	C	No CTV product is mentioned on Liftoff Monetize's own page; absence from a marketing page is not the same as documented absence of any CTV product, so this is scored n/e rather than 0.	V1-S24	n/e	R1 agree
D11 Non-gaming vertical evidence	3	live	A	Named non-gaming Accelerate cases: Delivery Hero/PedidosYa, Acorns, Casumo, ReelShort. 10-Q cites 'scaled leadership in non-gaming verticals' and diversified end markets; qualitative only, no vertical revenue split.	V1-S21 RS-S36 RS-S39 V1-S22	3	R2 higher score has grade-A evidence

Unity Ads (Grow)

Unity Software Inc. · roles: ad network, mediation, DSP-like UA · profile group V1

OWNERSHIP

Public (ticker U). Unity reports one segment; Grow (advertising) and Create (engine software) are revenue categories [G2-S08]. Its Vector machine-learning platform rolled out in Q1 2025 [AB-S04]. Unity merged with ironSource in November 2022 [AB-S06].

BUYER AND APP FIT

Primarily game developers within the Unity/ironSource ecosystem; non-gaming vertical claims found in Unity's FY2025 10-K describe the separate Create Solutions/engine business

FINANCIAL SCALE

From filings: FY2025 revenue \$1,849.6m; adjusted EBITDA \$408.8m (22%), a whole-company figure that includes the engine business; GAAP net loss \$401.5m [AB-S05]. Grow revenue was \$1,228.2m in FY2025 and Create \$621.4m [AB-S04]. Q2 2026 Grow revenue was \$388.9m, up 35%, credited to the Vector-driven Unity Ad Network and partly offset by the older ironSource network [G2-S06; G2-S07]. Grow revenue is mostly booked net, and gross where Unity is the publisher [G2-S08].

ECONOMICS

CPI billing for install, event and ROAS goals; CPM for creative testing, for iOS and as a fallback when there are no installs. No media-cost versus fee split or

WHAT IS SOLD

Unity Ads UA (install/retention/ROAS bidding) plus LevelPlay mediation (auctioning Unity Ads and third-party networks/bidders for publishers), positioned as an integrated developer-ecosystem growth stack.

DATA REQUIREMENTS

Event and ROAS goals need post-install data from an MMP or server to server [RS-S40; RS-S41]. SKAdNetwork postbacks are forwarded to MMPs with SKAN 4.0 fields (source identifier, coarse value,

(automotive, retail, healthcare etc.), not the Grow/advertising product -- do not inherit one onto the other.

published fee [RS-S46; RS-S44]. Minimums: not disclosed.

postback sequence), and SKAN installs appear in the dashboard and Statistics API. No AdAttributionKit support is documented [V1-S35; RS-S43; RS-S44].

SUPPLY DEPENDENCIES

Unity Ads network plus Unity Exchange (via a single SDK, per the Unity Ads product page); LevelPlay mediates Unity Ads and third-party bidding networks for publishers.

MEASURABLE CONTROLS

Reporting dashboard, 'History and Performance' tool, CSV export, and a documented Advertising Management API plus an Advertising Statistics API (services.docs.unity.com).

IMPLEMENTATION BURDEN

Self-serve: a UA dashboard, bulk-management spreadsheets and the Advertising Management API [RS-S51; RS-S52]. Event and ROAS goals need MMP or server-to-server post-install data [RS-S40].

PROOF QUALITY

- Lihuhu, Qcplay: Unity case references on the Unity Ads product page, both gaming studios; no method disclosed [V1-S32]

SWITCHING LIMITS

Not documented in the sources we read.

RISKS AND CONFLICTS

- The older ironSource Ad Network is in decline (11% of Grow revenue in Q4 2025) and has been sunset, and Unity sold Supersonic on 4 August 2026. Headline Grow revenue still includes both; Unity reports 'strategic' Grow revenue of \$329.0m for Q2 2026, up 63% [AB-S05; G2-S07].
- Unity's non-gaming language in its 10-K describes the engine business or states an aim for ads; we found no named non-gaming case study for Unity Ads [V1-S29; V1-S32].

POOR FIT

Buyers who need re-engagement campaigns, a lift test, a CTV product or named non-gaming results: none is documented for Unity Ads [RS-S42; V1-S31; V1-S29]. Source apps are reported and can be blocked or bid on, but only through abstracted IDs; Unity says it cannot give out source IDs for specific apps [RS-S47; RS-S48].

DIMENSION	FINAL	AVAILABILITY	GRADE	RATIONALE	SOURCES	SECOND SCORER	RECONCILIATION
D1 Optimization objectives	3	live	A	Official documentation defines and lets advertisers configure Install (CPI) bids, Retention bids, and ROAS targets.	V1-S33	3	R1 agree
D2 Re-engagement & suppression	n/e	unconfirmed	C	No documentation of a retargeting/re-engagement campaign product for Unity Ads advertisers was found or independently opened in this research.		n/e	R1 agree
D3 iOS privacy measurement (SKAN / AAK)	1	live	A	Privacy-compliance documentation confirms SKAdNetwork coverage (conversion values, ad-slot limits, dashboard support, postback integration) at an overview level, but no specific SKAN version (4 vs earlier) or AdAttributionKit mention was found in the pages retrieved.	V1-S35	2	R2 conservative: lower score
D4 Creative workflow	3	live	A	Creative Testing campaign goal distributes impressions equally across creative packs (disclosed allocation, CPM billed) with per-pack reporting; video, playable and end-card formats documented.	RS-S45 RS-S46	3	R2 higher score has grade-A evidence
D5 App & placement visibility	2	live	A	Reporting by Source App ID and source app category, with allow/blocklists and source bidding by ID; but IDs are abstracted and Unity states it 'cannot give out source IDs for specific apps'. No placement dimension.	RS-S47 RS-S48	2	R3 evidence found by one scorer
D6 Fee & cost transparency	1	live	A	Billing model documented (CPI for install/event/ROAS goals; CPM for creative testing, iOS and zero-install fallback). No media-cost versus fee separation or published fee.	RS-S46 RS-S44	1	R3 evidence found by one scorer
D7 Experiment / incrementality support	n/e	unconfirmed	C	No incrementality/holdout documentation for Unity Ads UA was found in this research.		n/e	R1 agree
D8 Reporting export	2	live	A	Reporting/analytics documentation lists a dashboard, History and Performance tool, CSV export, and REST APIs; no confirmation of log-level (vs aggregated) granularity was found.	V1-S34	2	R1 agree
D9 Buyer control (API & self-serve)	3	live	A	Documentation links to a distinct Advertising Management API (services.docs.unity.com/advertise/v1) alongside the Statistics API, indicating a documented campaign-management API rather than reporting-only access.	V1-S34	3	R1 agree
D10 CTV-to-app	n/e	unconfirmed	C	No CTV product documented for Unity Ads/Grow was found in this research.		n/e	R1 agree

DIMENSION	FINAL	AVAILABILITY	GRADE	RATIONALE	SOURCES	SECOND SCORER	RECONCILIATIO
D11 Non-gaming vertical evidence	1	live	A	10-K frames capturing ad spend 'outside of our core gaming audience' as an aspiration and risk; UA docs and product page are game-focused. No named non-gaming UA case studies found.	V1-S29 V1-S32	1	R3 evidence found by one scorer

Mintegral

Mobvista Limited (HKEX: 1860) · roles: DSP, ad network, ADX, SSP · profile group V1

OWNERSHIP

Mintegral is Mobvista's programmatic ad platform and is not listed separately. Mobvista is listed in Hong Kong (HKEX: 1860) [V1-S37; AB-S08].

FINANCIAL SCALE

From filings: Mobvista FY2025 revenue \$2,046.7m, up 35.7%, of which Mintegral \$1,960.9m. Revenue is booked gross, as principal, with a 21.2% gross margin; adjusted EBITDA was \$190.9m, about 9% of gross revenue [AB-S08; FX1-S11]. Mobvista's own ad-tech net revenue after payments to traffic publishers was about \$519m for 2025 [FX1-S11; FX4-S07]. H1 2026: group revenue \$1,155.5m, up 23.2%; Mintegral \$1,115.2m, up 24.3% [V1-S37].

WHAT IS SOLD

A programmatic in-app advertising platform (AppGrowth for UA, Retargeting for re-engagement, Monetization for publishers) with AI/smart bidding (Target ROAS spanning IAA/IAP/Hybrid, and Target CPE) plus an in-house creative studio (Mindworks) and playable-ad automation (Playturbo).

BUYER AND APP FIT

Predominantly gaming advertisers (75.6% of Mintegral H1 2026 revenue) with a growing non-gaming share (24.4%, +15.5% YoY, incl. e-commerce and utilities per the interim report's own vertical breakdown).

ECONOMICS

Billing types CPI, CPM, CPE and oCPI are documented. Target ROAS and Target CPE smart-bidding products made up more than 90% of Mintegral revenue in H1 2026. No media-cost versus fee split for advertisers [RS-S53; V1-S37]. Minimums: not disclosed.

DATA REQUIREMENTS

Postback-based data sharing for retargeting; IDFA and GAID audience lists can be uploaded by API and included or excluded [V1-S40; RS-S55; RS-S56]. No SKAdNetwork or AdAttributionKit documentation appears in about 60 English advertiser pages or the 2026 interim report [RS-S57; V1-S37].

SUPPLY DEPENDENCIES

Direct publisher/app relationships (10,000+ developers, 120,000+ apps per the interim report) plus programmatic DSP/ADX/SSP coverage of the mid-stream ecosystem.

MEASURABLE CONTROLS

Self-serve dashboard for setting ROAS goals [V1-S37]. Mintegral's Open API documentation describes creating campaigns and offers and updating budget, bid and status; some bid types and API access are gated by the account manager [RS-S63; RS-S53; RS-S64]. The reporting API returns aggregated data by offer, creative, sub-publisher, package and geography, and publishers can be allow- or block-listed by ID [RS-S60; RS-S61; RS-S62].

IMPLEMENTATION BURDEN

Hybrid: self-serve dashboard and Open API, with some features gated by the account manager [V1-S37; RS-S63].

PROOF QUALITY

- Winedrops (e-commerce): Mintegral testimonial (10X increase in trials); no comparator [V1-S41]

SWITCHING LIMITS

Not documented in the sources we read.

RISKS AND CONFLICTS

- Mobvista books revenue gross, as principal, so its headline revenue is not comparable with net-reporting peers such as AppLovin or Liftoff. On Mobvista's own net measure, after payments to traffic publishers, H1 2026 group net revenue was \$308.5m against \$1,146.6m of gross ad-tech revenue [FX1-S11; G2-S14; V1-S37].

POOR FIT

Buyers who need documented SKAdNetwork or AdAttributionKit handling, a lift test, log-level export or a CTV product: none is documented [RS-S57; RS-S60; V1-S39].

DIMENSION	FINAL	AVAILABILITY	GRADE	RATIONALE	SOURCES	SECOND SCORER	RECONCILIATIO
D1 Optimization objectives	3	live	A	H1 2026 interim report (filing) states smart-bidding products (IAA ROAS since May 2023, Target CPE since Jul 2024, Hybrid ROAS since Apr 2025, IAP ROAS since Jul 2025) account for >90% of Mintegral revenue; official AppGrowth page documents Target ROAS and Target CPE as selectable models.	V1-S37 V1-S39	3	R1 agree
D2 Re-engagement & suppression	2	live	B	Official Retargeting page documents postback-based re-engagement of active/dormant/lapsed users; no suppression/exclusion control or holdout/incrementality option is documented.	V1-S40	2	R1 agree

DIMENSION	FINAL	AVAILABILITY	GRADE	RATIONALE	SOURCES	SECOND SCORER	RECONCILIATIO
D3 iOS privacy measurement (SKAN / AAK)	n/e	unconfirmed	C	No SKAdNetwork documentation was found on the official marketing pages accessible in this research; the developer wiki (dev.mintegral.com) exists but its SKAN-specific pages could not be located within this research's fetch budget.		n/e	R1 agree
D4 Creative workflow	2	live	B	Mindworks (in-house creative studio) and Playturbo (playable-ad automation) are documented; no disclosed randomized/allocation testing methodology was found.	V1-S38 V1-S42	2	R1 agree
D5 App & placement visibility	3	live	A	Report API segments by Sub (mtgid publisher ID), Package (sub package name), AdType and AdOutputType; API and guide support whitelist/blacklist by mtgid (API access gated by AM).	RS-S60 RS-S61 RS-S62	3	R3 evidence found by one scorer
D6 Fee & cost transparency	1	live	A	Billing types CPI, CPM, CPE and oCPI documented; interim report describes performance-based ad fees. No media-cost versus platform-fee breakdown for advertisers.	RS-S53 V1-S37	1	R3 evidence found by one scorer
D7 Experiment / incrementality support	n/e	unconfirmed	C	No incrementality/holdout documentation was found in this research.		n/e	R1 agree
D8 Reporting export	2	live	A	Advanced Performance Reporting API returns aggregated TSV by offer, creative, sub-publisher, package, geo (daily/hourly, 7-day span). No log-level export documented.	RS-S60	2	R3 evidence found by one scorer
D9 Buyer control (API & self-serve)	3	live	A	Mintegral's Open API creates campaigns and offers with geo, bid type and goal, bid rate, daily cap and budget, creatives and audiences, and updates budget, bid and status; some bid types need permission.	RS-S63 RS-S53 RS-S64	3	R6 fact-check override: Round-2 review follow-up: rule R2 kept the lower first-pass score although the second scorer's grade-A documentation meets the anchor fo 3.
D10 CTV-to-app	n/e	unconfirmed	C	No CTV product for Mintegral was found in this research.		n/e	R1 agree
D11 Non-gaming vertical evidence	3	live	A	H1 2026 interim report (filing) discloses non-gaming revenue at 24.4% of Mintegral revenue (+15.5% YoY) versus gaming at 75.6% (+27.5% YoY) -- filing-level quantification of material non-gaming revenue -- plus a named non-gaming case study (Winedrops, e-commerce).	V1-S37 V1-S41	3	R1 agree

Digital Turbine · DT Exchange / AGP

Digital Turbine, Inc. · roles: ad network, exchange, mediation, OEM/on-device distribution · profile group V1

OWNERSHIP

status: public; ticker: APPS; note: Fiscal year ends 31 March. Two reportable segments: On Device Solutions (ODS) and App Growth Platform (AGP, comprising Advertising Solutions and Ad Monetization Solutions).

BUYER AND APP FIT

Advertisers who want on-device reach (preloads) or offerwall acquisition billed on CPI, CPE or CPA; bid control by country and source app is self-serve on the Offer Wall only [V1-S46; RS-S77]. Publishers use DT for SDK monetization [AB-S07]. Not a fit for CTV or documented re-engagement campaigns [V1-S43; RS-S68].

SUPPLY DEPENDENCIES

DT Exchange in-app supply plus ODS OEM/carrier preload placements (ignite); offerwall supply via ACP Edge.

FINANCIAL SCALE

From filings: FY2026 (year to 31 March 2026) revenue \$565.3m, up 15%; On Device Solutions \$382.4m and App Growth Platform \$185.7m, before a \$2.9m intersegment elimination. Adjusted EBITDA \$122m, about 22% of revenue [AB-S07; V1-S44]. Revenue mixes bases: the exchange marketplace is booked net, as agent, while brand, performance and on-device media are booked gross [G2-S12; FX1-S09].

ECONOMICS

pricing_model: CPI / CPE / CPA billing models; ROAS = Advertiser IAP Revenue / Advertiser Spend, reported at D3/D7/D30.; minimums: not disclosed

MEASURABLE CONTROLS

Offer Wall campaigns: self-serve Micro Bidding by country and source app, a Blocked Apps Tool, suppression of existing installers, and a GraphQL Advertiser Management API for bids, daily budgets and activation [V1-S46; RS-S73; RS-S77]. Unified campaigns: an aggregated Reporting API with supply source and placement type, but no self-serve or API

WHAT IS SOLD

For advertisers: unified DT campaigns, including preloads and dynamic installs on phones, the DT DSP and the Offer Wall, whose self-serve 'Micro Bidding' sets bids by country and source app [AB-S07; V1-S46; RS-S67]. DT Exchange is the supply-side exchange where outside DSPs buy DT inventory [V1-S46]. On Device Solutions distributes apps through phone makers and carriers [AB-S07].

DATA REQUIREMENTS

MMP guides documented for AppsFlyer, Adjust, Singular, Kochava, Branch, and Tenjin; SKAdNetwork bid-request support documented including version 4.0 (sourceidentifier field); no AdAttributionKit mention found.

IMPLEMENTATION BURDEN

Managed for on-device and unified campaigns; self-serve console and API for Offer Wall campaigns. MMP integration is needed for attribution and cost data [V1-S46; RS-S77].

control documented [RS-S71; RS-S72; RS-S75].
Impression-level price data reaches publishers
through SDK callbacks, not buyers [V1-S46].

PROOF QUALITY

- Magazine Luiza, Kenvue (Stayfree), Co-operative Group: DT case studies; narrative walkthroughs with no comparator [V1-S45]
- Playrix: DT case reporting that 97% of dynamic installs were incremental; no design, control or method described [RS-S74]
- Experimenting with Placements: DT documentation of a control-versus-variant comparison for placement changes; not a lift product for re-engagement [V1-S46]

SWITCHING LIMITS

Not documented in the sources we read.

RISKS AND CONFLICTS

- The FY2026 10-K lists significant indebtedness as a risk factor; covenant terms are not covered in the sources we read [V1-S43].
- App Growth Platform is the smaller segment (\$185.7m against \$382.4m for On Device Solutions in FY2026). Most of the company's revenue comes from phone-maker and carrier distribution, a different role from its advertising products [V1-S44].

POOR FIT

Buyers who need documented re-engagement campaigns, a creative-testing method, log-level buyer reporting or a CTV product: none is documented [RS-S68; V1-S43; RS-S75].

DIMENSION	FINAL	AVAILABILITY	GRADE	RATIONALE	SOURCES	SECOND SCORER	RECONCILIATION
D1 Optimization objectives	3	live	A	Documentation defines ROAS (Advertiser IAP Revenue / Advertiser Spend, D3/D7/D30) and documents CPI/CPE/CPA billing models for advertiser campaigns.	V1-S46	3	R1 agree
D2 Re-engagement & suppression	1	live	B	Magazine Luiza case mentions 'push notifications for retargeting' after on-device installs; no documented re-engagement campaign using advertiser/MMP audiences or suppression controls.	RS-S68	1	R3 evidence found by one scorer
D3 iOS privacy measurement (SKAN / AAK)	2	live	A	Bid-request documentation confirms SKAdNetwork version support 'including 4.0' with the sourceidentifier field for SKAN 4.0+, but conversion-value/fine-vs-coarse/crowd-anonymity mechanics and AdAttributionKit are not documented in the pages found in this research.	V1-S46	n/e	R3 evidence found by one scorer
D4 Creative workflow	n/e	unconfirmed	C	No creative-testing documentation with a disclosed method was found in this research; case studies reference 'creative innovation' without method.		n/e	R1 agree
D5 App & placement visibility	3	live	A	Documentation describes a Blocked Apps Tool (block apps at advertiser/product/campaign level) and DT Offer Suppression (excludes existing installers), combined with a Buyer's Report exposing average bid price.	V1-S46	3	R1 agree
D6 Fee & cost transparency	1	live	A	Reporting shows 'gross advertiser spend' only; 10-K describes CPM/CPC/CPI-type pricing and revenue-share supply costs. No advertiser-facing fee or media-cost breakdown.	RS-S71 V1-S43	1	R2 conservative: lower score
D7 Experiment / incrementality support	1	live	A	Documentation describes 'Experimenting with Placements' using a control-group benchmark, but explicitly states no dedicated holdout/incrementality documentation exists for re-engagement/retargeting specifically.	V1-S46	1	R1 agree
D8 Reporting export	2	live	A	Reporting API returns async, aggregated report files; impression-level pricing data is exposed via SDK callbacks (ImpressionData) to publishers, not documented as a buyer-side log-level export product.	V1-S46	2	R1 agree
D9 Buyer control (API & self-serve)	2	limited	A	Self-serve Micro Bidding adjusts bids by country and source app; an Advertiser Management API covers Offer Wall campaigns only (bids, daily budgets, activation). The main DSP and on-device products show no campaign API, so the anchor for 3 is not met for the product as a whole.	V1-S46 RS-S77 RS-S67	3	R6 fact-check override: Round-2 review follow-up: rationale now reflects both scorers' evidence.
D10 CTV-to-app	n/e	unconfirmed	C	No CTV product was found in Digital Turbine documentation reviewed in this research.		n/e	R1 agree
D11 Non-gaming vertical evidence	2	live	B	Case studies name non-gaming advertisers/publishers including Magazine Luiza, Kenvue (Stayfree), Co-operative Group, Moovit, 365Scores, OneFootball; FY2026 filing discloses AGP segment revenue (\$185.7M) but does not break out gaming vs non-gaming within it.	V1-S45	2	R1 agree

Kayzen

Ioniq Group · roles: DSP · profile group V2

OWNERSHIP

Private. Kayzen operates under parent Ioniq Group; Shackleton Ventures states its Victoria Fund 'acquired its interest in Ioniq Group, the parent company of Kayzen, in 2024' (V2-S03, an investor/vendor case-study page, not a filing — medium confidence). No IPO or further M&A found.

FINANCIAL SCALE

Not disclosed. Kayzen files no accounts, and third-party aggregator estimates are not used.

WHAT IS SOLD

A mobile-first, self-serve programmatic DSP for user acquisition, retargeting and 'brand performance' campaigns, sold to apps, agencies, media buyers and e-commerce advertisers, with a managed option [V2-S02]. Documented bidding models: CPM, CPC, CPI, CPA on post-install events and CPA for retargeting; no ROAS or LTV bidding is documented [RS-S79; RS-S80].

BUYER AND APP FIT

In-house UA teams and agencies that want direct, self-serve programmatic control with API access [V2-S02; V2-S04]. Named non-gaming cases include Greggs (published by an investor in Kayzen's parent) and Albertsons, Co-op and Maniko Nails [V2-S03; RS-S89].

ECONOMICS

Published SaaS pricing: a fixed platform fee plus a 2.5% billing fee and overage, billed separately from media cost with no bid markup; the platform reports media costs and CPMs [RS-S87; V2-S02]. Kayzen is the only product in the panel that publishes a fee schedule. Minimums: not stated in the sources we read.

DATA REQUIREMENTS

MMP postback integration, documented for AppsFlyer [V2-S05]. Retargeting audiences are streamed from the MMP or sent through the Audience API, and can be used as allow or suppression lists [RS-S81; RS-S82; V2-S04]. On iOS, 'SKAN campaigns' and the AppsFlyer 'Share SKAN transaction id' setting are referenced; no SKAN 4 or AdAttributionKit handling is documented [V2-S05; RS-S83].

SUPPLY DEPENDENCIES

In-app and mobile-web inventory across 'hundreds of thousands' of publisher apps and sites [V2-S02]. CTV campaigns are also documented, with installs attributed through AppsFlyer's probabilistic cross-platform view-through matching (guide dated September 2026) [RS-S88].

MEASURABLE CONTROLS

Campaign-management API confirmed (edit target bid, daily/total budget, bid multipliers by targeting field, pause/resume status; V2-S04) plus a separate read-only Reporting API.

IMPLEMENTATION BURDEN

Self-serve API/dashboard model per the product page, implying buyer-side setup and ongoing campaign management; managed option also offered per the product page ('leverages its intuitive platform for managed clients and offers it to self-service users') but not detailed further.

PROOF QUALITY

Named cases (Greggs, published by an investor in Kayzen's parent, plus Albertsons, Co-op and Maniko Nails) are vendor or investor case studies with no comparator or design disclosed; no independent measurement [V2-S03; RS-S89].

SWITCHING LIMITS

Not documented in the sources we read.

RISKS AND CONFLICTS

The ownership chain (Ioniq Group, with Shackleton's Victoria Fund as investor since 2024) rests on one investor page, not a filing or Kayzen's own statement [V2-S03]. No conflicts of interest identified.

POOR FIT

Buyers who need a documented lift or holdout test (Kayzen mentions measuring retargeting incrementality with non-attributed events, but documents no design) or ROAS and LTV bidding (not documented) [V2-S04; RS-S79]. Kayzen's CTV campaigns rely on probabilistic cross-device matching with no causal test [RS-S88].

DIMENSION	FINAL	AVAILABILITY	GRADE	RATIONALE	SOURCES	SECOND SCORER	RECONCILIATIO
D1 Optimization objectives	2	live	B	Product page documents post-install/'conversion event' optimization; ROAS language exists in seed-carried 'kAI Suite' claims but no technical doc confirming live ROAS/LTV bidding was opened in this research.	V2-S02	2	R1 agree
D2 Re-engagement & suppression	2	live	A	Retargeting via MMP-streamed audiences and Audience API; audiences whitelisted/blacklisted as suppression lists. Incrementality for retargeting only mentioned in passing, no holdout method documented.	RS-S81 RS-S82 V2-S04	2	R2 higher score has grade-A evidence
D3 iOS privacy measurement (SKAN / AAK)	1	live	A	References to 'SKAN campaigns' and enabling 'Share SKAN transaction id' in AppsFlyer; no SKAN 4, conversion-value handling or AdAttributionKit documentation found.	V2-S05 RS-S83	1	R3 evidence found by one scorer
D4 Creative workflow	2	live	B	Documents a proprietary 'Video+' interactive/fullscreen-video creative format and 'all major creative formats'; no disclosed testing design (split/randomized) or playable-specific format confirmed.	V2-S02	2	R1 agree

DIMENSION	FINAL	AVAILABILITY	GRADE	RATIONALE	SOURCES	SECOND SCORER	RECONCILIATIO
D5 App & placement visibility	3	live	A	Reports support placement-level performance analysis; app-bundle whitelists/blacklists via UI and App Lists API; per-app bid multipliers.	RS-S86 V2-S04 V2-S02	3	R2 higher score has grade-A evidence
D6 Fee & cost transparency	3	live	A	Published SaaS model: fixed platform fee plus 2.5% billing fee and overage, billed separately from media cost with no bid markup; platform reports media costs and CPMs.	RS-S87 V2-S02	3	R3 evidence found by one scorer
D7 Experiment / incrementality support	1	unconfirmed	A	Events API article says non-attributed events can be used 'to measure incrementality in your retargeting campaigns'; no holdout design or experiment product documented.	V2-S04	1	R3 evidence found by one scorer
D8 Reporting export	2	live	A	Help-centre API documentation explicitly describes a read-only Reporting API for programmatic access to campaign performance data (aggregated); no log-level/impression-level export confirmed.	V2-S04	3	R2 conservative: lower score
D9 Buyer control (API & self-serve)	3	live	A	Help-centre API documentation directly describes campaign-management endpoints: edit target bid, daily/total budget, bid multipliers, and pause/resume campaign or creative status.	V2-S04	3	R1 agree
D10 CTV-to-app	2	live	A	CTV campaigns attributed through AppsFlyer cross-platform view-through (probabilistic CTV-to-mobile matching); guide dated September 2026. No causal CTV measurement option found.	RS-S88	2	R5 value over n/a
D11 Non-gaming vertical evidence	2	live	B	Named non-gaming case studies: Greggs (with Dentsu), Albertsons, Co-op and Maniko Nails, plus betting apps (vendor case-study page and an investor page). No filing-level vertical revenue evidence, so the anchor for 3 is not met.	RS-S89 V2-S03	2	R6 fact-check override: Correction of 28 Sep 2026 (COR-117): the second scorer's documented named non-gaming cases meet the anchor; R2 had kept the lower first-pass score.

Smadex

Entravision Communications Corporation (NYSE: EVC) · roles: DSP, managed service · profile group V2

OWNERSHIP

Owned by Entravision Communications (NYSE: EVC) and reported in its Advertising Technology & Services segment [V2-S08]. The acquisition date is not confirmed in the sources we read.

FINANCIAL SCALE

ATS segment (Smadex + Adwake) was approximately 61% of Entravision's FY2025 total net revenue of \$447.6M — implying roughly \$273M, though the 10-K does not break out Smadex alone from Adwake within ATS (derived figure, not a disclosed line item).

WHAT IS SOLD

A managed-service DSP. Entravision says Smadex teams 'act as an extension of our client's marketing team', configuring bidding parameters and fraud-prevention protocols, for mobile user acquisition, retargeting and Connected TV performance campaigns; it processes about 500 billion bid requests a day [V2-S08].

BUYER AND APP FIT

Filing states customers are 'primarily developers of mobile games, fintech apps, and entertainment services' seeking global user acquisition — best fit for buyers wanting a hands-off, managed-service DSP rather than self-serve control; not a fit for buyers who require direct campaign-level API access.

ECONOMICS

Clients pay for outcomes such as installs or in-app purchases, and Smadex teams share media clearing prices and budget allocation; one customer story cites a 'transparent dCPM' model. No published fee schedule or minimum spend [V2-S08; RS-S95].

DATA REQUIREMENTS

Models use advertisers' post-install events, and retargeting uses advertisers' first-party data with match-rate analysis [V2-S08; RS-S91]. No SKAdNetwork, AdAttributionKit or MMP setup documentation was found, and there is no public help center [V2-S08; RS-S92].

SUPPLY DEPENDENCIES

Filing describes Smadex as buying 'advertising inventory in mobile apps, mobile websites and internet-connected televisions (Connected TVs)' in real time via proprietary AI — an exchange-buying model rather than owned/direct SDK supply, per the filing's own framing ('a gateway to the rest of the internet').

MEASURABLE CONTROLS

Smadex teams run the campaigns. Clients receive data on where ads ran and on media clearing prices, and the brand-safety policy provides client allow lists and publisher block lists [V2-S08; RS-S94]. No reporting API, scheduled export or log specification is documented [V2-S08].

IMPLEMENTATION BURDEN

Low buyer-side operational burden by design — Smadex is explicitly positioned as doing the console-level work for the client (managed-service model) rather than requiring in-house trafficking.

PROOF QUALITY

Named customer stories (for example Foodpanda, Talabat, Cabify, Babbal and Exness) are published by Smadex, with no design or comparator disclosed; no third-party study was found [RS-S97; RS-S98].

SWITCHING LIMITS

Not documented in the sources we read.

RISKS AND CONFLICTS

None identified beyond standard managed-service opacity (buyer cedes bidding/fraud-control configuration to Entravision's team, per the filing's own description) — no allegations or litigation found.

POOR FIT

Buyers requiring self-serve campaign control, log-level bid data, or a published fee schedule — none of these were found documented; the model is explicitly managed-only per the filing.

DIMENSION	FINAL	AVAILABILITY	GRADE	RATIONALE	SOURCES	SECOND SCORER	RECONCILIATIO
D1 Optimization objectives	2	live	A	Filing describes proprietary AI calculating 'the precise value of a specific ad placement for a specific advertiser' and performance-based, action-based pricing — documents post-install/event-level optimization; filing text reviewed does not use explicit 'ROAS' or 'LTV' terminology, so not scored 3.	V2-S08	3	R2 conservative: lower score
D2 Re-engagement & suppression	2	live	A	Filing explicitly names retargeting ('re-engaging users who have previously installed an app but have since stopped using it') as a core Smadex capability; no suppression/exclusion or holdout-for-re-engagement detail found.	V2-S08	2	R1 agree
D3 iOS privacy measurement (SKAN / AAK)	n/e	unconfirmed	C	The 10-K text reviewed does not address SKAN or AdAttributionKit; Smadex's own technical pages could not be read via WebFetch in this research (JS-rendered).		n/e	R1 agree
D4 Creative workflow	2	live	B	In-house Creative Studio produces UGC, video and playable ads; 'Creative Selector Algorithm' shifts spend to winners; new concepts piloted 'in a controlled environment' without disclosed design.	RS-S93 V2-S08	2	R3 evidence found by one scorer
D5 App & placement visibility	2	live	A	10-K: clients receive granular data on exactly where ads ran; brand-safety policy documents client whitelists and publisher blacklists. Placement/format-level reporting not explicitly documented.	V2-S08 RS-S94	2	R3 evidence found by one scorer
D6 Fee & cost transparency	2	live	A	10-K: clients pay for outcomes (install, IAP) and teams share the clearing price of media and budget allocation; Cabify story cites a 'transparent dCPM' model. No published fee schedule.	V2-S08 RS-S95	2	R3 evidence found by one scorer
D7 Experiment / incrementality support	n/e	unconfirmed	C	Not addressed in the filing text reviewed.		n/e	R1 agree
D8 Reporting export	n/e	unconfirmed	C	Not addressed in the filing text reviewed; managed-service framing makes a buyer-facing export product plausible but unconfirmed.		n/e	R1 agree
D9 Buyer control (API & self-serve)	1	live	A	Direct filing quote: Smadex teams 'act as an extension of our client's marketing team. Our teams configure bidding parameters, manage fraud prevention protocols' — buyer does not directly change campaigns.	V2-S08	1	R1 agree
D10 CTV-to-app	1	live	A	Filing confirms CTV inventory buying is a real, named capability ('transforms television from a broad branding tool into a precise performance channel') but does not state an identity method (deterministic/probabilistic/household) or MMP-based app-outcome reporting mechanism — scored as announced/claimed rather than 2 or 3.	V2-S08	2	R2 conservative: lower score
D11 Non-gaming vertical evidence	3	live	A	Named non-gaming stories (Foodpanda, Talabat, Cabify, Babbel, Exness), and Entravision's 10-K names Smadex's customers as primarily games, fintech and entertainment app developers: filing-level evidence of material non-gaming customers (qualitative; no revenue split).	V2-S08 RS-S97 RS-S98	3	R6 fact-check override: Correction of 28 Sep 2026 (COR-117): the second scorer's documented named non-gaming cases meet the anchor; rul R2 had kept the lower first-pass score.

Remerge

none disclosed (independent) · roles: DSP, retargeting · profile group V2

OWNERSHIP

Private. Ownership is not disclosed in the sources we read, and no filing, funding round or acquisition was found. AppsFlyer named Remerge a Premier Partner on 13 May 2025, a partner tier, not an ownership fact [V2-S06].

FINANCIAL SCALE

not disclosed

WHAT IS SOLD

An app-retargeting-specialist DSP: re-engagement/reactivation campaigns for gaming, e-commerce, delivery, and finance apps, explicitly combining 'fully managed service capabilities alongside programmatic advertising' (V2-S06) rather than positioning as purely self-serve.

BUYER AND APP FIT

ECONOMICS

DATA REQUIREMENTS

Ranked by AppsFlyer's Performance Marketing Index as a top retargeting DSP for Android gaming and among the top 6 e-commerce DSPs globally (V2-S06) — vendor-cited third-party ranking, best read as a relative-performance signal within AppsFlyer's own panel, not a market-wide claim.

CPM only: 'Remerge exclusively offers a CPM pricing model'. Reported cost is the campaign fee paid by the client, with no media-cost versus margin split [RS-S100; RS-S99].

Retargeting is built on MMP event forwarding and MMP audience segments; user lists can exclude users or hold them out as a control [RS-S101; RS-S102; RS-S103]. Remerge states compatibility with SKAdNetwork 2.0 to 4.0 and forwards install-validation postbacks to the MMP; SKAdNetwork is not used for retargeting, and no AdAttributionKit support is documented [RS-S104].

SUPPLY DEPENDENCIES

Not documented in the sources we read; no exchange or SDK detail is given.

MEASURABLE CONTROLS

Managed: Remerge's team builds campaigns after an insertion order, and clients get a reports dashboard and a Reporting API (aggregated JSON by hour or day, six-month retention) [RS-S101; RS-S111; RS-S108]. Its incrementality page presents Ghost Bids (always-on, for retargeting) and Causal Impact (without device IDs, using sub-market econometrics, for installs and re-engagement); intent-to-treat and PSA designs appear only in linked explainers [G1-S05]. The Reporting API has no publisher dimension; publisher breakdowns come through the MMP [RS-S107; RS-S108].

IMPLEMENTATION BURDEN

Described as a 'fully managed service' by Remerge's own May 2025 announcement (V2-S06), implying lower buyer-side operational burden but also less direct self-serve control.

PROOF QUALITY

Four case summaries on Remerge's incrementality page give no quantified lift: Miniclip retargeting, checked by Miniclip's own distribution matching, and Socialpoint, PhotoSi and HungerStation user acquisition [G1-S05]. AppsFlyer's Performance Marketing Index ranks Remerge highly for retargeting; that is a vendor-cited panel ranking, not a causal result [V2-S06].

SWITCHING LIMITS

Not documented in the sources we read.

RISKS AND CONFLICTS

Managed model: buyers rely on Remerge's team for setup and on its own uplift reporting; we found no independent audit of results [RS-S101; G1-S05].

POOR FIT

Buyers who want self-serve control (Remerge is managed) or whose main job is large-scale user acquisition: Remerge positions itself on retargeting, though its incrementality page includes user-acquisition cases [RS-S101; V2-S06; G1-S05]. No CTV product was found [RS-S112].

DIMENSION	FINAL	AVAILABILITY	GRADE	RATIONALE	SOURCES	SECOND SCORER	RECONCILIATIO
D1 Optimization objectives	3	live	A	Glossary defines ROAS and CPA 'achievement' against a Target ROAS/Target CPA for campaigns; onboarding asks for the in-app conversion event to optimize. CPM is the only pricing model.	RS-S99 RS-S100	3	R3 evidence found by one scorer
D2 Re-engagement & suppression	3	live	A	Retargeting core product built on MMP event forwarding and MMP audience segments; exclusion/blocklist user lists documented, and holdout groups supported via separate holdout lists.	RS-S101 RS-S102 RS-S103	3	R2 higher score has grade-A evidence
D3 iOS privacy measurement (SKAN / AAK)	2	live	A	States compatibility with SKAdNetwork 2.0-4.0; forwards install-validation postbacks with conversion value and source app ID to the MMP. No AdAttributionKit support found; SKAN not used for retargeting.	RS-S104	2	R3 evidence found by one scorer
D4 Creative workflow	3	limited	A	In-house design team offers curated A/B testing incl. no-ID incrementality tests; reports carry a 'Creative AB Test Group' (test vs control) dimension; rich media and Luna playables supported. Allocation details not published.	RS-S105 RS-S99 RS-S106	3	R3 evidence found by one scorer
D5 App & placement visibility	2	live	A	Tracking links pass a {publisher} macro to the MMP, enabling publisher breakdowns there; Remerge's own Reporting API has no publisher dimension. No advertiser block/allow-list control documented.	RS-S107 RS-S108	2	R3 evidence found by one scorer
D6 Fee & cost transparency	1	live	A	'Remerge exclusively offers a CPM pricing model'; reported cost is the campaign fee paid by the client. No media-cost versus margin breakdown.	RS-S100 RS-S99	1	R3 evidence found by one scorer

DIMENSION	FINAL	AVAILABILITY	GRADE	RATIONALE	SOURCES	SECOND SCORER	RECONCILIATIO
D7 Experiment / incrementality support	2	limited	B	Own blog names and defines three specific holdout/incrementality designs (Ghost Bids, PSA, intent-to-treat) — meets the anchor's 'documented holdout/lift method' bar, but the framing reads as educational/criteria-setting rather than an explicit statement that all three are offered as a standardized, self-serve product with results reporting to every client, so not scored 3.	V2-S07	3	R2 conservative: lower score
D8 Reporting export	2	live	A	Reporting API returns aggregated JSON (hour/day) by audience, country, campaign, ad; six-month retention. No log-level export documented.	RS-S108	2	R3 evidence found by one scorer
D9 Buyer control (API & self-serve)	1	live	B	Remerge's own May 2025 announcement states it offers 'fully managed service capabilities alongside programmatic advertising' — managed model emphasized; no self-serve UI or campaign-management API documentation found to support a higher score.	V2-S06	1	R1 agree
D10 CTV-to-app	n/e	unconfirmed	C	No CTV product or claim found for Remerge in this research (unlike Jampp, YouAppi, Adikteev, and Verve Dataseat, all of which document CTV-to-app products).		n/e	R1 agree
D11 Non-gaming vertical evidence	2	live	B	Named non-gaming case studies: Delivery Hero, Gojek, Swiggy Instamart, AliExpress, Trendyol, Zooplus and Omio. Private company with no filing-level revenue evidence, so the anchor for 3 is not met.	RS-S114 RS-S109 RS-S115	2	R6 fact-check override: Correction of 28 Sep 2026 (COR-117): the second scorer's documented named non-gaming cases meet the anchor; rul R2 had kept the lower first-pass score.

Jampp

Affle 3i Limited (NSE/BSE: AFFLE) · roles: DSP · profile group V2

OWNERSHIP

Wholly owned by Affle. The 100% acquisition was announced on 9 June 2021 and completed on 1 July 2021; the price was not disclosed [V2-S23]. Affle owns at least four app-marketing brands: Jampp, YouAppi, RevX and mediasmart [FX1-S20; FX1-S19; V2-S31; V2-S30].

FINANCIAL SCALE

Not disclosed. Jampp is not broken out separately in the Affle sources we read.

WHAT IS SOLD

A programmatic mobile DSP for user acquisition and retargeting, explicitly optimizing toward advertiser-set ROAS goals ('Set your goals and our algorithms will adjust everything to ensure you are ROAS positive'), plus a separate Jampp CTV product (launched 2023-07-20) linking CTV exposure to app downloads/re-engagement via MMP integration.

BUYER AND APP FIT

Performance app marketers who want a self-serve dashboard alongside managed creative production. Jampp advertises always-on lift measurement with 'Ghost Bids' for acquisition and retargeting, free from day one, but does not document the method [V2-S22].

ECONOMICS

Pricing model (CPI/CPA/CPM/% of spend) not explicitly disclosed on the product page reviewed; ROAS-goal framing implies performance-based billing but no rate card found.

DATA REQUIREMENTS

Documented support for SKAN, probabilistic attribution, and IDFA (when available) (V2-S22); AdAttributionKit support not mentioned in the page reviewed.

SUPPLY DEPENDENCIES

Not detailed beyond 'programmatic in-app' access; no owned-SDK-supply claim found (contrasts with sibling YouAppi/mediasmart, which document broader cross-screen supply).

MEASURABLE CONTROLS

Dashboard for direct campaign management plus a documented Reporting API and log-level data access 'for data science teams' (V2-S22); Dynamic Creative Optimization (DCO), feed-based creative, and HTML5 interactive-format testing also documented.

IMPLEMENTATION BURDEN

Hybrid model: buyer can self-serve via dashboard/API, while Jampp's in-house design team handles creative production — moderate burden, split between buyer and vendor.

PROOF QUALITY

Jampp's blog reports an 86% rise in new riders, 122% more rides and a 30% campaign lift for FREENOW, and a 92% conversion lift for Wallapop, under its always-on Ghost Bids measurement; it gives no period, baseline or uncertainty [G1-S06]. The lift product itself is advertised without a documented method [V2-S22]. Other named non-gaming cases (Wallapop, Instacart, Centauro, CashNow, Fetch) are Jampp's own case studies [RS-S121; RS-S122].

SWITCHING LIMITS

Not documented in the sources we read.

RISKS AND CONFLICTS

Jampp is one of at least four Affle app-marketing brands, with YouAppi, RevX and mediasmart; Jampp CTV runs on mediasmart. None of the materials we read says whether data or products are shared across the portfolio [FX1-S19; V2-S31; V2-S30; V2-S21].

POOR FIT

Buyers requiring documented AdAttributionKit support (not found) or CTV causal-measurement (Jampp CTV's own launch post does not claim an incrementality/causal method specific to CTV).

DIMENSION	FINAL	AVAILABILITY	GRADE	RATIONALE	SOURCES	SECOND SCORER	RECONCILIATIO
D1 Optimization objectives	3	live	B	Product page explicitly documents live ROAS-goal optimization ('ensure you are ROAS positive on your campaigns') as a standing, advertiser-facing feature.	V2-S22	3	R1 agree
D2 Re-engagement & suppression	2	live	B	Retargeting is a named, documented use case across the user journey; no explicit suppression/exclusion-list feature found to support score 3 (the documented Ghost Bids incrementality product is scored separately under D7, not treated here as re-engagement-specific suppression).	V2-S22	2	R1 agree
D3 iOS privacy measurement (SKAN / AAK)	1	live	B	Claims SKAN campaign-level reporting plus statistical modelling and conversion-value-to-ROAS modelling; no SKAN 4 or AdAttributionKit specifics documented.	V2-S22	1	R2 conservative: lower score
D4 Creative workflow	2	live	B	Documents Dynamic Creative Optimization, feed-based creative, and HTML5 interactive formats with a stated 'Build & Test, Measure, Iterate' process; no explicit randomized/disclosed-allocation testing design or named playable format confirmed, so not scored 3.	V2-S22	2	R1 agree
D5 App & placement visibility	2	live	B	Documents real-time dashboards (spend, conversions, ROAS) plus log-level data access; no explicit block/allow-list control feature found to support score 3.	V2-S22	2	R1 agree
D6 Fee & cost transparency	n/e	unconfirmed	C	No explicit media-cost-vs-platform-fee breakdown or published fee schedule found on the page reviewed.		n/e	R1 agree
D7 Experiment / incrementality support	2	live	B	Jampp advertises always-on lift measurement with ghost bids for acquisition and retargeting at no extra cost, but does not document the method, so the anchor for 3 (stated method) is not met.	V2-S22	3	R6 fact-check override: Fact-check! FX1 qualified V2-F0! the method is not documented and the 'modelled on Ghost Ads' wording is not on the page.
D8 Reporting export	3	live	B	Explicitly documents a 'reporting API' plus 'log-level data' provided to data-science teams — matches the anchor for log-level/impression-level export.	V2-S22	3	R1 agree
D9 Buyer control (API & self-serve)	2	live	B	Documents a dashboard for direct campaign management (self-serve); a documented reporting API exists, but no explicit campaign-management API (create/edit campaigns via API, not just reporting) was found to support score 3.	V2-S22	n/e	R3 evidence found by one scorer
D10 CTV-to-app	2	live	B	Jampp CTV documents tracking app downloads/re-engagements tied to CTV campaigns via MMP integration for real-time cost-per-action optimization; no stated identity method (deterministic/probabilistic/household) or CTV-specific causal/incrementality option found, so not scored 3.	V2-S21	2	R1 agree
D11 Non-gaming vertical evidence	2	live	B	Named non-gaming cases: Wallapop, Instacart, Centauro, CashNow, Fetch; Rakuten and Bigo cited in help center. No Jampp-level filing split in sources opened.	RS-S121 RS-S122	2	R3 evidence found by one scorer

Verve Dataseat

none (public company) · roles: exchange/SSP, DSP · profile group V2

OWNERSHIP

Public: Verve Group Media SE, formerly Verve Group SE and, before June 2024, MGI – Media and Games Invest SE (ticker VER; listed in Stockholm and Frankfurt) [V2-S10; FX1-S13]. Its 5 June 2026 AGM approved moving the registered office from Stockholm to Dublin; the move had not taken effect at the cutoff and was expected on 2 October 2026 [FX1-S13; FX1-S14]. The author was an executive at Verve Group from 2022 to 2024; Dataseat was scored under exactly the same rules as every other product.

BUYER AND APP FIT

Buyers who want SKAN-based iOS optimization and CTV-to-app campaigns [V2-S13]. Dataseat can be run in self-service mode or as a managed service [RS-

FINANCIAL SCALE

Group figures only: FY2025 revenue €550.9m and adjusted EBITDA €134.4m, about 24% of reported revenue and 22% of like-for-like revenue; some revenue moved from net to gross recognition in Q3 2025 [FX1-S12; AB-S09]. No Dataseat revenue is disclosed in the sources we read.

ECONOMICS

Dataseat says it charges 'a flat rate on your media spend, with no hidden costs', so the fee is separate

WHAT IS SOLD

Verve Dataseat is described on Verve's own site as a 'contextual performance DSP built to supercharge app growth — optimizing for installs, subscriptions, purchases, and high-LTV users without relying on device IDs,' positioned as one product line within Verve's broader omnichannel supply-side/exchange business (Jun Group, Smaato/Hybrid SDK, Captify).

DATA REQUIREMENTS

Documents SKAN postback ingestion and conversion-value mapping as inputs to its 'advertiser-adaptive AI'; when 'running with SKAN and MMP

S127]. It is the mobile DSP of an omnichannel group whose other businesses include supply-side units such as Smaato and Hybid [V2-S13].

from media; the rate is not published, and no minimum is stated [RS-S125].

data,' the platform is stated to use 'attributed installs, post-install events, and postbacks' — implying MMP dependency, though no specific MMP integrations were named in the pages opened.

SUPPLY DEPENDENCIES

Buyers can choose SSPs and publishers through inventory discovery [RS-S127]. Whether Dataseat buys the group's own Smaato and Hybid SDK supply is not confirmed in the sources we read.

MEASURABLE CONTROLS

Self-service mode or managed service covering setup and optimization [RS-S127]. No campaign-management API, reporting API or log specification is published; the site promises data 'at any level' without a documented report schema or block-list control [RS-S125; RS-S126; RS-S127].

IMPLEMENTATION BURDEN

Self-service or managed [RS-S127]. Sales run through a contact form, and we found no public help center or API documentation [V2-S13; RS-S125].

PROOF QUALITY

Two named non-gaming cases, both published by Verve: LinkedIn (app activations) and OTTO (a ROAS-focused SKAN campaign with a rebuilt conversion schema). Neither states a comparator, sample or independent check [RS-S126; RS-S123].

SWITCHING LIMITS

Not documented in the sources we read.

RISKS AND CONFLICTS

Structural: Verve sells supply through its exchange and SDKs and also buys media through Dataseat, a potential conflict of interest. No allegation of misconduct was found [V2-S13].

POOR FIT

Buyers wanting a neutral, demand-side-only DSP benchmark, or requiring a documented, named causal-measurement design (geo/PSA/ghost) for CTV-to-app claims — Verve's own page claims 'controlled incrementality testing' without naming the design.

DIMENSION	FINAL	AVAILABILITY	GRADE	RATIONALE	SOURCES	SECOND SCORER	RECONCILIATION
D1 Optimization objectives	3	live	B	Own product page explicitly documents optimization toward 'installs, subscriptions, purchases, ROAS, or LTV' as live, advertiser-selectable goals — meets the anchor for value/ROAS/LTV optimization documented as live; graded B (vendor page) and not upgraded despite specificity, per the author's stricter-scrutiny instruction.	V2-S13	3	R1 agree
D2 Re-engagement & suppression	2	live	B	Dataseat's own About page names Retargeting as one of three core solution areas (alongside User Acquisition and SKAN campaigns); no suppression/exclusion-control documentation found to support score 3.	V2-S09	2	R1 agree
D3 iOS privacy measurement (SKAN / AAK)	2	live	B	Documents SKAN postback ingestion, conversion-value mapping, and 'value-based outcome modeling' in detail; no AdAttributionKit mention found in the pages opened, so not scored 3.	V2-S13	2	R1 agree
D4 Creative workflow	n/e	unconfirmed	C	No creative-workflow or testing-design documentation found in this research.		n/e	R1 agree
D5 App & placement visibility	2	live	B	LinkedIn case cites optimizations 'based on publisher-level data'; inventory discovery lets buyers select SSPs and publishers; site promises data 'at any level'. No documented report schema or blocklist control.	RS-S126 RS-S127 RS-S125	2	R3 evidence found by one scorer
D6 Fee & cost transparency	2	live	B	States 'a flat rate on your media spend, with no hidden costs', i.e. fee separated from media; the rate itself is not published.	RS-S125	2	R3 evidence found by one scorer
D7 Experiment / incrementality support	1	live	B	Own page claims CTV spend is 'validated via controlled incrementality testing' but does not name a specific method (geo/PSA/ghost/ITT) or describe results reporting — meets only the anchor for 'claims incrementality without method.' Scored conservatively per the author's stricter-scrutiny instruction for Verve.	V2-S13	1	R1 agree
D8 Reporting export	n/e	unconfirmed	C	No reporting-export (API/log-level) documentation found in this research.		n/e	R1 agree
D9 Buyer control (API & self-serve)	2	live	B	Dataseat can be used 'in self-service mode' or managed service covering setup and optimization; no campaign-management API documented.	RS-S127	2	R3 evidence found by one scorer
D10 CTV-to-app	2	live	B	Own page documents CTV inventory tied to 'mobile installs and post-install revenue' with 'spend optimized through outcome-driven decisioning' and a reference to 'controlled incrementality testing' — meets the anchor for documented CTV-to-app reporting, but no identity method (deterministic/probabilistic/household matching) is named, so withheld from score 3 under stricter scrutiny.	V2-S13	2	R1 agree

DIMENSION	FINAL	AVAILABILITY	GRADE	RATIONALE	SOURCES	SECOND SCORER	RECONCILIATIO
D11 Non-gaming vertical evidence	2	live	B	Named non-gaming Dataseat cases: LinkedIn (app activations) and OTTO (e-commerce return). No Dataseat-level revenue split, so the anchor for 3 is not met.	RS-S126 RS-S123	2	R6 fact-check override: Correction of 28 Sep 2026 (COR-117): the second scorer's documented named non-gaming cases meet the anchor; rul R2 had kept the lower first-pass score.

References

Sources cited in the text, in order of first citation, with the address used. The full register of every source consulted is in the appendix and in data/tables/source-register.csv.

1. Apple Inc.. *Receiving postbacks in multiple conversion windows*. 2026. <https://developer.apple.com/documentation/storekit/receiving-postbacks-in-multiple-conversion-windows> [C-S02]
2. Google (Google Ads Help). *Choose a bid strategy for your App campaign*. n.d.. <https://support.google.com/google-ads/answer/12073727?hl=en> [EF-S02]
3. Apple Inc.. *Auto-Renewable Subscriptions (App Store)*. 2026. <https://developer.apple.com/app-store/subscriptions/> [GH-S02]
4. GameAnalytics. *2026 Mobile & PC Gaming Benchmarks*. 2026. <https://www.gameanalytics.com/reports/2026-mobile-pc-gaming-benchmarks> [GH-S11]
5. RevenueCat. *The State of Subscription Apps in 10 minutes: lessons, trends, and benchmarks for 2026*. 2026. <https://www.revenuecat.com/blog/growth/subscription-app-trends-benchmarks-2026> [GH-S08]
6. AppsFlyer. *Top 5 Data Trends of 2025 and Predictions for 2026*. 2026. <https://www.appsflyer.com/resources/reports/top-5-data-trends-report/> [MAIN-S01]
7. Sensor Tower. *State of Mobile 2026: App Spending Reaches \$167 Billion (State of Mobile 2026 report summary)*. 2026-01. <https://sensortower.com/blog/state-of-mobile-2026> [AB-S12]
8. TechCrunch, reporting Appfigures data. *App downloads declined again in 2025, but consumer spending soared to nearly \$156B*. 2026-01-14. <https://techcrunch.com/2026/01/14/app-downloads-declined-again-in-2025-but-consumer-spending-soared-to-nearly-156b/> [AB-S13]
9. AppLovin Corporation (Investor Relations). *AppLovin Announces Fourth Quarter and Full Year 2025 Financial Results*. 2026-02-11. <https://investors.applovin.com/news/news-details/2026/AppLovin-Announces-Fourth-Quarter-and-Full-Year-2025-Financial-Results/default.aspx> [AB-S02]
10. Mobvista Inc.. *Mobvista (1860.HK) Announces 2025 Full-Year Results: Revenue Surpasses \$2 Billion Driven by AI Innovation*. 2026-03-11. <https://www.mobvista.com/en/press-release/mobvista-2025-annual-results-mintegral-milestone-revenue-en> [AB-S08]
11. Unity Software Inc. / SEC EDGAR. *Unity Reports Fourth Quarter and Fiscal Year 2025 Financial Results (Exhibit 99.1 to Form 8-K)*. 2026-02-11. <https://www.sec.gov/Archives/edgar/data/1810806/000181080626000010/a2025q4ex-991.htm> [AB-S05]
12. Liftoff Mobile, Inc. / SEC EDGAR. *Liftoff Mobile, Inc. Form S-1/A (registration statement, round 2)*. 2026. https://www.sec.gov/Archives/edgar/data/0001850351/000119312526202384/iron_s-1a_1_round_2.htm [AB-S10]
13. Digital Turbine, Inc. (Investor Relations). *Digital Turbine Reports Fiscal 2026 Fourth Quarter and Fiscal Year 2026 Financial Results*. 2026-05. <https://ir.digitalturbine.com/news-events/press-releases/detail/700/digital-turbine-reports-fiscal-2026-fourth-quarter-and> [AB-S07]
14. Digital Turbine, Inc. / SEC EDGAR. *Digital Turbine FY2026 10-K, XBRL exhibit R31 (revenue recognition — AGP Marketplace vs Brand/Performance)*. 2026-05-26. <https://www.sec.gov/Archives/edgar/data/317788/000162828026038115/R31.htm> [G2-S12]
15. Liftoff Mobile, Inc.. *Liftoff Mobile, Inc. final prospectus (Rule 424(b)(4))*. 2026-06-04. https://www.sec.gov/Archives/edgar/data/1850351/000119312526257410/iron_424b4_round_2.htm [FX1-S05]
16. AppLovin Corporation. *AppLovin Corporation Form 10-Q for quarter ended June 30, 2026*. 2026-08-05. <https://www.sec.gov/Archives/edgar/data/1751008/000175100826000059/app-20260630.htm> [FX1-S01]
17. Mobvista Inc.. *Mobvista Inc. 2025 Annual Report (HKEX 1860)*. 2026-04-29. <https://assets-official.mobvista.com/v3/file-link/2026/05/14/2026042905634.pdf> [FX1-S11]
18. Liftoff Mobile, Inc. / SEC EDGAR. *Liftoff Mobile, Inc. Prospectus (Form 424B4)*. 2026-06. https://www.sec.gov/Archives/edgar/data/1850351/000119312526257410/iron_424b4_round_2.htm [FX4-S06]
19. Verve Group SE. *Verve Group SE delivers strong operational performance in Q4 2025... and publishes financial guidance for 2026*. 2026-01-26. <https://press.verve.com/verve-group-se-delivers-strong-operational-performance-in-q4-2025-driving-organic-growth-and-gross-margin-expansion-and-publishes-financial-guidance-for-2026> [AB-S09]
20. Verve Group SE. *Verve Group SE Annual and Sustainability Report 2025*. 2026-04. https://investors.verve.com/wp-content/uploads/2026/04/Verve_Annual_and_Sustainability_Report_2025_English.pdf [FX1-S12]
21. AppLovin Corporation / SEC EDGAR. *AppLovin Corporation Form 10-Q for the quarterly period ended June 30, 2026*. 2026-08-05. <https://www.sec.gov/Archives/edgar/data/1751008/000175100826000059/app-20260630.htm> [G2-S01]
22. Liftoff Mobile, Inc. / SEC EDGAR. *Liftoff Mobile, Inc. Form 10-Q for the quarterly period ended June 30, 2026*. 2026-08-13. <https://www.sec.gov/Archives/edgar/data/1850351/000162828026056398/lfto-20260630.htm> [G2-S09]
23. Mobvista Inc. / HKEXnews. *Mobvista Inc. 2026 Interim Report (used for G2-F15; same document as V1-S37)*. 2026-09-18. <https://www1.hkexnews.hk/listedco/listconews/sehk/2026/0918/2026091801490.pdf> [FX4-S07]
24. AppLovin Corporation (investor relations). *AppLovin Announces Second Quarter 2026 Financial Results*. 2026-08-05. <https://investors.applovin.com/news/news-details/2026/AppLovin-Announces-Second-Quarter-2026-Financial-Results/default.aspx> [G2-S02]
25. Unity Software Inc. / SEC EDGAR. *Unity Reports Second Quarter 2026 Financial Results (Exhibit 99.1 to Form 8-K)*. 2026-08-06. <https://www.sec.gov/Archives/edgar/data/1810806/000181080626000041/a2026q2ex-991.htm> [G2-S07]
26. AppLovin Corporation / SEC EDGAR. *AppLovin Corporation Form 10-K for fiscal year ended December 31, 2025*. 2026-02-19. <https://www.sec.gov/Archives/edgar/data/1751008/000175100826000010/app-20251231.htm> [AB-S01]
27. Google (Google Ads Help). *About Target ROAS bidding*. 2026. <https://support.google.com/google-ads/answer/6268637> [MAIN-S03]
28. Google LLC. *Google AdMob homepage*. n.d.. <https://admob.google.com/home/> [V3-S41]
29. AppLovin. *About AppLovin's Axon AI (legal disclosure page)*. 2026. <https://legal.applovin.com/about-applovin-axon-ai/> [KL-S12]
30. Muddy Waters Research. *MW is Short AppLovin (APP US)*. 2025-03-27. <https://muddywatersresearch.com/research/2025/mw-short-app/> [G3-S05]
31. Muddy Waters Research. *APP's Persistent Lies, Denying Use of Persistent IDs (APP US)*. 2025-05-07. <https://muddywatersresearch.com/research/2025/app-persistent-lies/> [G3-S06]
32. Lead Plaintiffs Northern California Pipe Trades Trust Funds and Monroe County Employees' Retirement System, via counsel. *Amended Complaint for Violations of the Federal Securities Laws, Brownback v. AppLovin Corporation, No. 4:25-cv-02772-HSG (N.D. Cal.)*. 2025-09-12. <https://storage.courtlistener.com/recap/gov.uscourts.cand.446917/gov.uscourts.cand.446917.61.0.pdf> [G3-S07]
33. Adam Foroughi / AppLovin (archived blog). *A Note from Our CEO: Discussing Web Advertising Opportunity and Unpacking Pixels*. 2025-03-27. <https://www.applovin.com/en/archived-blog/note-from-our-ceo-2> [FX5-S05]
34. Basil Shikin / AppLovin (archived blog). *Examination of e-commerce data practices*. 2025-03-31. <https://www.applovin.com/en/archived-blog/examination-of-e-commerce-data-practices> [FX5-S07]

35. AppLovin Corporation / U.S. SEC EDGAR. *AppLovin Corporation Form 10-Q for the quarterly period ended June 30, 2026*. 2026-08-05. <https://www.sec.gov/Archives/edgar/data/1751008/000175100826000059/app-20260630.htm> [G3-S10]
36. CourtListener / RECAP (PACER-sourced). *Brownback v. AppLovin Corporation, 4:25-cv-02772-HSG (N.D. Cal.) docket*. n.d.. <https://www.courtlistener.com/docket/69781362/brownback-v-applovin-corporation/> [FX5-S10]
37. Lead Plaintiffs (Robbins Geller Rudman & Dowd). *Lead Plaintiffs' Notice of Motion and Motion to Supplement the Amended Complaint (ECF 76), Brownback v. AppLovin*. 2026-01-12. <https://storage.courtlistener.com/recap/gov.uscourts.cand.446917/gov.uscourts.cand.446917.76.0.pdf> [FX5-S11]
38. Reuters (via Yahoo Finance). *AppLovin probed by US SEC over data-collection practices, Bloomberg News reports*. 2025-10-06. <https://finance.yahoo.com/news/applovin-probed-us-sec-over-195441930.html> [FX5-S13]
39. AppLovin Corporation / U.S. SEC EDGAR. *AppLovin Corporation Form 10-Q for the quarterly period ended September 30, 2025*. 2025-11-05. <https://www.sec.gov/Archives/edgar/data/1751008/000175100825000081/app-20250930.htm> [G3-S08]
40. AppLovin Corporation / U.S. SEC EDGAR. *AppLovin Corporation Form 10-K for fiscal year ended December 31, 2025*. 2026-02-19. <https://www.sec.gov/Archives/edgar/data/1751008/000175100826000010/app-20251231.htm> [G3-S09]
41. AppLovin Corporation / SEC EDGAR. *AppLovin Corporation Form 10-Q for the quarter ended March 31, 2026*. 2026-05-06. <https://www.sec.gov/Archives/edgar/data/1751008/000175100826000044/app-20260331.htm> [FX5-S08]
42. Stephen Talbot via Pomerantz LLP. *Class Action Complaint, Talbot v. AppLovin Corporation, 3:26-cv-10584 (N.D. Cal.)*. 2026-09-16. <https://storage.courtlistener.com/recap/gov.uscourts.cand.479176/gov.uscourts.cand.479176.1.0.pdf> [FX5-S12]
43. Unity Software Inc. / SEC EDGAR. *Unity Software Inc. Form 8-K reporting completion of ironSource merger, and Exhibit 99.1 press release*. 2022-11-07. <https://www.sec.gov/Archives/edgar/data/1810806/00011931252279161/d402948d8k.htm> [AB-S06]
44. AppLovin Corporation / SEC EDGAR. *AppLovin Corp Form 8-K, Exhibit 99.1: proposal to acquire Unity Software*. 2022-08-09. <https://www.sec.gov/Archives/edgar/data/1751008/00011931252215724/d385952dex991.htm> [AB-S18]
45. AppLovin Corporation (Investor Relations). *AppLovin Completes Sale of Mobile Gaming Business to Tripledot Studios*. 2025-07-01. <https://investors.applovin.com/news/news-details/2025/AppLovin-Completes-Sale-of-Mobile-Gaming-Business-to-Triplicated-Studios/default.aspx> [AB-S03]
46. Liftoff Mobile, Inc. (via PR Newswire, hosted on Yahoo Finance). *Liftoff Announces Pricing of Initial Public Offering*. 2026-06-03. <https://finance.yahoo.com/markets/stocks/articles/liftoff-announces-pricing-initial-public-002200678.html> [AB-S11]
47. Liftoff Mobile, Inc.. *Liftoff Mobile, Inc. Request to withdraw registration statement (Form RW, Rule 477)*. 2026-02-17. https://www.sec.gov/Archives/edgar/data/1850351/000119312526052860/iron_rw_-_rule_477.htm [FX1-S08]
48. LoopMe. *LoopMe acquires Chartboost from Zynga, accelerating mission to power brand advertising across digital ecosystem*. 2024-12-10. <https://loopme.ai/press-releases/loopme-acquires-chartboost-from-zynga-accelerating-mission-to-power-brand-advertising-across-digital-ecosystem/> [V2-S27]
49. RZR (formerly Aarki). *RZR homepage*. n.d.. <https://www.rzr.com/> [V2-S18]
50. Internet Archive. *Wayback Machine capture of rzr.com, 12 March 2026*. 2026-03-12. <https://web.archive.org/web/20260312121200/https://rzr.com/> [FX1-S22]
51. Affle. *Affle announces completion of Jampp acquisition*. 2021-07-01. https://affle.com/affle_news/affle-announces-completion-of-jampp-acquisition [FX1-S20]
52. Affle / YouAppi. *Nisan Schitrit Appointed CEO of YouAppi*. 2026-05-06. https://affle.com/affle_news/nisan-schitrit-appointed-ceo-of-youappi-to-lead-next-phase-of-ai-led-growth-and-global-expansion [FX1-S19]
53. RevX (Affle group). *RevX homepage*. n.d.. <https://www.revix.io/> [V2-S31]
54. mediasmart (Affle Iberia, S.L.). *mediasmart homepage*. n.d.. <https://www.mediasmart.io/> [V2-S30]
55. Affle. *Affle Announces Strategic Acquisition of AdColony Technology Assets and Trademark from DT*. 2026-06-15. https://affle.com/affle_news/affle-announces-strategic-acquisition-of-adcolony-technology-assets-and-trademark-from-dt [FX1-S18]
56. Chris Yu, Statsig/Amplitude blog. *Statsig + Amplitude: The drop on Phase 1*. 2026-06-17. <https://www.statsig.com/blog/statsig-amplitude-phase-1> [V3-S19]
57. Amplitude, Inc.. *Amplitude, Inc. Form 10-Q for quarter ended June 30, 2026*. 2026-08-06. <https://www.sec.gov/Archives/edgar/data/1866692/000119312526335706/ampl-20260630.htm> [FX1-S16]
58. Apple Inc.. *App Tracking Transparency (framework overview)*. 2026. <https://developer.apple.com/documentation/apptrackingtransparency> [C-S07]
59. Apple. *App Store Review Guidelines*. 2026. <https://developer.apple.com/app-store/review/guidelines/> [KL-S28]
60. Adjust. *Mobile app trends 2026 (blog announcement)*. 2026-02-18. <https://www.adjust.com/blog/mobile-app-trends-2026/> [C-S24]
61. ppc.land (reporting on Adjust's Mobile App Trends 2026 report). *Adjust's 2026 mobile app report: finance sessions up 21%, gaming CPI jumps 30%*. 2026-02. <https://ppc.land/adjusts-2026-mobile-app-report-finance-sessions-up-21-gaming-cpi-jumps-30/> [C-S25]
62. Adjust (via Business Wire, syndicated on Yahoo Finance). *Adjust: Mobile App Use Grew Globally In 2025 With Continued Move To Multi-Platform (press release)*. 2026-02-18. <https://finance.yahoo.com/news/adjust-mobile-app-grew-globally-130000772.html> [OA4-S03]
63. Adjust (Tiahn Wetzler). *ATT opt-in rates: The latest benchmarks by app category and country*. 2025-07-15. <https://www.adjust.com/blog/att-opt-in-rates-2025/> [OA4-S04]
64. Sagar Baviskar, Iffat Chowdhury, Daniel Deisenroth, Beibei Li, Daniel Sokol. *ATT vs. Personalized Ads: User's Data Sharing Choices Under Apple's Divergent Consent Strategies*. 2024-06. <https://www.bu.edu/dbi/files/2024/09/ssrn-4887872-ATT.pdf> [D-S07]
65. Apple Inc.. *AdAttributionKit (framework overview)*. 2026. <https://developer.apple.com/documentation/adattributionkit> [C-S04]
66. Apple Inc. (App Store developer site). *User Privacy and Data Use*. n.d.. <https://developer.apple.com/app-store/user-privacy-and-data-use/> [OA5-S01]
67. ads.apple.com. <https://ads.apple.com/app-store/help/attribution/0028-measuring-ad-performance>. n.d.. <https://ads.apple.com/app-store/help/attribution/0028-measuring-ad-performance> [RS-S151]
68. AppsFlyer. *AppsFlyer attribution model*. n.d.. <https://support.appsflyer.com/hc/en-us/articles/207447053-AppsFlyer-attribution-model> [D-S13]
69. Jampp. *Jampp product page*. n.d.. <https://www.jampp.com/product> [V2-S22]
70. Apple Inc.. *App Store user privacy and data use FAQ ('Can I fingerprint or use signals from the device to try to identify the device or a user?')*. n.d.. <https://developer.apple.com/app-store/user-privacy-and-data-use/> [G3-S13]
71. Adjust GmbH. *ATT & SKAN solutions (Help Center article comparing SKAdNetwork vs. ATT-based attribution, including probabilistic modeling)*. n.d.. <https://help.adjust.com/en/article/ios-att-and-skadnetwork> [G3-S15]
72. AppsFlyer. *AppsFlyer attribution model (Knowledge Base article, including the 'Probabilistic modeling' attribution method table)*. n.d.. <https://support.appsflyer.com/hc/en-us/articles/207447053-AppsFlyer-attribution-model> [G3-S14]
73. AppsFlyer. *AppsFlyer attribution model (Zendesk Help Center API record, article 207447053)*. 2026-08-12. https://support.appsflyer.com/api/v2/help_center/en-us/articles/207447053.json [FX5-S15]
74. Apple Inc.. *AdAttributionKit Updates (changelog)*. 2026. <https://developer.apple.com/documentation/updates/adattributionkit> [C-S06]
75. Apple Inc.. *SKAdNetwork (class reference)*. 2026. <https://developer.apple.com/documentation/storekit/skadnetwork/> [C-S01]

76. Apple Inc.. *Understanding AdAttributionKit and SKAdNetwork interoperability*. 2026. <https://developer.apple.com/documentation/adattributionkit/adattributionkit-skadnetwork-interoperability> [C-S05]
77. Apple Inc.. *SKAdNetwork release notes*. 2026. <https://developer.apple.com/documentation/storekit/skadnetwork-release-notes> [C-S03]
78. Apple. *AppImpression.handleView() (AdAttributionKit symbol, iOS 26.2+) and AAK symbol crawl*. 2025-12. [https://developer.apple.com/tutorials/data/documentation/adattributionkit/appimpression/handleview\(\).json](https://developer.apple.com/tutorials/data/documentation/adattributionkit/appimpression/handleview().json) [FX2-S05]
79. AppsFlyer. *Send SKAN and AdAttributionKit postback copies directly to AppsFlyer (iOS 15+)*. 2025-12-31. <https://support.appsflyer.com/hc/en-us/articles/4402320969617-Send-SKAN-and-AdAttributionKit-postback-copies-directly-to-AppsFlyer-iOS-15> [C-S19]
30. Adjust. *Set up SKAdNetwork and conversion values (iOS SDK docs)*. 2026. <https://dev.adjust.com/en/sdk/ios/features/skad/> [C-S20]
31. Singular. *How to Send SKAdNetwork & AdAttributionKit Postbacks to Singular*. 2025-11-13. <https://support.singular.net/hc/en-us/articles/4405749380379-How-to-Send-SKAdNetwork-AdAttributionKit-Postbacks-to-Singular> [C-S21]
32. Apple (Apple Developer Documentation). *AdServices*. n.d.. <https://developer.apple.com/documentation/adservices> [D-S20]
33. Apple (Apple Developer Documentation). *AAAttribution.attributionToken()*. n.d.. [https://developer.apple.com/documentation/adservices/aaattribution/attributiontoken\(\)](https://developer.apple.com/documentation/adservices/aaattribution/attributiontoken()) [D-S21]
34. Apple Inc.. *AdServices Changelog*. 2026-09. <https://developer.apple.com/documentation/adservices/changelog> [C-S10]
35. Apple Ads Help. *App ad attribution overview*. n.d.. <https://ads.apple.com/app-store/help/attribution/0094-ad-attribution-overview> [V3-S02]
36. Autorite de la concurrence (France). *Targeted advertising: the Autorite de la concurrence imposes a fine of EUR150,000,000 on Apple for the implementation of the App Tracking Transparency framework*. 2025-03-31. <https://www.autoritedelaconcurrence.fr/en/press-release/targeted-advertising-autorite-de-la-concurrence-imposes-fine-eu150000000-apple> [C-S11]
37. AGCM (Autorita Garante della Concorrenza e del Mercato, Italy). *The Italian Competition Authority fines Apple over 98 million euro for abuse of a dominant position (A561)*. 2025-12-22. <https://en.agcm.it/en/media/press-releases/2025/12/A561> [C-S13]
38. Bundeskartellamt (Germany). *Apple changes its rules for personalised advertising in apps*. 2026-08-17. https://www.bundeskartellamt.de/SharedDocs/Meldung/EN/Pressemitteilungen/2026/08_17_2026_Apple_ATTf.html [C-S12]
39. Apple Inc.. *requestTrackingAuthorization(usingExpandedInterface:additionalInformationAction:completionHandler:)*. 2026. [https://developer.apple.com/documentation/apptrackingtransparency/attrackingmanager/requesttrackingauthorization\(usingexpandedinterface:additionalinformationaction:completionhandler:\)](https://developer.apple.com/documentation/apptrackingtransparency/attrackingmanager/requesttrackingauthorization(usingexpandedinterface:additionalinformationaction:completionhandler:)) [C-S08]
30. Apple. *iOS & iPadOS 27.2 Beta 2 Release Notes*. 2026-09. <https://developer.apple.com/tutorials/data/documentation/ios-ipados-release-notes/ios-ipados-27.2-release-notes.json> [FX2-S01]
31. Google (Android Developers). *Advertising ID (training article)*. 2026. <https://developer.android.com/training/articles/ad-id> [C-S16]
32. Google (Google Play Help). *Google Play policy answer 6048248 (Advertising ID / AD_ID policy)*. 2026. <https://support.google.com/googleplay/android-developer/answer/6048248> [C-S18]
33. Google (Android Developers). *Google Play Install Referrer API*. 2026. <https://developer.android.com/google/play/installreferrer> [C-S17]
34. Google. *Update on Plans for Privacy Sandbox Technologies*. 2025-10-17. <https://privacysandbox.google.com/blog/update-on-plans-for-privacy-sandbox-technologies> [C-S15]
35. Google. *Privacy Sandbox feature status*. 2026. <https://privacysandbox.google.com/overview/status> [C-S14]
36. Evgeny Popov / No Fluff Advisory (author's own work). *Mobile App Growth Playbook*. 2026-06-05. <https://nofluffadvisory.com/services/playbooks/mobile-app-growth/> [CP-01]
37. U.S. Court of Appeals for the Ninth Circuit. *Epic Games Inc. v. Apple Inc., Opinion, No. 25-2935*. 2025-12-11. <https://cdn.ca9.uscourts.gov/datastore/opinions/2025/12/11/25-2935.pdf> [EF-S06]
38. Google (Play Console Help). *An update regarding Google Play's policies for developers serving users in the US*. n.d.. <https://support.google.com/googleplay/android-developer/answer/15582165?hl=en> [EF-S09]
39. Supreme Court of the United States. *Docket for No. 25-1311, Apple Inc. v. Epic Games, Inc.*. n.d.. <https://www.supremecourt.gov/search.aspx?filename=/docket/docketfiles/html/public/25-1311.html> [EF-S17]
30. Google LLC / Play Console Help. *Understanding Google Play's lower service fees*. 2026. <https://support.google.com/googleplay/android-developer/answer/1695462?hl=en> [GH-S06]
31. Federal Trade Commission / Federal Register. *Children's Online Privacy Protection Rule (2025 amendments)*. 2025-04-22. <https://www.federalregister.gov/documents/2025/04/22/2025-05904/childrens-online-privacy-protection-rule> [KL-S17]
32. Morrison Foerster (law firm alert). *Update: Texas App Store Law Takes Effect After Fifth Circuit Stays Preliminary Injunction*. 2026-06. <https://www.mofo.com/resources/insights/251111-texas-targets-app-stores-with-new-accountability-law> [KL-S24]
33. Meta (Meta for Developers). *App Events API*. n.d.. <https://developers.facebook.com/docs/marketing-api/app-event-api> [D-S16]
34. Meta (Meta for Developers). *Ad Campaign Group (Marketing API reference)*. n.d.. <https://developers.facebook.com/documentation/ads-commerce/marketing-api/reference/ad-campaign-group> [OA1-S06]
35. Meta (Meta Business Help Center). *About campaign attribution methods*. n.d.. <https://www.facebook.com/business/help/5574845785948422> [OA1-S05]
36. Meta (Meta Business Help Center). *Compare attribution settings in Meta Ads Manager*. n.d.. <https://www.facebook.com/business/help/854500742637772> [OA1-S04]
37. TikTok (TikTok for Business Help Center). *Attribution windows at the ad group level*. n.d.. <https://ads.tiktok.com/help/article/about-attribution-windows-at-the-ad-group-level> [D-S22]
38. Singular (Singular Help Center). *Facebook (Meta) Ads Attribution Integration*. n.d.. <https://support.singular.net/hc/en-us/articles/115003252706-Facebook-Meta-Ads-Attribution-Integration> [D-S24]
39. TikTok (TikTok for Business Help Center). *About Self-Attributing Network (SAN) integration with MMPs*. n.d.. <https://ads.tiktok.com/help/article/about-self-attribution-transition> [D-S23]
30. AppsFlyer. *What is reattribution window? (glossary)*. n.d.. <https://www.appsflyer.com/glossary/reattribution-window/> [D-S14]
31. Brett R. Gordon, Robert Moakler, Florian Zettelmeyer. *Close Enough? A Large-Scale Exploration of Non-Experimental Approaches to Advertising Measurement*. 2022-09. <https://arxiv.org/abs/2201.07055> [D-S02]
32. Brett R. Gordon, Florian Zettelmeyer, Neha Bhargava, Dan Chapsky. *A Comparison of Approaches to Advertising Measurement: Evidence from Big Field Experiments at Facebook*. 2019-04. <https://doi.org/10.1287/mksc.2018.1135> [D-S01]
33. Tom Blake; Chris Nosko; Steven Tadelis. *Consumer Heterogeneity and Paid Search Effectiveness: A Large Scale Field Experiment (NBER Working Paper 20171)*. 2014-05. https://www.nber.org/system/files/working_papers/w20171/w20171.pdf [OA3-S05]

14. Thomas Blake, Chris Nosko, Steven Tadelis. *Consumer Heterogeneity and Paid Search Effectiveness: A Large-Scale Field Experiment*. 2015-01. <https://doi.org/10.3982/ECTA12423> [D-S04]
15. Randall A. Lewis, Justin M. Rao. *The Unfavorable Economics of Measuring the Returns to Advertising*. 2015-11. <https://doi.org/10.1093/qje/qjv023> [D-S03]
16. Nils Wernerfelt, Anna Tuchman, Bradley T. Shapiro, Robert Moakler. *Estimating the Value of Offsite Data to Advertisers on Meta*. 2022-08. https://bfi.uchicago.edu/wp-content/uploads/2022/08/BFI_WP_2022-114.pdf [D-S06]
17. Guy Aridor, Yeon-Koo Che, Brett Hollenbeck, Maximilian Kaiser, Daniel McCarthy. *Evaluating the Impact of Privacy Regulation on E-Commerce Firms: Evidence from Apple's App Tracking Transparency*. 2025-05. <https://doi.org/10.1287/mnsc.2024.06600> [D-S05]
18. Nils Wernerfelt; Anna Tuchman; Bradley T. Shapiro; Robert Moakler. *Estimating the Value of Offsite Tracking Data to Advertisers: Evidence from Meta (Marketing Science 44(2):268-286)*. 2025-03. <https://doi.org/10.1287/mksc.2023.0274> [OA3-S02]
19. Reinhold Kesler. *The Impact of Apple's App Tracking Transparency on App Monetization*. 2023-08. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4090786 [D-S08]
20. Ali Goli, Jason Huang, David Reiley, Nickolai M. Riabov. *Measuring Consumer Sensitivity to Audio Advertising: A Long-Run Field Experiment on Pandora Internet Radio*. 2024-12. <https://arxiv.org/abs/2412.05516> [G1-S02]
21. Pei Pei Chen, Anna Guitart, Ana Fernández del Río, África Periañez (Yokozuna Data, a Keywords Studio). *Customer Lifetime Value in Video Games Using Deep Learning and Parametric Models*. 2018-11-28. <https://arxiv.org/abs/1811.12799> [G1-S04]
22. Guido Karmel, Jampp. *Why on-and-off lift measurement tests cost more than you think*. 2026-08-27. <https://www.jampp.com/blog/why-on-and-off-lift-measurement-tests-cost-more-than-you-think> [G1-S06]
23. Bradley T. Shapiro, Günter J. Hitsch, Anna E. Tuchman. *TV Advertising Effectiveness and Profitability: Generalizable Results From 288 Brands*. 2021. <https://doi.org/10.3982/ECTA17674> [G1-S09]
24. Nils Wernerfelt; Anna Tuchman; Bradley Shapiro; Robert Moakler. *Estimating the Value of Offsite Tracking Data to Advertisers: Evidence from Meta (NBER Working Paper 32765)*. 2024-08. https://www.nber.org/system/files/working_papers/w32765/w32765.pdf [OA3-S01]
25. Meta for Developers. *Conversion Lift Measurement (Marketing API guide)*. n.d.. <https://developers.facebook.com/docs/marketing-api/guides/lift-studies/v2.9> [V3-S06]
26. Google Ads Help. *About Conversion Lift*. n.d.. <https://support.google.com/google-ads/answer/12003020?hl=en> [V3-S07]
27. TikTok Ads Manager Help Center. *About Conversion Lift Study*. n.d.. <https://ads.tiktok.com/help/article/about-conversion-lift-study?lang=en> [V3-S09]
28. Google. *Google Ads API Release Notes (v25.1, including Conversion Lift and Brand Lift measurement support)*. 2026-08-19. <https://developers.google.com/google-ads/api/docs/release-notes> [G3-S20]
29. Moloco. *Moloco Performance CTV (solutions page)*. 2026. <https://www.moloco.com/solutions/ctv> [V1-S13]
30. Moloco. *Maximize lifetime value with Moloco's intelligent re-engagement*. 2026. <https://www.moloco.com/solutions/moloco-ads/re-engagement> [V1-S14]
31. Remerge. *Remerge Incrementality (product/methodology page with case studies: Miniclip 8 Ball Pool, Socialpoint Dragon City, PhotoSi, Delivery Hero HungerStation)*. n.d.. <https://www.remerge.io/incrementality> [G1-S05]
32. AppLovin (company blog). *Making sense of AppLovin through a measurement lens*. 2026. <https://applovin.com/en/blog/measurement-lens> [V1-S04]
33. Garrett A. Johnson, Randall A. Lewis, Elmar I. Nubbemeyer. *Ghost Ads: Improving the Economics of Measuring Online Ad Effectiveness*. 2017-12. <https://doi.org/10.1509/jmr.15.0297> [D-S09]
34. Garrett A. Johnson; Randall A. Lewis; Elmar I. Nubbemeyer. *Ghost Ads: Improving the Economics of Measuring Ad Effectiveness (conference draft)*. 2016-02-18. https://conference.nber.org/confer/2016/EoDs16/Johnson_Lewis_Nubbemeyer.pdf [OA3-S04]
35. Google (Google for Developers). *Meridian introduction*. n.d.. <https://developers.google.com/meridian/docs/basics/meridian-introduction> [D-S17]
36. Meta (facebookexperimental/Robyn). *An Analyst's Guide to MMM (calibration section)*. n.d.. <https://facebookexperimental.github.io/Robyn/docs/analysts-guide-to-MMM/> [D-S18]
37. PyMC Labs (PyMC-Marketing docs). *MMM library comparison*. n.d.. <https://www.pymc-marketing.io/en/stable/guide/mmm/comparison.html> [D-S19]
38. Google (Google for Developers). *Amount of data needed (Meridian pre-modeling guide, including 'Can I use campaign-level data?')*. 2026-06-03 (last updated). <https://developers.google.com/meridian/docs/pre-modeling/amount-data-needed> [OA3-S03]
39. Google (Google Ads Help). *Best practices for App campaigns*. n.d.. <https://support.google.com/google-ads/answer/14104492> [EF-S03]
40. Google (Google Ads Help). *About data-driven attribution*. n.d.. <https://support.google.com/google-ads/answer/6394265> [D-S10]
41. TikTok (TikTok Ads Manager Help). *About Smart+ App Campaigns*. n.d.. <https://ads.tiktok.com/help/article/about-smart-plus-app-campaigns> [EF-S04]
42. TikTok Ads Manager Help Center. *About Smart+ App Campaigns*. n.d.. <https://ads.tiktok.com/help/article/about-smart-plus-app-campaigns> [V3-S01]
43. TikTok (TikTok Ads Manager Help). *About Value-based Optimization for app*. 2026-05. <https://ads.tiktok.com/help/article/value-based-optimization-app?lang=en> [OA1-S13]
44. TikTok (TikTok Ads Manager Help). *Tips for Value-based Optimization for app*. 2026-05. <https://ads.tiktok.com/help/article/tips-for-value-based-optimization-for-app> [OA1-S15]
45. Meta (Meta Business Help Center). *About the learning phase*. 2026. <https://www.facebook.com/business/help/112167992830700> [MAIN-S05]
46. AppLovin. *Scale your campaign (Track & optimize)*. n/a. <https://support.applovin.com/en/growth/promoting-your-apps/track-and-optimize/scale-your-campaign> [V1-S07]
47. AppsFlyer. *The State of App Monetization - 2026 Edition*. 2026. <https://www.appsflyer.com/resources/reports/app-marketing-monetization-report/> [GH-S10]
48. RevenueCat. *State of Subscription Apps 2026*. 2026. <https://www.revenuecat.com/state-of-subscription-apps> [GH-S07]
49. RevenueCat. *RevenueCat blog: Ad Channel Diversification for Apps*. 2026. <https://www.revenuecat.com/blog/growth/ad-channel-diversification/> [EF-S22]
50. Digital Remedy. *How Cross-Platform Audience Duplication Analysis Reveals Overlap*. 2026. <https://www.digitalremedy.com/blog/how-cross-platform-audience-duplication-analysis-exposes-hidden-waste/> [EF-S23]
51. Garrett A. Johnson, Randall A. Lewis, Elmar I. Nubbemeyer. *Ghost Ads: Improving the Economics of Measuring Online Ad Effectiveness*. 2017-12. <https://doi.org/10.1509/jmr.15.0297> [G3-S01]
52. Garrett A. Johnson, Randall A. Lewis, Elmar I. Nubbemeyer. *Ghost Ads: Improving the Economics of Measuring Ad Effectiveness (working paper, 18 Feb 2016, NBER Economics of Digitization meeting)*. 2016-02-18. https://conference.nber.org/confer/2016/EoDs16/Johnson_Lewis_Nubbemeyer.pdf [FX5-S01]
53. Navdeep S. Sahni, Sridhar Narayanan, Kirithi Kalyanam. *An Experimental Investigation of the Effects of Retargeted Advertising: The Role of Frequency and Timing*. 2019-06. <https://doi.org/10.1177/0022243718813987> [G3-S02]
54. Andrey Simonov, Chris Nosko, Justin M. Rao. *Competition and Crowd-Out for Brand Keywords in Sponsored Search*. 2018-03. <https://doi.org/10.1287/mksc.2017.1065> [G3-S03]
55. Andrey Simonov, Chris Nosko, Justin M. Rao. *Competition and Crowd-Out for Brand Keywords in Sponsored Search (Marketing Science 37(2):200-215)*. 2018-03. https://business.columbia.edu/sites/default/files-efs/pubfiles/26238/competition_crowdout.pdf [FX5-S02]
56. Apple Inc. (App Store developer site). *Product page optimization*. n.d.. <https://developer.apple.com/app-store/product-page-optimization/> [OA5-S02]

57. Apple Inc.. *Product Page Optimization — Acquisition — App Store Connect Analytics*. n.d.. <https://developer.apple.com/help/app-store-connect-analytics/acquisition/product-page-optimization/> [EF-S01]
58. Google (Play Console Help). *Run A/B tests on your Store Listing*. n.d.. <https://support.google.com/googleplay/android-developer/answer/12053285?hl=en> [EF-S05]
59. Apple Inc. (App Store developer site). *Custom Product Pages*. 2026. <https://developer.apple.com/app-store/custom-product-pages/> [MAIN-S04]
60. Meta (Meta Business Help Center). *About A/B testing*. n.d.. <https://www.facebook.com/business/help/1738164643098669> [OA1-S08]
61. Meta (Meta Business Help Center). *About Dynamic Creative*. n.d.. <https://www.facebook.com/business/help/170372403538781> [OA1-S09]
62. Michael Braun, Eric M. Schwartz. *Where A/B Testing Goes Wrong: How Divergent Delivery Affects What Online Experiments Cannot (and Can) Tell You About How Customers Respond to Advertising*. 2025-01-09. <https://doi.org/10.1177/00222429241275886> [G1-S01]
63. Michael Braun; Eric M. Schwartz. *Where A-B Testing Goes Wrong: How Divergent Delivery Affects What Online Experiments Cannot (and Can) Tell You About How Customers Respond to Advertising (accepted preprint)*. 2024-08-01. https://braunm.github.io/assets/documents/papers/BraunSchwartz2025_preprint.pdf [FX4-S01]
64. Charles Shaw (arXiv preprint 2509.09758). *A Path Signature Framework for Detecting Creative Fatigue in Digital Advertising*. 2025-09-11. <https://arxiv.org/abs/2509.09758> [EF-S10]
65. Charles Shaw (WPP Data Science). *Creative Fatigue Screening with Path Signatures (arXiv 2509.09758v5)*. 2026-09-01. <https://arxiv.org/pdf/2509.09758v5> [OA3-S07]
66. Apple Inc.. *App Store Review Guidelines (section 3.2.2)*. n.d.. <https://developer.apple.com/app-store/review/guidelines/> [IJ-S13]
67. Google LLC / Play Console Help. *User Ratings, Reviews, and Installs (Google Play Developer Program Policy)*. n.d.. <https://support.google.com/googleplay/android-developer/answer/9898684?hl=en> [OA2-S04]
68. Google / Android Developers Blog. *Google Play's policy on incentivized ratings, reviews, and installs*. 2017-06-05. <https://android-developers.googleblog.com/2017/06/google-plays-policy-on-incentivized.html> [IJ-S14]
69. AppsFlyer Ltd.. *Deferred Deep Linking: how it works*. n.d.. <https://www.appsflyer.com/glossary/deferred-deep-linking/> [EF-S20]
70. Supreme Court of the United States. *Apple v. Epic Games, No. 25-1311, question presented (granted limited to Q1)*. 2026-06-30. <https://www.supremecourt.gov/qp/25-01311qp.pdf> [FX2-S21]
71. US District Court, N.D. Cal. (Gonzalez Rogers, J.), via CourtListener RECAP. *Order denying Apple's motion to stay proceedings (Dkt. 1706), Epic Games v. Apple, 4:20-cv-05640-YGR*. 2026-08-11. <https://storage.courtlistener.com/recap/gov.uscourts.cand.364265/gov.uscourts.cand.364265.1706.0.pdf> [FX3-S04]
72. Supreme Court of the United States (Kagan, J.), via CourtListener RECAP. *Order administratively staying proceedings, Apple, Inc. v. Epic Games, Inc., No. 26A194 (N.D. Cal. Dkt. 1707)*. 2026-08-12. <https://storage.courtlistener.com/recap/gov.uscourts.cand.364265/gov.uscourts.cand.364265.1707.0.pdf> [FX3-S03]
73. CourtListener (RECAP of PACER). *Epic Games v. Google, No. 3:20-cv-05671-JD (N.D. Cal.) docket*. 2026-09. <https://www.courtlistener.com/docket/17443962/epic-games-inc-v-google-llc/> [FX2-S24]
74. Android Developers Blog. *I/O 2026: What's new in Google Play (re-read)*. 2026-05-19. <https://android-developers.googleblog.com/2026/05/io-2026-whats-new-in-google-play.html> [FX2-S27]
75. OpenAI. *Developers can now submit apps to ChatGPT*. 2025-12-17. <https://openai.com/index/developers-can-now-submit-apps-to-chatgpt/> [FX2-S25]
76. Adapty. *State of In-App Subscriptions 2026*. 2026. <https://adapty.io/state-of-in-app-subscriptions/> [GH-S09]
77. Unity Technologies. *Understanding the Impact of Rewarded Ads on IAP, Retention, and Engagement*. n.d.. <https://unity.com/blog/understanding-the-impact-of-rewarded-ads-on-iap-retention-and-engagement> [GH-S15]
78. Ali Goli; Jason Huang; David Reiley; Nickolai M. Riabov. *Measuring Consumer Sensitivity to Audio Advertising: A Long-Run Field Experiment on Pandora Internet Radio (full PDF of G1-S02)*. 2024-12-07. <https://arxiv.org/pdf/2412.05516> [FX4-S02]
79. United States Court of Appeals for the Eighth Circuit (Judges Loken, Erickson, Kobes, per curiam). *Custom Communications, Inc. v. Federal Trade Commission, Nos. 24-3137, 24-3388, 24-3415, 24-3442 (8th Cir.)*. 2025-07-08. <https://ecf.ca8.uscourts.gov/opndir/25/07/243137P.pdf> [G3-S18]
80. Federal Trade Commission. *Negative Option Rule*. n.d.. <https://www.ftc.gov/legal-library/browse/rules/negative-option-rule> [G3-S19]
81. Federal Trade Commission / Federal Register. *Revision of the Negative Option Rule, Withdrawal of the Non-Compete Rule To Conform These Rules to Federal Court Decisions (FR Doc 2026-02866)*. 2026-02-12. <https://www.federalregister.gov/documents/2026/02/12/2026-02866/revision-of-the-negative-option-rule-withdrawal-of-the-cars-rule-removal-of-the-non-compete-rule-to> [FX5-S19]
82. Federal Trade Commission / Federal Register. *Rule Concerning the Use of Prenotification Negative Option Plans; ANPRM (FR Doc 2026-04952)*. 2026-03-13. <https://www.federalregister.gov/documents/2026/03/13/2026-04952/rule-concerning-the-use-of-prenotification-negative-option-plans> [FX5-S20]
83. Apple Inc.. *App Store Small Business Program*. 2026. <https://developer.apple.com/app-store/small-business-program/> [GH-S01]
84. Apple Inc.. *Alternative Terms Addendum for Apps in the EU (to the Apple Developer Program License Agreement), version dated December 17, 2025*. 2025-12-17. https://developer.apple.com/contact/request/download/alternate_eu_terms_addendum.pdf [OA2-S03]
85. Apple Inc.. *Changes for apps in the European Union (Apple Developer support page describing the Developer Program License Agreement update of 18 August 2026)*. 2026. <https://developer.apple.com/support/dma-and-apps-in-the-eu/> [GH-S03]
86. Google LLC / Play Console Help. *Service fees (Play Console Help, general schedule)*. 2026. <https://support.google.com/googleplay/android-developer/answer/112622> [GH-S05]
87. Stripe, Inc.. *Stripe Pricing (US)*. 2026. <https://stripe.com/pricing> [G2-S15]
88. RevenueCat. *State of Subscription Apps 2026*. 2026. <https://www.revenuecat.com/state-of-subscription-apps/> [G2-S16]
89. Xiaojing Wang, Tianqi Liu, Jingang Miao (Google). *A Deep Probabilistic Model for Customer Lifetime Value Prediction*. 2019-12-16. <https://arxiv.org/abs/1912.07753> [GH-S12]
90. Google Research (open-source repo). *google/lifetime_value (README)*. n.d.. https://github.com/google/lifetime_value/blob/master/README.md [GH-S13]
91. Singular. *SKAN predictive (9 steps to make SKAN predictive)*. n.d.. <https://www.singular.net/blog/skan-predictive/> [GH-S19]
92. Eric Benjamin Seufert, Mobile Dev Memo. *Back-testing LTV models*. 2015-11-02. <https://mobiledevmemo.com/back-testing-ltv-models/> [GH-S14]
93. Pei Pei Chen; Anna Guitart; Ana Fernandez del Rio; Africa Perianez (Yokozuna Data). *Customer Lifetime Value in Video Games Using Deep Learning and Parametric Models (full PDF of G1-S04)*. 2018-11-28. <https://arxiv.org/pdf/1811.12799> [FX4-S03]
94. Xinzhe Cao, Yadong Xu, Xiaofeng Yang. *Customer Lifetime Value Prediction with Uncertainty Estimation Using Monte Carlo Dropout*. 2024-11. <https://arxiv.org/abs/2411.15944> [G1-S07]
95. Xinzhe Cao; Yadong Xu; Xiaofeng Yang (Tencent). *Customer Lifetime Value Prediction with Uncertainty Estimation Using Monte Carlo Dropout (full PDF of G1-S07)*. 2024-11-24. <https://arxiv.org/pdf/2411.15944> [OA3-S06]
96. Congde Yuan. *SHORE: A Long-term User Lifetime Value Prediction Model in Digital Games*. 2025-06. <https://arxiv.org/abs/2506.10487> [G1-S08]
97. AppsFlyer. *Ad Fraud Protection Solution*. n.d.. <https://www.appsflyer.com/products/measurement/fraud-protection/> [IJ-S17]
98. Adjust. *What is SDK spoofing fraud? / Why do you need mobile fraud prevention?*. n.d.. <https://www.adjust.com/glossary/sdk-spoofing/> [IJ-S18]

39. Farooqi, Feal, Lauinger, McCoy, Shafiq, Vallina-Rodriguez (ACM Internet Measurement Conference). *Understanding Incentivized Mobile App Installs on Google Play Store*. 2020. <https://arxiv.org/abs/2010.01497> [IJ-S15]
40. Media Rating Council (MRC). *Digital Accreditation listing*. n.d.. <https://mediaratingcouncil.org/accreditation/digital> [IJ-S27]
41. CNBC. *Uber sues Fetch for ad fraud*. 2017-09-19. <https://www.cnbc.com/2017/09/19/uber-sues-fetch-for-ad-fraud.html> [IJ-S05]
42. PYMNTS. *Uber Sued Over Payment For Alleged Fraudulent Ads*. 2018-01. <https://www.pymnts.com/legal/2018/uber-lawsuit-fetch-media-ad-fraud/> [IJ-S06]
43. Phunware, Inc. (SEC EDGAR exhibit). *Settlement Agreement and Mutual General Release (Uber Technologies, Inc. v. Phunware, Inc. et al.)*. 2020-10-09. <https://www.sec.gov/Archives/edgar/data/1665300/000162828020016344/ex1012-settlementagreement.htm> [IJ-S04]
44. HUMAN Security. *State of the Union: IAB Tech Lab Supply Chain Standards Adoption*. n.d.. <https://www.humansecurity.com/learn/blog/state-of-the-union-iab-tech-lab-supply-chain-standards-adoption/> [IJ-S21]
45. HUMAN Security (Braedon Vickers); Internet Archive. *State of the Union: IAB Tech Lab Supply Chain Standards Adoption (Wayback Machine capture, 9 Oct 2025)*. 2023-05-09. <https://web.archive.org/web/20251009141800/https://www.humansecurity.com/learn/blog/state-of-the-union-iab-tech-lab-supply-chain-standards-adoption/> [OA4-S07]
46. AppsFlyer (Knowledge Base). *AppsFlyer attribution model / CTV, PC, and console platform attribution concepts*. n.d.. <https://support.appsflyer.com/hc/en-us/articles/4404083608849-CTV-PC-and-console-platform-attribution-concepts> [IJ-S25]
47. AppsFlyer. *CTV Attribution Platform*. n.d.. <https://www.appsflyer.com/products/measurement/ctv-attribution/> [IJ-S09]
48. Kochava. *Beyond Fragmented Signals: Unifying Linear TV and CTV Measurement (Kochava + Samba TV)*. n.d.. <https://www.kochava.com/blog/beyond-fragmented-signals-unifying-linear-tv-and-ctv-measurement/> [IJ-S10]
49. Federal Trade Commission. *VIZIO to Pay \$2.2 Million to FTC, State of New Jersey to Settle Charges It Collected Viewing Histories on 11 Million Smart Televisions without Users' Consent*. 2017-02-06. <https://www.ftc.gov/news-events/news/press-releases/2017/02/vizio-pay-22-million-ftc-state-new-jersey-settle-charges-it-collected-viewing-histories-11-million> [MAIN-S06]
50. MNTN. *Verified Visits - Precise CTV Attribution Technology*. n.d.. <https://mountain.com/performance-tv/attribution/> [IJ-S12]
51. Roku Advertising. *Measuring performance on Roku: pixels & events*. n.d.. <https://advertising.roku.com/learn/resources/measuring-performance-on-roku-pixels-events> [IJ-S11]
52. Moloco. *Moloco Launches AI-Powered Performance CTV for App Marketers*. 2026-04. <https://www.moloco.com/press-releases/ai-powered-performance-ctv> [IJ-S08]
53. Modern Retail. *AppLovin rebrands its ad platform as Axon, launches ads manager on referral-only basis*. 2025. <https://www.modernretail.co/marketing/applovin-rebrands-its-ad-platform-as-axon-launches-ads-manager-on-referral-only-basis/> [IJ-S23]
54. AppLovin (Adam Foroughi, CEO). *AppLovin Ads is now open to all advertisers*. 2026-06-22. <https://www.applovin.com/en/blog/applovin-ads-now-open> [OA4-S06]
55. ppc.land. *AppLovin's \$1.84B Q1 Beats Guidance as Axon Platform Opens to All in June*. 2026. <https://ppc.land/applovin-1-84b-q1-beats-guidance-as-axon-platform-opens-to-all-in-june/> [IJ-S22]
56. Moloco. *Moloco Commerce Media (product pages)*. n.d.. <https://www.moloco.com/solutions/mcm> [IJ-S24]
57. Evgeny Popov / No Fluff Advisory. *Proof and Story: The State of Contextual Advertising in 2026 (Contextual Quotient method)*. 2026-08-30. <https://nofluffadvisory.com/research/contextual-2026/> [MAIN-S07]
58. Google Ads API developer documentation. *App campaigns overview*. n.d.. <https://developers.google.com/google-ads/api/docs/app-campaigns/overview> [V3-S11]
59. Apple Ads Help. *Use the Campaign Management API*. n.d.. <https://ads.apple.com/app-store/help/campaigns/0022-use-the-campaign-management-api> [V3-S12]
60. TikTok (TikTok API for Business). *Create a campaign (campaign/create, API v1.3)*. n.d.. <https://business-api.tiktok.com/portal/docs/create-a-campaign/v1.3> [OA1-S12]
61. Amazon (Amazon Ads API documentation). *Getting started with the Amazon DSP Campaign Management API*. n.d.. <https://advertising.amazon.com/API/docs/en-us/guides/dsp/developer-guide> [OA1-S16]
62. Google (Google Ads Help). *Find or create placement reports for your App campaigns*. n.d.. <https://support.google.com/google-ads/answer/9141542?hl=en> [OA1-S01]
63. Meta. *Meta Business Help Center: About targeting and reporting for Advantage+ app campaigns (rendered in browser)*. n.d.. <https://www.facebook.com/business/help/1153577308409919> [FX1-S23]
64. TikTok Ads Manager Help Center. *About SKAN 4.0 and TikTok*. 2025-02. <https://ads.tiktok.com/help/article/about-skan-4-0-and-tiktok?lang=en> [V3-S08]
65. Amazon Ads. *Unlock the potential of app campaigns with Amazon DSP events manager*. 2024-02-13. <https://advertising.amazon.com/resources/whats-new/unlock-potential-of-app-campaigns-with-amazon-dsp-events-manager> [V3-S10]
66. AppsFlyer (Help Center). *[Beta] SKAN interoperability with Amazon*. n.d.. <https://support.appsflyer.com/hc/en-us/articles/27548444974353--Closed-beta-Special-access-SKAN-interoperability-with-Amazon> [OA1-S17]
67. AppLovin. *AppLovin - Adjust product page*. n.d.. <https://www.applovin.com/adjust/> [V3-S15]
68. Google / Firebase. *Firebase A/B Testing documentation*. n.d.. <https://firebase.google.com/docs/ab-testing> [V3-S29]
69. Sensor Tower (Oliver Yeh). *Sensor Tower acquires market intelligence platform data.ai*. 2024-03-18. <https://sensortower.com/blog/data-ai-joins-sensor-tower> [FX1-S17]
70. AppsFlyer Ltd. *AppsFlyer - About*. n.d.. <https://www.appsflyer.com/about/> [V3-S21]
71. Kochava Inc.. *Kochava homepage*. n.d.. <https://www.kochava.com/> [V3-S23]
72. Adjust. *Adjust InSight (Help Center article)*. n.d.. <https://help.adjust.com/en/article/insight> [OA4-S01]
73. Singular. *Audiences FAQ (Audience Incrementality section)*. 2026-09-02. <https://support.singular.net/hc/en-us/articles/360025454492-Audiences-FAQ> [OA4-S02]
74. Kayzen. *API Documentation*. n.d.. <https://help.kayzen.io/en/articles/3247742-api-documentation> [V2-S04]
75. Entravision Communications Corporation / SEC EDGAR. *Entravision Communications Corp — Form 10-K FY2025*. 2026. <https://www.sec.gov/Archives/edgar/data/1109116/000119312526093993/evc-20251231.htm> [V2-S08]
76. Remerge. *Remerge becomes an official AppsFlyer Premier Partner*. 2025-05-13. <https://www.remerge.io/blog-post/remerge-becomes-an-official-appsflyer-premier-partner> [V2-S06]
77. Association of National Advertisers (ANA). *ANA: In-house agency trend continues to gain steam (2023 study)*. 2023-05-02. <https://www.ana.net/content/show/id/79185> [KL-S16]
78. Association of National Advertisers (ANA). *ANA 2026 State of In-Housing report (via IHALC summary)*. 2026-07. <https://www.ihalc.com/insights/ana-report-finds-ihas-more-capable-more-strategic/> [KL-S15]
79. Google. *Google Ads API: App campaign reporting*. n.d.. <https://developers.google.com/google-ads/api/docs/app-campaigns/reporting> [FX1-S24]
80. Mobvista Limited / HKEX. *Mobvista Limited 2026 Interim Report*. 2026-09-19. <https://www1.hkexnews.hk/listedco/listconews/sehk/2026/0918/2026091801490.pdf> [V1-S37]
81. Google Ads Help. *How AI Max for Search campaigns works*. 2026. <https://support.google.com/google-ads/answer/15910187?hl=en> [KL-S11]
82. Unity Technologies. *Unity Ads newsletter — What's New (June 2026) / Vector product page*. 2026-06. <https://unity-ads-newsletter.github.io/> [KL-S14]
83. Google. *Google Ads API MCP server developer documentation*. 2026. <https://developers.google.com/google-ads/api/docs/developer-toolkit/mcp-server> [KL-S06]

14. Google Ads Developer Blog. *Open Source Google Ads API MCP Server (blog post)*. 2025-10-07. <https://ads-developers.googleblog.com/2025/10/open-source-google-ads-api-mcp-server.html> [FX2-S09]
15. Meta (Meta for Developers). *Ads MCP Server: Get started*. 2026-09-04. <https://developers.facebook.com/documentation/ads-commerce/ads-ai-connectors/ads-mcp-server/ads-mcp-server-get-started> [OA1-S11]
16. Meta. *Ads MCP Server overview (Meta for Developers)*. 2026. <https://developers.facebook.com/documentation/ads-commerce/ads-ai-connectors/ads-mcp-server/ads-mcp-server-overview> [KL-S07]
17. Meta (Meta for Developers). *Ads MCP Server: Ad creation and management (tools)*. 2026-07-14. <https://developers.facebook.com/documentation/ads-commerce/ads-ai-connectors/ads-mcp-server/ads-mcp-server-tools-ad-creation-and-management> [OA1-S10]
18. TikTok for Business. *About TikTok for Business Agentic Hub and MCP Server*. 2026-08. <https://ads.tiktok.com/help/article/about-tiktok-for-business-agentic-hub-and-mcp-server?lang=en> [KL-S08]
19. AgenticAdvertising.org. *AgenticAdvertising.org (AdCP homepage)*. 2026. <https://adcontextprotocol.org/> [KL-S01]
50. AgenticAdvertising.org. *AdCP CHANGELOG.md*. 2026-09. <https://raw.githubusercontent.com/adcontextprotocol/adcp/main/CHANGELOG.md> [FX2-S12]
51. AgenticAdvertising.org / adcontextprotocol. *AdCP release v3.1.24 (GitHub release tag)*. 2026-09-23. <https://github.com/adcontextprotocol/adcp/releases/tag/v3.1.24> [OA2-S05]
52. AgenticAdvertising.org. *AdCP docs — planning/execution and audit trail description*. 2026. <https://docs.adcontextprotocol.org/> [KL-S03]
53. AgenticAdvertising.org / adcontextprotocol. *AdCP Campaign Governance specification (docs/governance/campaign/specification.mdx) at tag v3.1.24*. 2026-09-23. <https://raw.githubusercontent.com/adcontextprotocol/adcp/v3.1.24/docs/governance/campaign/specification.mdx> [OA2-S06]
54. Ad Context Protocol (AdCP) / Agentic Advertising governance working group. *AdCP Governance protocol (docs version 3.1.24)*. n.d.. <https://docs.adcontextprotocol.org/dist/docs/3.1.24/governance/overview> [G3-S16]
55. AgenticAdvertising.org / adcontextprotocol. *AdCP CHANGELOG.md at tag v3.1.24*. 2026-09-23. <https://raw.githubusercontent.com/adcontextprotocol/adcp/v3.1.24/CHANGELOG.md> [OA2-S07]
56. IAB Tech Lab. *AAMP: Agentic Advertising Management Protocols*. 2026-09-22. <https://iabtechlab.com/standards/aamp-agentic-advertising-management-protocols/> [KL-S35]
57. IAB Tech Lab. *IABTechLab/AAMP repository (component list)*. 2026. <https://github.com/IABTechLab/AAMP> [FX2-S11]
58. IAB Tech Lab. *IAB Tech Lab Agentic RTB Framework (ARTF) — GitHub repository*. 2025-11. <https://github.com/IABTechLab/agentic-realtime-framework> [KL-S04]
59. IAB Tech Lab (Agentic Initiative). *AAMP 2.0 Release Brings Transaction-ready Buyer and Seller Agent SDKs*. 2026-04-23. <https://iabtechlab.com/aamp-2-0-release-brings-transaction-ready-buyer-and-seller-agent-sdks/> [G3-S17]
30. FTC / Federal Register. *Children's Online Privacy Protection Rule, final rule amendments (full text)*. 2025-04-22. https://www.federalregister.gov/documents/full_text/html/2025/04/22/2025-05904.html [FX2-S14]
31. Google (Play Console Help). *Google Play Data safety & Families policy (ads requirements)*. 2026. <https://support.google.com/googleplay/android-developer/answer/9893335> [KL-S30]
32. CourtListener (RECAP of PACER). *CCIA v. Paxton, No. 26-50001 (5th Cir.) docket*. 2026-09. <https://www.courtlistener.com/docket/72099057/computer-communications-industry-association-v-ken-paxton/> [FX2-S17]
33. Utah Legislature. *Utah H.B. 498 (2026) enrolled text and status data*. 2026-03-12. <https://le.utah.gov/Session/2026/bills/enrolled/HB0498.xml> [FX2-S15]
34. Louisiana Legislature. *Louisiana Act No. 185 (2026 RS, HB977) enrolled act and Resume Digest*. 2026-05-15. <https://legis.la.gov/legis/ViewDocument.aspx?d=1475238> [FX2-S18]
35. California Legislature. *AB 1043 Digital Age Assurance Act (Chapter 675, Statutes of 2025)*. 2025-10-13. https://leginfo.ca.gov/faces/billNavClient.xhtml?bill_id=20250260AB1043 [FX2-S26]
36. Apple. *Declared Age Range framework documentation (JSON)*. 2026. <https://developer.apple.com/tutorials/data/documentation/declaredagerange.json> [FX2-S19]
37. Google (Android Developers). *Play Age Signals overview*. 2026. <https://developer.android.com/google/play/age-signals/overview> [KL-S29]
38. Federal Trade Commission. *FTC Finalizes Order with X-Mode and Successor Outlogic*. 2024-04. <https://www.ftc.gov/news-events/news/press-releases/2024/04/ftc-finalizes-order-x-mode-successor-outlogic-prohibiting-it-sharing-or-selling-sensitive-location> [KL-S20]
39. Federal Trade Commission. *FTC Finalizes Order with InMarket Prohibiting It from Selling or Sharing Precise Location Data*. 2024-05-01. <https://www.ftc.gov/news-events/news/press-releases/2024/05/ftc-finalizes-order-inmarket-prohibiting-it-selling-or-sharing-precise-location-data> [KL-S21]
70. Federal Trade Commission. *FTC to Ban Kochava and Subsidiary from Selling Sensitive Location Data*. 2026-05-04. <https://www.ftc.gov/news-events/news/press-releases/2026/05/ftc-ban-kochava-subsidiary-selling-sensitive-location-data-settle-charges-they-sold-location-data> [KL-S19]
71. U.S. District Court for the District of Idaho (Winmill, J.), hosted by the Federal Trade Commission. *Stipulated Order for Injunction and Other Relief, FTC v. Kochava, Inc., No. 2:22-cv-00377-BLW (D. Idaho), Dkt. 138*. 2026-06-25. https://www.ftc.gov/system/files/ftc_gov/pdf/Kochava-Order.pdf [OA2-S01]
72. Federal Trade Commission. *FTC v Kochava, Inc. (timeline item) - June 26, 2026*. 2026-06-26. <https://www.ftc.gov/legal-library/browse/cases-proceedings/ftc-v-kochava-inc-timeline-item-2026-06-26> [OA2-S02]
73. European Parliament and Council of the European Union / EUR-Lex. *Regulation (EU) 2022/1925 of the European Parliament and of the Council of 14 September 2022 on contestable and fair markets in the digital sector (Digital Markets Act), Article 6*. 2022-09-14. <https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:32022R1925> [G3-S12]

© 2026 Evgeny Popov · No Fluff Advisory. Research cutoff 27 September 2026. This is AI-assisted research by the author; ChatGPT (OpenAI) assisted with technical and editorial review and source verification. The author retains responsibility for the paper. This assistance does not constitute independent assurance or endorsement by OpenAI. It has not been peer reviewed. Built from the canonical dataset in the package. Every figure has a CSV; every chart also has a static SVG, while Figure 14 (the matrix) and Figure 16 (the scenarios) are HTML tables with CSVs. Readability of main prose: Flesch Reading Ease 57.6 over 19,791 words; the original target of 85 was not met and this is a disclosed exception for the technical paper. The plain-language edition carries the target of 85.