Advertising & Marketing: How Things Are Sold

Figures current as of November 2026

How advertising works, who pays for it, who sells it, and what the reader gives up in exchange for the free services it funds. The course covers what marketing is and how big the business is; its history, from the first newspaper agents to the feed sold by auction; brands, and the evidence on whether advertising works; and the advertisers, the agencies, and the platforms. It then turns to retail media, television, streaming, and creators; how ads are bought, measured, and faked; personal data and the end of tracking; the rules for truth, children, drugs, politics, and vice; the antitrust cases against Google, Meta, and Amazon; and what artificial intelligence is changing.

MODULE 1 — What Marketing Is, and How Big It Is

Goal: Learn what marketing covers, where advertising fits within it, and how much the world spends.

Marketing, advertising, promotion, publicity, and selling; the four Ps and the funnel; brand building and performance; owned, earned, and paid media; and the size of the business by country and channel.

MODULE 2 — Selling to the Masses, 1841–1994

Goal: Trace how advertising became an industry in the age of print, radio, and television.

The first agents and the commission; patent medicines and the first food and drug law; national brands; the copywriters, the psychologists, and public relations; radio's sponsors and television's spots; the creative revolution and its critics; and the first rules.

MODULE 3 — The Digital Turn, 1994–2026

Goal: Follow the shift from selling space and time to selling single impressions by auction.

The cookie and the banner; the pay-per-click search auction; social feeds and mobile; programmatic trading; the money's move away from print and television; the privacy backlash; and the arrival of AI.

MODULE 4 — Brands and Persuasion: Why People Pay More

Goal: Understand what brands are worth and how advertising works on memory, and weigh whether it informs or manipulates.

Trademarks, signals, and memories; the brand rankings and why they disagree; store brands and informed shoppers; mental availability and distinctive assets; and the economics of information and persuasion.

MODULE 5 — Does Advertising Work? The Evidence

Goal: Learn what experiments and large studies show about advertising's effects, and weigh whether most of it pays for itself.

Why returns are hard to measure; eBay's search experiment; Facebook's experiments; television advertising for 288 brands; the meta-analyses; Super Bowl ads; and the practitioners' evidence.

MODULE 6 — Who Buys Advertising, and How a Budget Is Set

Goal: Learn who spends the most and how budgets are set, and weigh whether marketing has tilted too far toward the short term.

The largest advertisers; small businesses on self-serve platforms; budget rules, from share of sales to share of voice; media mix models; brand against performance; and a launch budget worked across channels.

MODULE 7 — The Agencies and Media Buying

Goal: Understand what agencies do, how they are paid, and why the business is consolidating and shrinking.

The holding companies after Omnicom bought Interpublic; WPP's and Dentsu's troubles; from the commission to fees, principal trading, and rebates; in-house agencies; and the pressure from AI.

MODULE 8 — Search, Social, and Video: The Platforms

Goal: Learn how the largest sellers of advertising make their money, and how concentrated the market has become.

Google's search, YouTube, and network; Meta's apps; ByteDance and TikTok's American joint venture; Microsoft, Apple, and the smaller platforms; China's sellers; and how each sells its ads.

MODULE 9 — The Stores Become Media: Retail Media

Goal: Understand why retailers became advertising companies, and what it costs the brands that sell through them.

Sponsored listings and closed-loop measurement; Amazon's ad business; Walmart, Instacart, Uber, and the rest; why the margins matter; and a seller's economics on Amazon.

MODULE 10 — How Ad Tech Works: The Auction

Goal: Follow an ad from the advertiser's bid to the screen, and see where the money goes.

The search, social, and programmatic auctions; identity and targeting signals; the supply chain's fees; Google's place in it, as a court found; The Trade Desk and Amazon; and the suit over Amazon's ad auctions.

MODULE 11 — Television, Streaming, Audio, and Outdoor

Goal: Understand how television and the older media are bought and priced as viewing moves to streaming.

Linear television's decline; the upfronts; the Super Bowl; the streaming services' ad tiers and falling prices; YouTube on the television screen; and radio, podcasts, billboards, and print.

MODULE 12 — Influencers, Creators, and Sponsorship

Goal: Learn how brands pay creators and sponsor events, and what the disclosure rules require.

The creator economy and its spending estimates; what YouTube and Meta pay; influencer rates; the endorsement and review rules; and sponsorship, from jersey patches to the World Cup.

MODULE 13 — Counting: Ratings, Attribution, and Fraud

Goal: Understand how advertising is measured and faked, and weigh the fight over brand safety.

Television ratings and their rivals; viewability and attention; attribution models and lift tests; platforms grading their own results; ad fraud; and brand safety, from the advertisers' alliance to the FTC's orders.

MODULE 14 — Personal Data, Privacy, and the End of Tracking

Goal: Learn how targeting uses personal data and how laws and companies have limited it, and weigh whether targeting is worth its privacy cost.

First-party and third-party data; cookies and device IDs; Apple's tracking prompt; Chrome's reversal; Europe's privacy law and Meta's consent fights; America's state laws; data brokers; and discrimination in ad delivery.

MODULE 15 — The Rules: Truth, Children, Drugs, Politics, and Vice

Goal: Learn who polices advertising and what it forbids, and weigh whether prescription drugs should be advertised on television.

Deception and substantiation; competitors' suits and self-regulation; advertising to children; prescription drug ads and the stricter enforcement of 2025; political ads and deepfakes; tobacco, alcohol, and gambling; and Europe's platform rules.

MODULE 16 — Antitrust: The Cases Against the Platforms

Goal: Follow the monopoly cases against Google, Meta, and Amazon, and weigh whether the platforms overcharge advertisers.

Two-sided markets and self-preferencing; the two Google rulings and their remedies; Europe's adtech fine; the FTC's loss against Meta; Amazon's coming trial; and Britain, Japan, and Canada.

MODULE 17 — AI and the Next Shift

Goal: Understand how AI is changing who makes ads, where they appear, and who gets paid, and weigh whether shopping agents will shrink advertising.

Automated buying and machine-made ads; ads inside AI answers; ChatGPT's ad business; fewer clicks for publishers; and AI agents that shop.

MODULE 18 — Where Things Stand, and What to Watch

Goal: Close with the argument, the contradictions, and a dated watch list.

The course's argument on a page; the contradictions no single module states; what to watch from November 2026 on; and a last word.

A note on the sources: forecasters count advertising differently, some including retail media, marketplace listings, or political ads and some not, so the course names whose figure it gives. Companies' figures come from their earnings releases and filings, mostly for the third quarter of 2026; figures for private companies, such as ByteDance, are labeled as reported. Evidence on whether advertising works comes from academic papers, industry bodies, and the platforms themselves, and the course says which, since the platforms sell what they measure. Prices are averages from benchmark firms and vary widely by audience and season. Leaders' ages are as of November 1, 2026.

The course explains how the industry works and gives no advice on buying advertising or shares. Where the field argues, over whether advertising informs or manipulates, whether most of it pays, brand against performance spending, brand safety, targeting's privacy cost, drug advertising, the platforms' market power, and AI shopping agents, it makes each side's case as its advocates make it before giving its own reading once.

Advertising at a Glance, as of November 2026

The size of the market, who sells, who buys, and what it costs

The Market

Measure	Figure	Period	Source and release
Global advertising revenue	$1.14 trillion, up 8.8%, excluding US political ads; about $1.3 trillion forecast for 2026, up 8.9%	2025; 2026	WPP Media, *This Year Next Year*, December 2025 and June 2026
Global advertising, a second count	$1.34 trillion, up 11.9%; $1.46 trillion forecast for 2027	2026	WARC, October 8, 2026
United States	$431.2 billion, up 12.5%	2025	WPP Media, December 2025
China	$216 billion, up 6.8%, the second-largest market	2025	WPP Media, December 2025
US internet advertising	$294.6 billion, up 13.9%; the ten largest sellers took 84.1%	2025	The Interactive Advertising Bureau (IAB) and PwC, April 16, 2026
By channel, worldwide	Social $394.6 billion; search $295.7 billion; retail media $202.1 billion; television $176.2 billion; out-of-home $63.3 billion; print $44.4 billion; audio $41.9 billion	2026 forecast	WARC, October 8, 2026
The five largest sellers	About 58% of the world's ad revenue: Google, Meta, ByteDance, Amazon, and Alibaba	2026 forecast	WPP Media, June 2026
Marketing budgets	7.8% of company revenue; paid media 31.4% of the budget	2026	Gartner CMO Spend Survey, May 2026
US political advertising	$11.21 billion for the 2025-26 election cycle, a record	To October 27, 2026	AdImpact

WPP Media and WARC include retail media and marketplace listings in their counts; Dentsu and Magna count less and put 2026 near $1 trillion (Module 1).

The Sellers and the Buyers

Company	Figure	Period	Source and release
Alphabet (Google, YouTube)	$294.7 billion of advertising revenue; $82.0 billion in the third quarter of 2026, up 10.6%	2025; Q3 2026	Alphabet, February 4 and October 28, 2026
Meta (Facebook, Instagram, WhatsApp)	$196.2 billion; $63.1 billion in the third quarter, up 26%	2025; Q3 2026	Meta, January 28 and October 21, 2026
Amazon	$68.6 billion; $19.6 billion in the third quarter, up 11%	2025; Q3 2026	Amazon's quarterly figures, as reported
ByteDance (TikTok, Douyin)	Private: about $200 billion of total revenue, as reported	2025	*The Information*, September 15, 2026
Netflix	More than $1.5 billion of advertising; about $2.7 billion expected for 2026	2025; 2026	Netflix, January and October 20, 2026
The largest buyers	Alibaba $30.2 billion of advertising worldwide; Amazon $23.5 billion; PDD $16.3 billion; L'Oréal $16.0 billion; P&G $13.5 billion (estimates)	2025	Ad Age, *Brand Report 2026*, October 14, 2026

Company	Figure	Period	Source and release
The agency groups	Omnicom $26.5 billion; Accenture Song $20.0 billion; Publicis $19.7 billion; WPP $17.9 billion; Dentsu $9.6 billion; Havas $3.3 billion	2025 revenue	Ad Age Agency Report, June 2026

The Prices

What	Price	When	Source
A 30-second Super Bowl spot	About $8 million, or about $64 per thousand of the 125.6 million viewers	February 8, 2026	NBCUniversal, as reported; Nielsen
A 30-second spot on *Sunday Night Football*	About $1.0 million	2024-25 season	Ad Age
Prime-time television, per thousand adults	Broadcast $41.65; cable $17.70; streaming $25.90	2026-27 upfront	Media Dynamics, August 2026
A click on a search ad, Google and Microsoft	$5.42 on average, with 8.18% of clicks turning into a sale or lead	April 2025-March 2026	WordStream and LocaliQ, via Shopifreaks, May 19, 2026
A click on a Facebook ad seeking leads	$1.80 on average; $27.39 per lead	2026	WordStream and LocaliQ, September 23, 2026
A click on an Amazon ad	$1.07 at the median across ad types; Sponsored Products spending equal to 27.7% of the sales those ads brought	July-October 2026	Eva, 150 brand accounts, October 26, 2026
A host-read podcast ad, per thousand downloads	$25 to $40	2026	Castos, September 2026

Prices are averages from benchmark firms' client accounts and vary widely by audience, season, and format.

Who Leads, as of November 1, 2026

The executives and officials who lead the companies and regulators the course covers, with their ages and time in office

The Sellers

Office	Officeholder	Age	In office since (years)	Role in the course
CEO, Alphabet and Google	Sundar Pichai	54	Google Oct 2015 (11); Alphabet Dec 2019 (6)	Runs the largest seller of advertising (Module 8)
Chairman and CEO, Meta	Mark Zuckerberg	42	Founded Facebook Feb 2004 (22)	Runs the second-largest seller (Module 8)
President and CEO, Amazon	Andy Jassy	58	Jul 5, 2021 (5)	Runs the largest retail media business (Module 9)
Chairman and CEO, Microsoft	Satya Nadella	59	Feb 4, 2014 (12)	Bing and LinkedIn (Module 8)
CEO, Apple	John Ternus	50 or 51	Sep 1, 2026 (new)	Succeeded Tim Cook; Apple's tracking rules reshaped mobile advertising (Module 14)
CEO, TikTok	Shou Zi Chew	43	2021 (5)	Leads TikTok, whose American arm became a joint venture in January 2026 (Module 8)
CEO, TikTok USDS Joint Venture	Adam Presser	Not published	Jan 2026 (new)	Runs the American joint venture (Module 8)
Co-CEOs, Netflix	Ted Sarandos; Greg Peters	62; 55 or 56	Jul 2020 (6); Jan 2023 (3)	Netflix's ad tier (Module 11)
President and CEO, Walmart	John Furner	51 or 52	Feb 1, 2026 (new)	Walmart's ad business and Vizio (Module 9)
Chairman and CEO, The Trade Desk	Jeff Green	49 or 50	Founded it in Oct 2009 (17)	The largest independent DSP (Module 10)

The Agency Groups

Office	Officeholder	Age	In office since (years)	Role in the course
President and CEO, Omnicom	John Wren	74	Jan 1, 1997 (29)	Leads the largest group by revenue since buying Interpublic (Module 7)
Chairman and CEO, Publicis Groupe	Arthur Sadoun	55	Jun 1, 2017 (9)	Leads the most valuable group (Module 7)
CEO, WPP	Cindy Rose	61	Sep 1, 2025 (1)	Leads WPP's turnaround plan (Module 7)
President and CEO, Dentsu	Takeshi Sano	56	Mar 27, 2026 (new)	Took over after Dentsu abandoned the sale of its international business (Module 7)
Chairman and CEO, Havas	Yannick Bolloré	46	2013 (13)	Leads the smallest of the big six groups (Module 7)

The Regulators

Office	Officeholder	Age	In office since (years)	Role in the course
Chairman, Federal Trade Commission	Andrew Ferguson	40	Jan 20, 2025 (1)	Leads America's advertising and competition regulator (Modules 13, 15, and 16)
Executive Vice-President for competition, European Commission	Teresa Ribera	57	Dec 1, 2024 (1)	Runs the EU's adtech case against Google (Module 16)
Commissioner for justice and consumer protection, European Commission	Michael McGrath	50	Dec 1, 2024 (1)	Leads the EU's proposed Digital Fairness Act (Module 15)

Ages and time in office are as of November 1, 2026, from official biographies, company filings, and news reports; years are completed years. Where only a birth year or an age at an earlier date is published, the age is given as a range.

Who's Who, as of November 1, 2026

The companies, bodies, laws, and people the course names; the leaders are on the page before this one

Sellers of Advertising

Name	What it is
Alphabet: Google Search, YouTube, and Google's ad network and ad tech	The largest seller of advertising, through search ads, video ads, and the tools that place ads on other sites (Modules 3, 8, 10, and 16)
Meta: Facebook, Instagram, WhatsApp, and Messenger	The second-largest seller, through ads in its apps' feeds, stories, and short videos (Modules 3, 8, 14, and 16)
Amazon Ads	Amazon's advertising arm: sponsored listings on its store, ads on Prime Video and Twitch, and a DSP that buys ads elsewhere (Modules 9, 10, and 16)
ByteDance; TikTok; Douyin; the TikTok USDS Joint Venture	The Chinese owner of TikTok and its Chinese twin Douyin, and the American company that has run TikTok's US app and data since January 2026 (Module 8)
Microsoft; Apple; Pinterest; Snap; Reddit; X	Smaller sellers: Bing and LinkedIn; Apple's App Store search ads; and four social platforms (Module 8)
Netflix; Disney; Amazon Prime Video; Peacock	Streaming services that sell ads in or around shows (Module 11)
Walmart Connect; Instacart; Uber; DoorDash; Target's Roundel	Retailers' and delivery companies' ad businesses (Module 9)
Alibaba; PDD; Tencent; Baidu; Kuaishou	China's largest sellers of advertising, through shopping apps, social apps, search, and short video (Module 8)
OpenAI; Perplexity	AI companies: OpenAI began selling ads in ChatGPT in 2026; Perplexity tried ads and dropped them (Module 17)

Agencies, Ad Tech, and Measurement

Name	What it is
Omnicom (with Interpublic); Publicis Groupe; WPP; Dentsu; Havas; Accenture Song; S4 Capital	The agency holding companies and their rivals, which plan, make, and buy advertising for clients (Module 7)
The Trade Desk; Prebid	The largest independent DSP, and the open-source software behind header bidding (Module 10)
Nielsen; VideoAmp; iSpot; Comscore	Measurers of television and video audiences; Nielsen's ratings remain the main currency (Modules 11 and 13)
The Media Rating Council (MRC)	An industry body that accredits audience measurements and sets standards such as viewability (Module 13)
WPP Media; WARC; Magna; eMarketer; MoffettNathanson; Media Dynamics; AdImpact	Forecasters and analysts of ad spending and prices (Modules 1, 8, 11, and 15)
Gartner; Ad Age; Interbrand; Kantar BrandZ; Circana; the Private Label Manufacturers Association (PLMA)	Sources for marketing budgets, the largest advertisers and agencies, brand values, and store brands' share (Modules 1, 4, 6, and 7)

Trade Bodies and Self-Regulators

Name	What it is
The Association of National Advertisers (ANA)	The main US body of advertisers, whose studies tracked agency pay, rebates, and where programmatic money goes (Modules 7 and 10)
The Interactive Advertising Bureau (IAB)	The digital ad industry's trade body, which counts US internet ad revenue with PwC and writes technical standards (Modules 1, 3, and 10)
The American Association of Advertising Agencies (4A's)	The trade body of US agencies, founded in 1917 (Modules 2 and 13)
The World Federation of Advertisers (WFA) and its Global Alliance for Responsible Media (GARM)	The global body of advertisers, and its brand-safety alliance of 2019-24 (Module 13)
The Institute of Practitioners in Advertising (IPA)	The British agencies' body, whose databank of campaigns underlies much practitioners' evidence (Module 5)
ISBA	The British advertisers' body, which commissioned studies of the programmatic supply chain (Module 10)
The National Advertising Division (NAD) of BBB National Programs	American advertising's self-regulatory court, which hears challenges to ad claims (Module 15)

Regulators, Courts, and Laws

Name	What it is
The Federal Trade Commission (FTC)	The US agency that polices deceptive advertising and enforces competition law (Modules 12 to 16)
The Food and Drug Administration (FDA); the Federal Election Commission (FEC); the Federal Communications Commission (FCC)	US regulators of prescription drug ads, of campaign ads' disclaimers, and of broadcasting (Module 15)
The Department of Justice (DOJ) Antitrust Division	Brought the two monopoly cases against Google (Module 16)
The European Commission	The EU's executive, which enforces competition law, the Digital Services Act (DSA), and the Digital Markets Act (DMA) (Modules 14 to 16)
The General Data Protection Regulation (GDPR)	The EU's privacy law of 2018, which governs the use of personal data for ads (Module 14)
The Children's Online Privacy Protection Act (COPPA)	The US law limiting the collection of data from children under 13 (Module 15)
Apple's App Tracking Transparency (ATT)	Apple's rule of 2021 requiring apps to ask before tracking users across other companies' apps (Module 14)
The Competition and Markets Authority (CMA); the Advertising Standards Authority (ASA)	Britain's competition regulator and its advertising self-regulator (Modules 12, 15, and 16)
Judges Amit Mehta, Leonie Brinkema, and James Boasberg	The federal judges who decided the search and ad tech cases against Google and the FTC's case against Meta (Module 16)

Researchers and Writers

Name	What they did
Claude Hopkins; Albert Lasker; David Ogilvy; Rosser Reeves; Bill Bernbach	Copywriters and agency heads who shaped advertising's methods (Module 2)
Edward Bernays; Vance Packard; John Kenneth Galbraith	The founder of public relations and two critics of advertising's power (Modules 2 and 4)
Byron Sharp and the Ehrenberg-Bass Institute; Les Binet and Peter Field	Researchers whose findings on reach, memory, and the balance of brand and short-term ads guide practitioners (Modules 4, 5, and 6)
Randall Lewis and Justin Rao; Thomas Blake, Chris Nosko, and Steven Tadelis; Bradley Shapiro, Günter Hitsch, and Anna Tuchman	Economists whose experiments and large studies measured advertising's effects (Module 5)
Bart Bronnenberg, Jean-Pierre Dubé, Matthew Gentzkow, and Jesse Shapiro	Economists who showed that informed shoppers buy store brands (Module 4)
Alessandro Acquisti; Garrett Johnson	Researchers who measured what tracking is worth to publishers and advertisers (Module 14)

Terms

The industry's vocabulary, defined once; each module defines its other terms where they first appear

Term	Meaning	Where it matters
Impression	One showing of an ad to one person or device	Throughout
CPM, CPC, and CPA	Cost per thousand impressions (from the Latin *mille*), cost per click, and cost per acquisition, meaning per sale, sign-up, or other action	Modules 6, 10, and 11
Click-through rate (CTR)	The share of impressions that get a click	Modules 3, 8, and 14
Return on ad spend (ROAS) and return on investment (ROI)	Revenue per dollar of advertising, against profit after the advertising's cost; a campaign can show a high ROAS and a negative ROI	Modules 5 and 13
Reach, frequency, and gross rating points (GRPs)	The share of a target audience that sees an ad at least once; the average number of times each of them sees it; and reach times frequency, in points of the audience	Modules 6 and 11
Share of voice (SOV) and excess share of voice (ESOV)	A brand's share of all advertising in its category, and how far that share exceeds its share of the market	Module 6
Brand building and performance (activation) advertising	Ads meant to build memories and future demand, against ads meant to prompt a purchase or click now	Modules 1, 5, and 6
Programmatic advertising and real-time bidding (RTB)	Buying and selling impressions by software; RTB auctions each impression in the fraction of a second a page or app takes to load	Modules 3 and 10
Demand-side platform (DSP) and supply-side platform (SSP)	Software that advertisers use to bid for impressions, and software that publishers use to offer them	Module 10
Ad exchange and ad server	The marketplace where impressions are auctioned, and the software that decides which ad a page shows and counts the result	Module 10
Header bidding	Code on a publisher's page that offers each impression to several exchanges at once before the ad server decides	Module 10
Made-for-advertising (MFA) site	A website built to show many ads to traffic it buys, with little content of its own	Module 10
First-party and third-party data	Data a company gathers from its own customers, against data gathered about people elsewhere and bought or shared	Module 14
Cookie and mobile advertising ID	A small file a browser keeps for a website, which in its third-party form lets ad firms follow a browser across sites; and the identifier a phone's operating system assigns for ad tracking	Modules 3 and 14
Attribution	Assigning credit for a sale among the ads a customer saw or clicked	Module 13
Incrementality, or lift	The sales an ad causes beyond those that would have happened anyway, measured by comparing people who saw it with a similar group who did not	Modules 5 and 13
Media mix model (MMM)	A statistical model that estimates each channel's contribution to sales from spending and sales over time	Modules 6 and 13
Viewability	Whether an ad could be seen: by the industry's standard, half its pixels on screen for one second, or two seconds for video	Module 13
Connected TV (CTV)	Television watched over an internet connection, including streaming services on smart TVs	Module 11
Upfront and scatter	Television ad time sold months ahead for a season, against time sold closer to airing	Module 11

Term	Meaning	Where it matters
Retail media	Advertising sold by retailers on their own sites and apps, and elsewhere using their shoppers' data	Module 9

Module 1 — What Marketing Is, and How Big It Is

The Four Ps, the Funnel, and a $1.3 Trillion Business

Recommended Reading / Reference

Book: Philip Kotler, Kevin Lane Keller and Alexander Chernev, *Marketing Management* (16th edition, 2021), the standard business-school text since Kotler's first edition of 1967, which sets out the marketing mix, segmentation and positioning at length

Report: WPP Media, *This Year Next Year: 2026 Global Midyear Forecast* (June 16, 2026), one of the standard counts of the world's advertising revenue by country and channel, best read with its December 2025 end-of-year edition

1a. Five Words for Selling

Five words describe the work of selling, and they are often confused; marketing is the widest of them. The American Marketing Association (AMA) defines it as “the activity, set of institutions, and processes for creating, communicating, delivering, and exchanging offerings that have value for customers, clients, partners, and society at large” (approved 2017). In practice it covers what a company makes, what it charges, where it sells, and how it tells people about it.

Advertising is one part of that last task: paid messages, identified as coming from the advertiser, in media the advertiser does not own, such as a television spot, a search listing, a billboard, or a sponsored post. Promotion, in the narrow sense the trade uses, means short-term incentives to buy: coupons, temporary price cuts, free samples, and payments to retailers for displays and shelf space. Publicity is coverage the company does not pay for, such as a news story or a review, and winning it is the work of public relations, whose origins Module 2 traces. Selling is the person-to-person part, done by sales forces, store staff, and call centers.

Most coupons go unused: Inmar Intelligence counted 67 billion coupons distributed in the United States in 2024 and 871 million redeemed in stores, a redemption rate of 1.3% (Coupons in the News, February 24, 2025). The management writer Peter Drucker drew the line between marketing and selling in 1973: “The aim of marketing is to make selling superfluous.” A product that fits what buyers want, at a price they accept, in places where they shop, needs less pushing.

1b. The Four Ps, and Why Price and Place Can Cost More Than Advertising

The standard map of a marketer's choices is the marketing mix, an idea with Harvard origins: in 1948 the professor James Culliton described marketers as “mixers of ingredients”. His colleague Neil Borden used the phrase “marketing mix” in his teaching from the late 1940s and in his presidential address to the American Marketing Association in 1953, and published “The Concept of the Marketing Mix” in 1964. In *Basic Marketing: A Managerial Approach* (1960), E. Jerome McCarthy grouped the ingredients into four Ps: product, price, place (where and how the product is sold), and promotion (how buyers are told about it).

Philip Kotler's textbooks made the four Ps the field's standard, and others have added to them. Bernard Booms and Mary Jo Bitner proposed three more for service businesses in 1981: people, process, and physical evidence. In 1990 Robert Lauterborn argued in *Advertising Age* for four Cs seen from the buyer's side: the consumer, cost, convenience, and communication.

Promotion, the P that contains advertising, is often not the largest cost. For makers of packaged groceries and household goods, the biggest marketing bill is usually trade promotion, meaning payments to retailers in the form of temporary price cuts, fees for displays and shelf space, and allowances that fund the store's own discounts. Strategy&, the consulting arm of PwC, put trade promotion at about 20% of gross sales in 2017. That made it the second-largest cost in a packaged-goods maker's accounts after the goods themselves, and more than $200 billion a year in the United States.

The most recent published split between trade promotion and advertising that could be found is from 2004, when Cannondale Associates reported that packaged-goods marketing budgets went 48% to trade promotion and 25% to advertising. McKinsey still puts trade promotion at about 20% of revenue, and finds that 72% of US promotions do not break even (as quoted by TELUS, November 24, 2025). Accounting rules treat most payments to retailers as a reduction of sales rather than as a marketing cost, so they rarely appear in counts of advertising, including those later in this module.

1c. Segmentation, Targeting, Positioning

Before choosing a mix, a company chooses whom to serve. In 1956 the marketing scholar Wendell Smith set out the idea of market segmentation in the *Journal of Marketing*: dividing a market into groups of buyers who want different things, instead of selling one product to everyone. Segments can be drawn by demographics such as age and income, by geography, by attitudes and ways of life, or by behavior, such as how often and how much people buy.

Targeting is the choice of which segments to pursue: a maker of luxury cars and a discount grocer both segment the market, then aim at opposite ends of it. Online advertising made targeting cheaper and finer, down to a single browser or phone; how that works, and what it costs in privacy, are the subjects of Modules 10 and 14.

Positioning is the place a brand tries to hold in buyers' minds, relative to its rivals. Al Ries and Jack Trout popularized the term in articles in *Advertising Age* in 1972 and in their book *Positioning: The Battle for Your Mind* (1981). A well-known example is Avis, the car-rental company, which from 1962 turned its second place in the market into a promise: “We try harder.” A position is a claim about the product, so it works only if the product can bear it.

1d. The Funnel, From 1898 to Today

The oldest model of how advertising works is a sequence of steps. In 1925 the psychologist Edward Strong Jr. credited the advertising man Elias St. Elmo Lewis with a slogan of 1898: an advertisement should “attract attention, maintain interest, create desire.” Lewis later added “get action”, and the four steps became known as AIDA (Attention, Interest, Desire, Action). In 1961 Robert Lavidge and Gary Steiner proposed a longer ladder in the *Journal of Marketing*: awareness, knowledge, liking, preference, conviction, and purchase.

These models gave the industry its funnel, a planning tool more than a description: many people become aware of a brand, fewer consider it, fewer still buy, and some buy again. In 2009 consultants at McKinsey argued that buyers loop rather than fall, adding and dropping brands as they research, then judging the purchase afterward and sometimes recommending it. Online, the stages can collapse into seconds, from an ad in a feed to a purchase in the same app.

The funnel also divides the industry's work. Brand-building advertising works at the top, aiming to make a brand come to mind easily and seem worth more, among many people, over months and years.

Performance advertising, also called activation, works at the bottom: it seeks an action now, such as a click, a sign-up, or a sale, from people already close to buying, and is judged by the results it can count. The evidence on how to balance the two is Module 5's subject, and how companies split their budgets between them is Module 6's.

Marketers also sort their channels by who controls them. Paid media is space bought from someone else, which is to say advertising, while owned media belongs to the company: its website, app, stores, packaging, e-mail lists, and social-media accounts. Earned media is attention that others give for free, such as press coverage, reviews, shares, and word of mouth. The three feed one another, as when an ad wins press coverage or a company uses its own customer list to aim its paid ads.

Marketing also divides by the buyer. Consumer marketing, known as B2C (business to consumer), sells to households, usually in large numbers at low prices, through mass media and stores. Business-to-business (B2B) marketing sells to other companies, often to a committee of buyers over a long sales cycle, in which a sales force does much of the work and advertising mainly opens doors; the market totals that follow include both.

1e. How Big Is Advertising? Four Counts

No one counts the world's advertising directly: forecasting firms estimate it from companies' reported revenue, trade associations' data, and their own surveys, and they publish new estimates once or twice a year. The four most quoted, set out in the table, differ by roughly $300 billion for 2026.

Forecaster	Editions used	2025	2026
WPP Media	December 10, 2025; June 16, 2026	$1.14 trillion (+8.8%)	About $1.3 trillion (+8.9%)
WARC Media	October 8, 2026	Growth of 10.0%	$1.34 trillion (+11.9%), with $1.46 trillion forecast for 2027
Dentsu	December 3, 2025; May 27, 2026	Growth of 5.8%	Over $1 trillion for the first time (+5.0%)
Magna	June 16, 2025	$979 billion (+4.9%)	Just over $1 trillion (+6.3%)

WPP Media's releases; WARC Media via MediaNews4U (October 8, 2026); Dentsu's December release via the ANA, and its May update, which cut 2026 growth from 5.1%, via Communicate Online; Magna's release.

Most of the gap comes from what each firm counts. Magna, the forecasting unit of the Interpublic Group, estimates the net advertising revenue of media owners in 70 markets from their financial reports and local trade data, and it shows growth with and without cyclical events such as elections and the Olympics. WPP Media also counts sellers' ad revenue, leaves out US political advertising, and in December 2025 added travel and financial-services companies' ad sales to its commerce category (Marketing Week). WARC and Dentsu describe their figures as advertising spend, and exchange rates and revisions move every total between editions.

No global forecast from Magna later than June 2025 could be found as of October 31, 2026; Omnicom has since bought Interpublic (Module 7). However it is counted, advertising is growing faster than the world economy: WPP Media says that in 2026 advertising will take its highest share of world GDP since 1999, above its peak of 2000, at the height of the dot-com boom (June 2026).

1f. Where the Money Is: Countries and Channels

The United States is by far the largest market: WPP Media put its advertising revenue at $431.2 billion in 2025, up 12.5%, about 38% of the world total. China came second, at $216 billion, up 6.8%, followed by the United Kingdom at $58.4 billion and Japan at $51.9 billion; India's market was $20.7 billion (WPP Media, December 2025, as reported by IndianTelevision.com). In June 2026 WPP Media forecast US growth in 2026 of 11.9% excluding political advertising, or 13.9% with the spending on that year's midterm elections.

Spending per person shows how uneven the market is: WARC Media estimates 2026 advertising at $1,395 per person in the United States, $935 in the United Kingdom, $170 in China, and $13 in India (October 8, 2026).

By channel, most of the money now goes to platforms that sell by auction. WARC Media's October 2026 forecast puts social media at $394.6 billion in 2026, up 21.3%, and search at $295.7 billion, up 14.2%. With retail media, the ads that stores and marketplaces sell (Module 9), these channels take 66.4% of all advertising, a share WARC expects to reach 70.0% by 2028. Television, once the largest medium, now ranks fourth among WARC's channels (Module 11).

The standard measure of US online advertising is the annual survey that the IAB, the online advertising industry's trade association, runs with the accounting firm PwC. It put US internet advertising revenue at $294.6 billion in 2025, up 13.9% (April 16, 2026). The ten biggest sellers took 84.1% of it, up from 80.8% in 2024, and $162.4 billion, up 20.5%, was bought programmatically, by automated auction (Module 10); the table breaks the total down by format.

Format	2025 revenue	Growth	What it includes
Social media	$117.7 billion	+32.6%	Ads in social apps' feeds, stories, and short videos
Search	$114.2 billion	+11.0%	Paid listings on search results pages; growth slowed from 15.9% in 2024
Display	$81.6 billion	+9.8%	Banners and other graphic ads on websites and in apps
Digital video	$78.0 billion	+25.4%	Video ads, including connected TV, social video, and short-form video
Commerce media	$63.4 billion	+18.0%	Ads sold by retailers and other commerce companies, on their own sites and beyond (Module 9)

IAB/PwC, *Internet Advertising Revenue Report: Full Year 2025* (April 16, 2026). The categories overlap, since social video counts as both social and video, so the rows sum to more than the total. Podcast advertising is in Module 11.

1g. What Companies Spend on Marketing

For a large company, marketing typically costs a little under a tenth of revenue. Gartner's 2026 CMO Spend Survey was fielded from January to March 2026 among 401 marketing leaders in North America, the United Kingdom, and Europe, most at companies with more than $1 billion of revenue. It found budgets averaging 7.8% of revenue, against 7.7% in 2025, a change Gartner called effectively flat (May 11, 2026); the worked example shows how such a budget divides.

Makers of branded consumer goods spend more: L'Oréal reported €14.18 billion of advertising and promotion in 2025, 32.2% of its sales of €44.05 billion. Unilever's “brand and marketing investment” was €8.14 billion, 16.1% of its turnover of €50.50 billion. Each company defines its own measure, so the two shares are not strictly comparable, but both are two to four times the average across industries.

Worked example — Where $100 of a Marketing Budget Goes

Inputs: in Gartner's 2026 survey, paid media takes 31.4% of the average marketing budget, labor (the marketing staff) 24.5%, and marketing technology 19.4% (“martech”: the software and data services marketers use). The published summaries of the 2026 survey give no share for agencies; in 2025 agencies took 21% (Marketing Brew, May 20, 2025), used here as an assumption. Step 1, the named items: $31.40 + $24.50 + $19.40 = $75.30 of every $100. Step 2, the rest: $100 − $75.30 = $24.70 for agencies and everything else. Step 3, at the 2025 agency share, about $21.00 goes to agencies and about $3.70 to other costs. Step 4, a company with $1 billion of revenue spending the average share: $1 billion × 0.078 = $78 million, of which about $24.5 million buys media, $19.1 million pays staff, $15.1 million buys software and data, and about $16.4 million goes to agencies. Step 5, per $100 of the company's sales: $7.80 × 0.314 = about $2.45 buys advertising space and time. Result: the media spending that the forecasts in 1e add up is only about a third of what marketers spend. The rest pays for the people, tools, and agencies that decide what to buy, and Gartner's figures show the balance shifting: paid media's share has risen from 25.1% in 2021, while martech's has fallen from 26.6%.

Module 1 Key Points

What Marketing Is, and How Big It Is

• Advertising is one part of marketing, which also decides the product, its price, and where it is sold: McCarthy's four Ps (1960).

• Packaged-goods makers spend about 20% of gross sales on trade promotion, paying retailers, which can exceed what they spend on advertising.

• Brand building works on many people over years; performance advertising seeks an action now from people already close to buying.

• Forecasts of 2026 world ad revenue run from about $1 trillion to $1.34 trillion, because the forecasters count different things.

• The United States is the largest market, with $431.2 billion of advertising revenue in 2025, about twice China's.

• Large companies spend 7.8% of revenue on marketing, but only about a third of that budget buys advertising.

Module 1 Quiz

Check Your Understanding

1. Who grouped the marketing mix into the four Ps, and in what year?

A) E. Jerome McCarthy, in 1960
B) Neil Borden, in his 1953 address to the American Marketing Association
C) Elias St. Elmo Lewis, in 1898
D) Philip Kotler, in the first edition of his textbook in 1967, which added people and process

2. Why can price and place cost a packaged-goods maker more than advertising does?

A) Advertising is recorded as a reduction of sales, so it looks smaller than it is
B) Payments to retailers for price cuts, displays, and shelf space run to about 20% of gross sales
C) Delivering goods to stores is counted as a marketing cost.
D) Retailers take a 15% commission on every product they stock.

3. What does WPP Media's headline count of world advertising leave out?

A) China's ad market
B) Ads sold by retailers and marketplaces
C) US political advertising
D) Television advertising sold by streaming services

4. On WARC Media's October 2026 forecast, which three channels take about two-thirds of world ad spending?

A) Television, search, and social media
B) Search, online video, and out-of-home advertising, led by digital billboards
C) Social media, television, and print
D) Social media, search, and retail media

5. In Gartner's 2026 survey, roughly what share of the average marketing budget buys paid media?

A) About two-thirds, with staff and software sharing the rest
B) About 8%, the same as marketing's share of revenue
C) About a third
D) About half, after agencies' fees are paid

Module 2 — Selling to the Masses, 1841–1994

From Patent Medicines to the Thirty-Second Spot

Recommended Reading / Reference

Book: Stephen Fox, *The Mirror Makers: A History of American Advertising and Its Creators* (1984), a narrative history of the agencies and the people who ran them, from Volney Palmer to the creative revolution of the 1960s and after

Reference: Duke University's Rubenstein Library, *Emergence of Advertising in America, 1850–1920*, a free online collection of early advertisements and agency records, with a timeline of the period

2a. Selling Space: The First Agents

American advertising began as a trade in newspaper space, and in 1841 Volney Palmer opened what is usually counted as the first American advertising agency, in Philadelphia. He found advertisers for newspapers and took a commission from the papers of up to 25%, according to a 2022 working paper on the history of agency pay by the marketing scholars Dan Horsky and Robert Zeithammer. Other agents worked as space brokers, buying columns wholesale and reselling them at whatever markup an advertiser would bear.

Francis Wayland Ayer changed the terms. In 1875 his Philadelphia agency, opened in 1869 and named for his father, offered advertisers an “open contract”: the agency would act for the advertiser rather than the newspaper, buy space at the rates the papers actually charged, and add a commission on that price, eventually 15% (Advertising Hall of Fame). The agency had become the buyer's agent, although the newspapers still paid its commission.

The 15% then hardened into custom, stabilizing by the early 1890s in Horsky and Zeithammer's account. In 1893 the newspapers' trade association resolved that commissions would be paid “only to recognized advertising agents and not to advertisers”, so an advertiser could not keep the commission by buying space directly. The American Association of Advertising Agencies, founded in 1917, helped hold the rate. A federal antitrust inquiry into agency recognition, begun in 1924, was dismissed, and further antitrust actions followed in the 1950s. How agencies moved from the commission to fees is Module 7's subject.

2b. Patent Medicines and National Brands

The biggest early advertisers sold patent medicines: tonics, bitters, and compounds that promised to cure almost anything and often contained alcohol or opiates. In October 1905 the journalist Samuel Hopkins Adams opened his series “The Great American Fraud” in Collier's Weekly: “Gullible America will spend this year some seventy-five millions of dollars in the purchase of patent medicines.” Of one famous remedy he wrote, “Lydia Pinkham's variety of drink depends for its popularity chiefly on its alcohol.”

Adams also showed how the makers kept newspapers on their side: the proprietor of Hall's Catarrh Cure printed in red on every advertising contract a clause that voided it if the newspaper's state passed a law hostile to patent medicines, and many large firms copied this “red clause”. Any paper that carried the ads had a reason to oppose regulation.

Congress passed the Pure Food and Drug Act in 1906, requiring labels to declare dangerous ingredients such as alcohol, morphine, and opium, and banning false labeling. It governed the label, not the advertisement; false advertising was left to a separate agency and a later law (2h).

Patent-medicine makers were among the first to sell one branded product across the whole country, and makers of soap, cereal, and biscuits followed. Railroads, national magazines, and factories that could supply a continent made it possible to sell the same packaged product everywhere under one name. A brand let the maker go over the retailer's head: if shoppers asked for it by name, the grocer had to stock it. An early example was the National Biscuit Company's Uneeda biscuit, sold in a sealed package instead of loose from the grocer's barrel and advertised nationally by Ayer; Module 4 explains what such brands are worth.

2c. “Salesmanship in Print”

By 1900 agencies wrote the advertisements as well as placing them, and two schools of copy emerged. The first came from Lord & Thomas of Chicago, run by Albert Lasker. In 1904 a former Canadian mounted policeman named John E. Kennedy told Lasker that advertising was “salesmanship in print”, and Lasker put him in charge of copy. Ads in this school gave reasons to buy, at length, as a good salesman would.

Claude Hopkins, whom Lasker hired in 1908, turned the idea into a method. He put coupons in ads, coded so that each newspaper or each version of an ad could be credited with the replies it drew, and he gave away samples to gauge demand. In *Scientific Advertising* (1923) he wrote that “almost any question can be answered, cheaply, quickly and finally, by a test campaign.” His coded coupons are the ancestors of today's online experiments (Module 5).

The second school came from a New York agency, J. Walter Thompson (JWT). Stanley Resor bought it in 1916 and ran it with Helen Lansdowne, who had joined in 1908 as its first woman copywriter and married him in 1917, and JWT became the first agency to open a research department. Lansdowne's 1911 advertisement for Woodbury soap, “A skin you love to touch”, sold a feeling rather than a list of reasons, and Woodbury used the slogan until the 1940s. JWT also hired, in 1920, the founder of behaviorist psychology, John B. Watson, after Johns Hopkins University forced him out, and it led American agencies in billings from 1922 to 1972.

2d. Bernays and Public Relations

Public relations grew up beside advertising, and its best-known early theorist was Edward Bernays, a nephew of Sigmund Freud, who set out its methods in *Crystallizing Public Opinion* (1923) and *Propaganda* (1928). The second book opens: “The conscious and intelligent manipulation of the organized habits and opinions of the masses is an important element in democratic society.” His technique was to stage events that the press would report as news, so that a client's message reached the public as publicity rather than as an advertisement.

His best-known event took place on March 31, 1929, when, working for George Washington Hill of the American Tobacco Company, he arranged for young women to light cigarettes as they walked in New York's Easter Parade, presented as “torches of freedom”, to break the taboo on women smoking in public. Historians have since argued that its effect was overstated (History Today, February 2019). Its method, a paid message presented as news, returns with influencer marketing (Module 12).

2e. Radio, the Soap Opera, and the Brand Manager

Radio made the advertiser a producer. On August 28, 1922, WEAF, the American Telephone and Telegraph Company's station in New York, aired what is usually counted as the first radio commercial: a ten-minute talk promoting the Queensboro Corporation's Hawthorne Court apartments in Jackson Heights, for which the company paid $50 (Saturday Evening Post, August 2024). AT&T called its plan toll broadcasting: like a toll telephone line, the station was open to anyone who paid.

Through the 1920s and 1930s the advertiser or its agency usually produced a whole program and put its name on it. Daytime serials written for housewives gave one form its name: *Painted Dreams*, written by Irna Phillips, began in October 1930 on WGN, a Chicago station. Procter & Gamble began sponsoring serials in 1932, soon including the long-running *Ma Perkins*, and invested more in the form than any other company; the serials became known as soap operas (The Irish Times, June 7, 2026). Phillips's *Guiding Light* (1937) later moved to television and ran until 2009.

Procter & Gamble also invented the job that still organizes consumer-goods marketing. On May 13, 1931, Neil McElroy, a young advertising manager, wrote a three-page memo arguing that each brand, starting with Camay soap, should have one manager responsible for its sales, advertising, and research, even when that meant competing with the company's own Ivory. P&G adopted the brand-management system, and its rivals copied it. McElroy later ran the company and served as President Eisenhower's secretary of defense.

2f. Television: From Sponsors to Spots

Television began with radio's model. The first television commercial, on Guinness World Records' account, ran on July 1, 1941, on WNBT, the New York station of the National Broadcasting Company (NBC), before a Dodgers-Phillies baseball game: a ten-second spot for Bulova watches that cost $9. After the war, single sponsors paid for whole shows and put their names on them, as in the *Colgate Comedy Hour* and the *Camel News Caravan*. A one-hour show cost its sponsor about $35,000 in 1952 and about $90,000 by the end of the decade (Campbell and others, *Media & Culture*).

Rising costs pushed the networks to sell time instead. Sylvester “Pat” Weaver, NBC's head of television and from 1953 its president, launched *Today* (1952) and *Tonight* (1954) as programs the network controlled, selling 30- and 60-second spots in them to many advertisers. A scandal then finished the sponsor era. Quiz shows had been the hits of the mid-1950s, among them *The $64,000 Question* on the Columbia Broadcasting System (CBS), sponsored by Revlon, and NBC's *Twenty-One*, sponsored by Geritol, on which Charles Van Doren won $129,000 in 1957.

In 1958, after accusations that the show *Dotto* was fixed, the networks dropped twenty quiz shows, and in 1959 Van Doren admitted to a congressional committee that he had been given answers. The networks took control of programming and sold commercial time to many advertisers, and in the 1960s the spot became the norm on network television. The thirty-second spot became the standard unit of the network era, and it is still the unit in which television's prices are quoted (Module 11).

2g. Hard Sell, Brand Image, and the Creative Revolution

Rosser Reeves of the Ted Bates agency stood for the hard sell. Every ad, he argued, must make one claim that rivals did not or could not make, and repeat it: the unique selling proposition (USP), set out in his *Reality in Advertising* (1961). In 1952 he made short television spots for Dwight Eisenhower's presidential campaign, in which the candidate answered a citizen's question on Korea, corruption, or the cost of living;

they ran in the campaign's last three weeks (PBS). The spots brought the methods of the packaged-goods account to a presidential race, and Module 15 covers political advertising's rules.

David Ogilvy, a Briton who opened his New York agency in 1948, joined research to brand image: in his view every advertisement was part of a long-term investment in a brand's personality. He set out his rules for copy, headlines, and dealing with clients in *Confessions of an Advertising Man* (1963).

Doyle Dane Bernbach (DDB), founded in 1949, led what the trade called the creative revolution of the 1960s, in which copywriters and art directors worked in pairs and wit took the place of the hard sell. Its 1959 campaign for the Volkswagen Beetle, “Think Small”, written by Julian Koenig and designed by Helmut Krone, sold a small German car to Americans by admitting its smallness. *Advertising Age* named it the best campaign of the twentieth century in 1999.

2h. The Critics and the Regulators

The postwar boom in advertising drew critics, among them Vance Packard, whose *The Hidden Persuaders* (1957) described agencies hiring psychologists to probe buyers' unconscious motives, a practice known as motivation research. In September 1957 the market researcher James Vicary announced that flashing “Eat Popcorn” and “Drink Coca-Cola” on a cinema screen for 1/3,000 of a second had raised sales of popcorn by 57.5% and of Coca-Cola by 18.1% among 45,699 filmgoers in Fort Lee, New Jersey. He failed to repeat the result, and in 1962 he told Advertising Age that the study had been “a gimmick”, with too little data to be meaningful. Module 4 explains what is known about ads that work below awareness.

The economist John Kenneth Galbraith made the broader case in *The Affluent Society* (1958). In a rich economy, he argued, production creates the wants it then satisfies, through advertising and salesmanship, a process he called the dependence effect. If so, the fact that people buy what is advertised shows little about whether it makes them better off; economists' replies are set out in Module 4.

Federal regulation came slowly: the Federal Trade Commission (FTC), created in 1914, was set up to police unfair methods of competition, and in 1931 the Supreme Court held in *FTC v. Raladam* that it had to show harm to competitors, not merely to buyers, to stop a false advertisement. The Wheeler-Lea Act of March 21, 1938 removed that limit: it added “unfair or deceptive acts or practices” to the FTC's remit, gave it power over false advertising of foods, drugs, and cosmetics, and set fines of up to $5,000 for breaking its orders.

The sharpest restriction fell on cigarettes: under the Public Health Cigarette Smoking Act, signed in 1970, cigarette advertising left American television and radio on January 2, 1971, and the tobacco companies moved their budgets to print, billboards, and sports sponsorship. Module 15 describes the law as it stands, including the later limits on tobacco.

2i. Lists, Holding Companies, and Cable

A different tradition sold without mass media: mail-order catalogs, starting with Montgomery Ward's in 1872, reached buyers far from stores, and their makers learned to judge each mailing by the orders it brought in. Lester Wunderman, whose agency was founded in 1958, named and defined “direct marketing” in a 1967 address at the Massachusetts Institute of Technology: selling to individuals through mail and telephone, and measuring every response. The ZIP (Zone Improvement Plan) code, introduced in 1963, and computers let marketers sort households and rent lists of them.

By the 1980s companies were building databases of their own customers, as airlines did with frequent-flyer programs such as American Airlines' AAdvantage (1981). The online advertising of Module 3 inherited this measured, list-based tradition as much as it inherited the mass media's.

Agencies, for their part, grew into groups, because one agency could not serve two rival advertisers: each feared its plans would reach the other. In 1961 Marion Harper of McCann-Erickson solved the problem by forming Interpublic, renamed the Interpublic Group of Companies in 1964: a holding company owning separate agencies that could serve competing clients (MediaVillage). The structure spread in the 1980s, as the table shows, when WPP, a British maker of wire baskets that Martin Sorrell turned into an agency group, took over two of the best-known American agencies.

Year	Deal	What it made
1986	BBDO (Batten, Barton, Durstine & Osborn), Doyle Dane Bernbach, and Needham Harper merge	Omnicom, a deal *Advertising Age* called “the Big Bang”
1986	Saatchi & Saatchi buys Ted Bates	Rosser Reeves's old agency passes to a British group
1987	WPP buys the JWT Group for $566 million, one of the trade's first hostile takeovers	J. Walter Thompson, whose roots go back to 1864, under a British holding company
1989	WPP buys the Ogilvy Group for $864 million	WPP owns both J. Walter Thompson and Ogilvy & Mather

MediaVillage, “The Birth of the Big 4 Holding Companies”; Campaign's history of agency mergers. Module 7 follows the groups to 2026, including Omnicom's purchase of Interpublic.

Television's audience, meanwhile, began to split: for much of the 1950s, 1960s, and 1970s the three networks, the American Broadcasting Company (ABC), CBS, and NBC, drew more than 90% of American viewers, but by the 1990–91 season their share had fallen to 61%. Cable channels such as the Entertainment and Sports Programming Network (ESPN, 1979), the Cable News Network (CNN, 1980), and Music Television (MTV, 1981) offered narrower audiences that advertisers could pick out, such as sports fans or teenagers. Fox began in 1986 as a fourth broadcast network. Videocassette recorders (VCRs), in 36% of households in 1986 and 70% by 1989, let viewers skip the commercials in recorded shows (EBSCO Research Starters).

By 1994 the mass audience that the thirty-second spot had been built to reach was fragmenting, and advertisers were already buying smaller, better-defined audiences. Module 3 begins in October of that year, with the first advertisement sold on the World Wide Web.

Module 2 Key Points

Selling to the Masses, 1841–1994

• Ayer's open contract of 1875 made the agency the advertiser's agent; its commission settled at 15% by the early 1890s.

• Patent-medicine makers were the first great advertisers; the 1906 food and drug law policed their labels, not their ads.

• Hopkins's coded coupons and test campaigns, set out in *Scientific Advertising* (1923), are the ancestors of online experiments.

• Sponsors produced whole radio and TV shows until the quiz-show scandals of 1958–59; networks then sold spots to many advertisers.

• The FTC gained clear power over false advertising only with the Wheeler-Lea Act of 1938.

• Holding companies, starting with Interpublic in 1961, let one owner serve rival clients through separate agencies.

Module 2 Quiz

Check Your Understanding

1. What did N. W. Ayer & Son's “open contract” of 1875 change?

A) Newspapers agreed to stop paying commissions to any agency
B) The agency acted for the advertiser and added a commission to the price the papers actually charged
C) Advertisers could cancel their contracts if a state restricted patent medicines
D) Agencies stopped taking commissions from the papers and charged advertisers a flat monthly fee instead

2. What did the Pure Food and Drug Act of 1906 regulate?

A) Newspaper and magazine advertisements for medicines
B) The prices that patent-medicine makers could charge
C) The alcohol content of medicines, which it capped at 10%
D) The labels of foods and drugs, which had to declare ingredients such as alcohol and opiates

3. What did Neil McElroy's memo of May 13, 1931 propose?

A) That each brand have its own manager, even against P&G's other brands
B) That P&G sponsor daytime radio serials
C) That Camay be merged into Ivory so that one campaign could sell both soaps
D) That P&G pay its agencies a flat fee instead of the customary 15% commission

4. What ended the era of single-sponsor network television programs?

A) The quiz-show scandals of 1958–59, after which the networks controlled programs and sold spots
B) The 1971 ban on cigarette advertising, which removed the largest sponsors
C) The arrival of cable channels such as CNN and MTV
D) An FTC rule of the 1950s limiting each sponsor to one program

5. What did James Vicary admit in 1962 about his subliminal-advertising claim of 1957?

A) That the effect was real but much smaller than he first reported
B) That the messages had been flashed for a full second, not 1/3,000 of one
C) That the cinema had cut the price of popcorn during the test, which explained the extra sales
D) That it was a gimmick, with too little data to be meaningful

Module 3 — The Digital Turn, 1994–2026

From the Banner Ad to the Feed, the Auction, and the Store

Recommended Reading / Reference

Book: Tim Wu, *The Attention Merchants: The Epic Scramble to Get Inside Our Heads* (2016), a Columbia law professor's history of the businesses that sell attention to advertisers, from the penny press to Google and Facebook

Report: Google Inc., Form S-1 registration statement (April 29, 2004), filed for its listing on the stock market, with the founders' letter and the first public figures on the search-advertising business; free on the Securities and Exchange Commission's website

3a. A Cookie and a Banner, 1994

Two inventions of 1994 made the web an advertising medium. The first was the cookie, a small file that a website stores in a visitor's browser so that it can recognize the browser on a later visit. Lou Montulli, a programmer at Netscape, devised it that year so that sites could remember such things as the contents of a shopping cart. Advertisers soon used cookies to recognize the same browser on many sites, and that became the basis for counting and aiming online ads; Module 14 explains how cookies are now being limited.

The second was the paid banner: on October 27, 1994, HotWired, the online offshoot of Wired magazine, launched carrying ads from about a dozen advertisers. AT&T's, which read “Have you ever clicked your mouse right here? You will.”, is usually counted as the first (Guinness World Records). About 44% of the people who saw it clicked on it, by the figure usually quoted (The Atlantic, April 2017).

HotWired sold its space the way magazines did: “We charged $10,000 a month for advertisers to have their ads run on a section of the website,” one of its early staff later told Digiday. With no cookies yet in use, the site counted page views in its server logs and e-mailed the totals to its advertisers.

3b. Selling by the Thousand

The pricing that spread was the one magazines and television already used, a price per thousand impressions, known as the cost per thousand (CPM). DoubleClick, founded in 1996, sold and served ads on this basis across many websites. Because it delivered the ads from its own computers, a single DoubleClick cookie could recognize a browser on every site in its network, so advertisers could limit how often one person saw an ad, aim ads by browsing history, and count results across many publishers. A cookie set by a company other than the site being visited became known as a third-party cookie, and it carried online targeting for two decades.

The first boom ended in the dot-com crash of 2000–01, when many of the start-ups that had bought banners failed. DoubleClick survived as the leading ad server for publishers, the software that decides which ad fills each space on a site. Google bought it for $3.1 billion, announcing the deal in April 2007 and closing it in March 2008 after American and European regulators cleared it. That purchase was the base of the publisher ad-tech business that a federal court later found Google had monopolized (Module 16).

3c. Paying by the Click: The Search Auction

The model that made online advertising rich came from search. In February 1998 GoTo.com, a start-up from Bill Gross's Idealab, let advertisers bid for the top places in its search results and charged the bid each time someone clicked; by July 1998 bids reached a dollar a click. An advertiser paid only when a searcher showed interest, and the search itself revealed what the searcher wanted at that moment. GoTo renamed itself Overture in 2001, and Yahoo bought it in 2003 for $1.63 billion.

Google launched AdWords in October 2000, at first selling text ads beside its results by the impression. In 2002 it moved to pricing by the click (CPC) and ranked ads not by bid alone but by bid combined with how relevant users found them, judged mainly by how often they clicked. The design rewarded ads that people wanted to click, and it also raised Google's expected revenue from each search, since a click is what Google is paid for. In 2003 Google's AdSense extended the system to other publishers' pages, placing ads that matched each page's content and sharing the revenue.

Search ads combined two things no earlier medium had joined so closely: each answered a question the buyer had just asked, and its results could be counted click by click. Module 8 describes Google's advertising business today, and Module 10 explains how such auctions set prices.

3d. The Feed

Social networks sold something search could not: who people were, as well as what they did. On November 6, 2007, Facebook launched an advertising system that included Beacon, which took purchases and other activity from 44 partner websites and published them in the News Feeds of users' friends. A protest group organized by MoveOn drew 50,000 members in under ten days, and Coca-Cola withdrew. On December 5 Mark Zuckerberg apologized and made the system opt-in; a class action ended in a $9.5 million settlement, and on September 21, 2009 Facebook announced that it would shut Beacon down.

Meanwhile, phones moved the audience. Apple's iPhone went on sale in June 2007 and its App Store opened in July 2008, and people's time online began shifting from web pages, where cookies worked, to apps, where they did not. Ad networks followed the audience: Google bought the mobile ad network AdMob for $750 million, a deal it announced in November 2009.

The format that lasted was the ad placed in the feed itself, among friends' posts, and sold by auction, and it mattered most on phones. Facebook's prospectus of February 2012 warned that it did not “directly generate any meaningful revenue” from its mobile products; that year it began placing ads in its mobile News Feed, and it listed its shares in May. By the fourth quarter, mobile brought in about 23% of its advertising revenue, up from about 14% in the third, and more people used Facebook each day on phones than on computers (Facebook, January 30, 2013). Across the whole US market, phones and tablets took 51% of internet advertising revenue in 2016 (IAB/PwC, April 26, 2017).

The feed changed what advertisers bought: instead of a place on a page, they bought a kind of person, defined by age, place, interests, and behavior, and increasingly by the platform's own prediction of who would act; the platform then found such people wherever they were scrolling (Module 8 describes how the platforms sell today).

3e. The Plumbing Becomes an Auction

Behind display advertising, selling moved from salespeople to software, in the steps the table lists. Exchanges began to auction each single view of an ad while the page loaded; publishers then let many exchanges bid at once rather than one after another; and the auctions moved from charging the winner the runner-up's price to charging its own bid, which is called a first-price auction. Module 10 explains how the machinery works and where the money goes.

Year	Milestone	What it changed
About 2009–10	Ad exchanges begin real-time bidding (RTB)	Each ad view can be priced separately, for a particular browser or device
2012	Apple gives iPhones a resettable advertising identifier (IDFA); Google later does the same on Android	Apps gain a shared tracking identifier, the phone's counterpart of the cookie
2014–15	Header bidding spreads	Rival exchanges bid for each impression on more equal terms with Google's
2019	Google completes its move to first-price auctions in its Ad Manager (September)	Buyers must decide for themselves how far below their value to bid

IAB Tech Lab (OpenRTB); PPC Land on header bidding (a secondary source); Google, September 5, 2019. Device identifiers and their limits are Module 14's subject.

3f. Where the Money Went

The money followed the users: US internet advertising grew from $7.3 billion in 2003 to $72.5 billion in 2016, about 19% a year, with a single dip, in the recession year of 2009 (IAB/PwC). In 2016 it passed television advertising in the United States for the first time; eMarketer put TV's revenue that year at $71.3 billion (Module 1 gives the internet total for 2025).

Worldwide, the shift went further: on Magna's count, companies that sell only digital advertising took 73% of the world's ad revenue in 2025, while the revenue of traditional media owners, the broadcasters, publishers, and outdoor and cinema companies, fell 3% (June 2025). Within digital, the newest money went to stores and other sellers of goods and services. WPP Media estimated that commerce advertising, sold by retailers, marketplaces, and travel and financial companies, reached $178.2 billion in 2025 and passed all television advertising worldwide, $167.5 billion, for the first time (December 2025). Module 9 covers retail media.

Television is itself moving to streaming, a shift Module 11 measures in the US prime-time upfront, the annual advance sale of the coming season's commercial time. Google bought YouTube in 2006 for $1.65 billion; its video advertising is covered in Modules 8 and 11.

3g. The Backlash, TikTok, and the Machines

The targeting that made digital advertising valuable also provoked a backlash. In March 2018 newspapers reported that Cambridge Analytica, a political consultancy, had obtained data on tens of millions of Facebook users through a personality-quiz app. Then the European Union's General Data Protection Regulation (GDPR) took effect; in July 2019 the FTC fined Facebook $5 billion over its privacy practices; and Apple began requiring apps to ask permission before tracking people across other companies' apps, under its App Tracking Transparency (ATT) rules. Critics gave the business model a name: in *The Age of Surveillance Capitalism* (2019), the Harvard scholar Shoshana Zuboff argued that the platforms profit by collecting behavioral data in order to predict and shape what people do. Module 14 covers what each of these events meant for advertising, and the case for and against targeting.

The newest giant came from China: ByteDance launched TikTok outside China in 2017 and merged it with the lip-sync app Musical.ly in 2018. Its main feed chose videos for each viewer from what that viewer watched, rather than from whom the viewer followed, and other platforms copied the design. TikTok said in September 2021 that it had passed one billion monthly users. In the United States its future became a question of law, and the American business passed to a US joint venture (Module 8).

The latest shift is toward machines that write ads, buy them, and answer shoppers' questions directly. Since 2022 generative AI has begun producing ad copy and images, the platforms are automating the buying, and ads have begun to appear inside AI chatbots' answers; Module 17 examines what that may do to advertising.

Across three decades the unit of sale shrank: from a section of a website for a month, to a thousand impressions, to a single click, to one view auctioned for one person in a fraction of a second. Each step made advertising easier to measure and to aim, and moved more of the money to the companies that owned the audience and the data about it: search engines, social networks, and stores.

Module 3 Key Points

The Digital Turn, 1994–2026

• The cookie and the first paid banner (HotWired, October 27, 1994) made the web an advertising medium.

• GoTo.com and Google's AdWords turned search ads into auctions priced per click, ranked by bid and relevance from 2002.

• Facebook's 2007 ad system included Beacon; by late 2012 phones brought it about 23% of its ad revenue.

• US internet advertising passed TV in 2016; worldwide, commerce advertising passed TV in 2025, at $178.2 billion.

• By 2019, exchanges, real-time bidding, header bidding, and first-price auctions had automated the selling of display ads.

• A privacy backlash followed from 2018: Cambridge Analytica, Europe's GDPR, the FTC's fine on Facebook, and Apple's tracking prompt.

Module 3 Quiz

Check Your Understanding

1. How did HotWired price its first banner ads in 1994?

A) A price per thousand impressions, set by auction
B) A price per click, paid only when a visitor clicked through to the advertiser's site
C) A flat fee per month for a section of the site
D) A share of the sales each ad produced

2. What did Google change in 2002 about how AdWords ranked ads?

A) It ranked them by bid alone, as GoTo.com had done
B) It ranked them by bid combined with how often users clicked them
C) It let advertisers buy the top position for a month at a fixed price, as directories did
D) It ranked them by each advertiser's total budget

3. What was Facebook's Beacon, launched in November 2007?

A) Facebook's first app for mobile phones
B) An auction for ads placed beside search results, modeled on Google's AdWords
C) A system that posted users' partner-site purchases to friends' News Feeds
D) A tool that let users block tracking cookies on partner sites

4. According to WPP Media, what happened to commerce advertising in 2025?

A) It passed all television advertising worldwide for the first time.
B) It fell as shoppers moved to AI assistants.
C) It passed search advertising worldwide.
D) It passed television in the United States but not worldwide.

5. What changes for the winning bidder when an exchange moves from second-price to first-price auctions?

A) It pays only when the viewer clicks through to its site
B) It must also pay the publisher's ad-server fee
C) It can no longer see which site or app the ad will appear on before it bids
D) It pays its own bid, not the runner-up's

Module 4 — Brands and Persuasion: Why People Pay More

What a Brand Is, What It Is Worth, and How Ads Work on Memory

Recommended Reading / Reference

Book: Byron Sharp, *How Brands Grow: What Marketers Don't Know* (2010), the Ehrenberg-Bass Institute's argument, drawn from decades of shoppers' purchase records, that brands grow by being easy to think of and easy to buy rather than by deepening loyalty

Report: Kantar, *Kantar BrandZ Most Valuable Global Brands 2026* (May 14, 2026), the survey-based ranking with its notes on method, best read beside Interbrand's finance-based *Best Global Brands 2026* (October 28, 2026)

4a. Three Ways to See a Brand

A brand begins as a legal right: in the United States the Lanham Act of 1946, the federal trademark law, defines a trademark as any word, name, symbol, or device used to identify one seller's goods, distinguish them from others', and indicate their source, “even if that source is unknown.” A registration can be renewed for as long as the mark stays in use, and its owner can stop rivals from using anything confusingly similar. The idea is older than the statute: the red triangle of Bass, a British brewer, was the first mark entered on Britain's trademark register when it opened in 1876. The same federal law lets competitors sue one another over false advertising (Module 15).

To an economist the name is a signal, telling a shopper that this box comes from the same maker as the last one, so a satisfied buyer can buy again without testing the product and a disappointed one can avoid it. A maker with a valuable name also has more to lose from a bad batch, so the name works as a pledge of consistent quality. Brands spread when railways and national markets carried goods far from the people who made them, and buyers could no longer judge the maker in person (Module 2).

Marketers now also describe a brand as a memory: the associations that a name, a logo, a color, or a sound calls up in a buyer's mind. On this view a brand lives less in the factory or the trademark file than in millions of heads, and advertising's job is to build and refresh those memories so that the brand comes to mind when someone is ready to buy (section 4d). The three views fit together: the law protects the mark, the signal gives it economic value, and memory makes it work at the moment of choice. Whatever buyers pay for a name, beyond what they would pay for the same product without it, is the brand premium.

4b. What Brands Are Worth, and Why the Rankings Disagree

Two consultancies publish the most-cited league tables of brand value, and their totals differ by a factor of more than three. Interbrand released its *Best Global Brands 2026* on October 28, 2026, and Kantar its BrandZ ranking on May 14, 2026, when Google displaced Apple at the top after Apple's four years there. Both lists are led by American technology companies, but they disagree most about scale: Kantar's figure for Google is almost four times Interbrand's.

Rank	Interbrand, October 2026	Kantar BrandZ, May 2026
1	Apple, $514.6 billion (+9%)	Google, $1.48 trillion (+57%)
2	Microsoft, $443.7 billion (+14%)	Apple, $1.38 trillion (+6%)

Rank	Interbrand, October 2026	Kantar BrandZ, May 2026
3	Google, $389.3 billion (+23%)	Microsoft, $1.11 trillion (+26%)
4	Amazon, $351.8 billion (+10%)	Amazon, $1.02 trillion (+18%)
5	Samsung, $105.2 billion (+16%)	NVIDIA, $814.9 billion (+60%)
Top 100	$3.9 trillion (+8.1%)	$13.1 trillion (+22%)

Interbrand, *Best Global Brands 2026* (October 28, 2026), as reported by Blog du Modérateur and Mercado & Consumo; Kantar, *Kantar BrandZ Most Valuable Global Brands 2026* (May 14, 2026). Changes are against each firm's own values for 2025.

Neither number is a price, because brands are almost never sold apart from their companies. Interbrand starts from a forecast of the profit a company earns on its branded products after a charge for the capital it uses. It takes the share of buying decisions it attributes to the brand and discounts the result at a rate set by a brand-strength score. Kantar starts from the earnings it attributes to the brand, applies a multiple drawn from financial forecasts, and scales the result by a “brand contribution” measured in its own consumer surveys. Every step takes judgment, and a change in one multiple can move a valuation by hundreds of billions of dollars.

Accountants avoid the question: under both US and international accounting rules, a company records a brand on its balance sheet only when it buys one; a brand built by decades of its own advertising appears nowhere in its accounts. The rankings also show brand value moving away from the packaged-goods names that built modern marketing: Coca-Cola, eighth on Interbrand's list at $60.3 billion, was flat in 2026, while NVIDIA's value rose 71% on Interbrand's count.

4c. Store Brands and the Price of a Name

Store brands, sold under a retailer's own name such as Walmart's Equate or Costco's Kirkland Signature, are the cleanest test of what a name is worth, because many are close copies of the national brands on the same shelf. In the United States they sold $282.8 billion in 2025, up 3.3% against 1.2% for national brands, according to the Private Label Manufacturers Association (PLMA) and the data firm Circana (January 20, 2026). That was 21.3% of dollar sales, up from 19.1% in 2021, and a record 23.5% of units, up from 21.6%. Their share of units exceeds their share of dollars because they cost less.

Europeans buy far more of them: across 17 countries tracked by PLMA International and NielsenIQ, store brands took 38.7% of sales by value in the year to early October 2025. Switzerland was highest, at 52.3%, and Germany, Britain, and France together came to 40.3%.

The people who know most about a product pay least for its name. Bart Bronnenberg, Jean-Pierre Dubé, Matthew Gentzkow, and Jesse Shapiro (2015) matched households' purchase records with their occupations. For headache remedies, whose store and national versions contain the same active ingredients, pharmacists chose the national brand only 9% of the time, against 26% for the average shopper. Chefs devoted 12 percentage points less of their purchases of pantry staples such as salt and sugar to national brands than similar shoppers who were not chefs. The authors concluded that missing information explains much of the premium on health products but little of it in food and drink.

The authors also estimated the total at stake: shoppers who bought the store brand wherever one was available would save about $44 billion a year, 22% of their $196 billion of spending in the 420 product groups studied. A brain scanner shows why a name can be worth paying for even when the contents match. Samuel McClure and colleagues (2004) found that when people tasted Coke and Pepsi blind, their stated preferences tracked the response of a brain region associated with reward. When the cups were labeled, knowing that a cup held Coke, but not knowing that one held Pepsi, changed both what people said they

preferred and how their brains responded. The worked example at the end of this module prices one such premium.

4d. How Advertising Works on Memory

The most influential account of how advertising works comes from the Ehrenberg-Bass Institute at the University of South Australia, set out for general readers in Byron Sharp's *How Brands Grow*. Its central idea is mental availability: the chance that a brand is noticed or comes to mind when a buyer is in a buying situation. Its companion is physical availability, being easy to find and buy. Most purchases, on this account, are quick and habitual, so the brand that comes to mind first and sits within reach usually wins.

Two regularities in purchase data shape the institute's advice. The first, double jeopardy, named by the sociologist William McPhee in 1963 and documented by Andrew Ehrenberg, holds that small brands suffer twice: they have fewer buyers, and those buyers are slightly less loyal. Since loyalty moves with size, a brand grows mainly by adding buyers rather than by deepening their loyalty. The second is the weight of light buyers: a brand's heaviest 20% of buyers “generally contribute not much more than half” of its sales, not the 80% that the Pareto rule implies, and almost half comes from the lightest 80% (Sharp, Romaniuk and Graham, 2019). Advertising aimed only at loyal customers misses where growth lies.

Memories need cues, which Jenni Romaniuk of the institute calls distinctive assets: the colors, logos, characters, shapes, and sounds that identify a brand without its name, such as Coca-Cola's red script or Intel's five-note chime. Assets work only if they are used consistently for years, which is why the institute warns against frequent redesigns. Emotion helps an ad be noticed and remembered, and Les Binet and Peter Field's analyses of the IPA's databank of award entries found that campaigns built mainly on emotion reported large business effects more often than rational ones, especially over the long run (Module 5 describes the databank). Plain information still works where there is news, such as a price cut or a new product.

Repetition does part of the work: the psychologist Robert Zajonc showed in 1968 that people come to like words, shapes, and faces more the more often they see them, the mere-exposure effect, and later studies found it even when people could not recall the exposures (Kunst-Wilson and Zajonc, 1980). Each repetition adds less than the one before, and an ad seen too often can irritate its audience. What repetition does not need is a hidden channel: the claim that flashed messages could make audiences buy began as a hoax (Module 2). Laboratory studies since have found subliminal brand cues nudging choice only at the margin, such as one drink brand chosen more often by people who were already thirsty (Karremans, Stroebe and Claus, 2006).

4e. The Economics: Information, Persuasion, or Complement

Economists have three theories of what advertising does, and each implies a different verdict on its worth. The first treats it as information: George Stigler (1961) argued that buyers search for prices and quality at a cost, and that advertising lowers the cost, which should narrow the spread of prices for the same good. Phillip Nelson (1970, 1974) divided products into search goods, whose qualities a buyer can check before buying, such as a dress's color, and experience goods, which can be judged only by use, such as a soup's taste. For search goods, ads carry checkable facts; for experience goods, Nelson argued, the spending itself is the message.

The logic is that of a costly signal: a maker spends heavily to win first purchases only if it expects repeat sales, and it can expect them only if the product is good. A lavish campaign therefore tells buyers something

even when its content is a jingle. Paul Milgrom and John Roberts (1986) worked out when such signals can be believed. On this theory, much advertising that looks empty is informative.

On the second theory, advertising is persuasion: Nicholas Kaldor (1950) argued that most advertising aims to change tastes rather than inform, wastes resources, and helps large firms dominate. John Kenneth Galbraith (Module 2) went further, arguing that producers create the wants their products then satisfy. William Comanor and Thomas Wilson (1967) found higher profits in industries that advertise heavily, and read advertising as a barrier to entry: a newcomer must match the incumbents' budgets to be noticed.

The third theory sidesteps the choice: Gary Becker and Kevin Murphy (1993) treated advertising as a good consumed along with the product, one that can raise the pleasure of using it, as an ad's imagery may make a beer more enjoyable to share. Tastes stay fixed; what changes is the product as experienced. On this view a premium paid for an advertised brand can buy something real to the buyer, even when the contents match a store brand's. Becker and Murphy also allowed that ads can be a nuisance, which is why some viewers pay to avoid them.

4f. Does Advertising Inform or Manipulate? The Cases as Their Advocates Make Them

The case that it informs: Economists in the tradition of Stigler and Nelson, and the advertising industry's trade bodies, argue that advertising is how sellers tell buyers what exists, what it costs, and where to find it. Their strongest evidence is what happens when it is banned. Lee Benham (1972) found that eyeglasses cost about 25% more in US states that barred optometrists from advertising than in states that allowed it. In *Virginia State Board of Pharmacy v. Virginia Citizens Consumer Council* (1976), the Supreme Court struck down a ban on advertising drug prices, holding that a consumer's interest in “the free flow of commercial information” may be as keen as his interest in “the day's most urgent political debate.”

The pattern recurred after the Court struck down Rhode Island's ban on liquor price advertising in *44 Liquormart v. Rhode Island* (1996): the prices of the products stores advertised fell, while the prices of other products did not (Milyo and Waldfogel, 1999). Daniel Ackerberg (2001) found that television ads for a new yogurt moved households that had never bought it but not those that had, the signature of information rather than image. Even ads with no facts in them, on Nelson's argument, tell buyers that the maker is betting on repeat sales.

The case that it manipulates: Kaldor, Galbraith, and their successors, now joined by behavioral economists and consumer advocates, argue that most advertising works by exploiting the shortcuts and feelings of busy people rather than by informing them. In a field experiment with a South African lender's mailed loan offers, adding a photo of an attractive woman raised loan demand about as much as cutting the interest rate by 25% (Bertrand, Karlan, Mullainathan, Shafir and Zinman, 2010). The photo said nothing about the loan. The Coke experiment shows a label changing the experience itself, and the pharmacists' choices show who pays: the shoppers who know least.

On this view the premium on identical products is a transfer from the less informed to the owners of famous names. Heavy advertising also shields incumbents, since a newcomer must outspend them to be noticed, and higher prices follow. The harm is greatest where buyers are least able to judge, as with children or products whose costs are hidden (Module 15).

The evidence: Both effects are real, and they show up in different places. Advertising most clearly informs and lowers prices where it carries prices or news of a new product, as the eyeglass, liquor, and yogurt studies show. Effects unrelated to information are also measurable, as in the loan mailing, but they work at

the margin of decisions people were already making. The premium on identical goods falls steeply with expertise, which points to missing information that advertising does not correct and may exploit. Becker and Murphy's point also holds, since some buyers value the brand itself. Experiments find advertising's effect on sales small on average (Module 5), which sits poorly with the picture of advertisers who can make people buy at will.

The course's reading, stated once: advertising informs more than its critics allow, especially about prices and new products, and its persuasive power is real but modest. The clearest cost falls on the least informed buyers, who pay large premiums for identical goods. The remedy the evidence supports is better information for buyers, such as clear unit prices, more than limits on advertising, which have tended to raise prices.

Worked example — The Brand Premium on Ibuprofen

Inputs: prices on Walmart.com on October 31, 2026, for Advil coated tablets, 200 mg of ibuprofen each, 100 count, at $10.98, and for Walmart's Equate ibuprofen tablets, 200 mg, 100 count, labeled “compare to Advil Tablets active ingredient,” at $1.98. Assume a household uses 300 tablets a year, about three a week for each of two adults. Step 1, the price per tablet: $10.98 ÷ 100 = 11.0 cents for Advil and $1.98 ÷ 100 = 2.0 cents for Equate. Step 2, the premium: $10.98 − $1.98 = $9.00 per 100 tablets, so the national brand costs 5.5 times as much, and 82% of its price is premium over the store brand's. Step 3, a year's supply: 3 × $10.98 = $32.94 for Advil against 3 × $1.98 = $5.94 for Equate, a difference of $27.00 a year, or $270 over ten years before inflation. Step 4, check whether buying in bulk closes the gap: Advil's 200-count bottle at $17.98 comes to 9.0 cents a tablet and Equate's 500-count at $7.98 to 1.6 cents, still 5.6 times as much. Result: the household pays about $27 a year for the name on a pill with the same active ingredient and dose. The pharmacists of section 4c, who know the contents match, mostly decline to pay it; the shoppers who do are buying reassurance, habit, and a memory that advertising built.

Module 4 Key Points

Brands and Persuasion: Why People Pay More

• A brand is a trademark in law, a pledge of consistent quality in economics, and a set of memories in buyers' minds.

• Brand values are model estimates, not prices: Kantar's 2026 top 100 is worth more than three times Interbrand's.

• Store brands took a record 23.5% of US unit sales in 2025, and Europeans buy them more still.

• Experts rarely pay for the name: pharmacists choose national-brand headache remedies only 9% of the time.

• Brands grow mainly by adding buyers, most of them light ones; the heaviest fifth supply only about half of sales.

• The course's reading: advertising both informs and persuades, and the premium on identical goods falls mostly on the least informed.

Module 4 Quiz

Check Your Understanding

1. In Bronnenberg and colleagues' study, how often did pharmacists choose a national-brand headache remedy over a store brand?

A) About 26% of the time, like other shoppers
B) 9% of the time, far less than the average shopper
C) More than half the time, because they trust the large makers' quality control
D) About 40% of the time, more often than chefs bought national brands of salt and sugar

2. Why do Interbrand's and Kantar's brand values differ by a factor of more than three?

A) Kantar counts only sales in the United States
B) Their models differ, and Kantar scales its values by a brand contribution from surveys
C) Interbrand leaves technology companies out of its top 100
D) One values brands alone, while the other reports each company's total stock-market value

3. What is double jeopardy in the Ehrenberg-Bass account of buying?

A) Small brands have fewer buyers, and those buyers are also slightly less loyal
B) A brand that cuts its price loses both margin and image
C) Heavy buyers defect twice as fast as light ones
D) An ad shown too often loses attention and goodwill at once

4. Which finding supports the case that advertising informs?

A) A photo of a woman on a loan offer raised demand as much as a 25% rate cut
B) Knowing a cup held Coke changed tasters' brain responses
C) Eyeglasses cost about 25% more where optometrists' ads were banned
D) Store brands reached a record share of US unit sales in 2025, at 23.5%

5. In the module's worked example, how did Advil's price per tablet compare with Equate's at Walmart.com on October 31, 2026?

A) About twice as much
B) About 20% more, a typical premium for a national brand
C) The same, once both are bought in 200-count bottles
D) 5.5 times as much

Module 5 — Does Advertising Work? The Evidence

Experiments, Elasticities, and Why Measuring Is Hard

Recommended Reading / Reference

Book: Les Binet and Peter Field, *The Long and the Short of It: Balancing Short and Long-Term Marketing Strategies* (2013), the IPA's study of its databank of award-entered campaigns and the source of the brand-to-activation ratios practitioners quote

Reference: Bradley Shapiro, Günter Hitsch and Anna Tuchman, “TV Advertising Effectiveness and Profitability: Generalizable Results From 288 Brands,” *Econometrica* (July 2021), the largest study of what television advertising does to the sales of packaged goods, technical in its methods but readable in its introduction and conclusions

5a. Why Measuring Is Hard

Advertising is among the largest purchases companies make whose effect they cannot see directly. The difficulty is not that ads do nothing but that their effect on any one person's spending is tiny next to everything else that moves it: payday, the weather, a move, a birthday. Randall Lewis and Justin Rao (2015) analyzed 25 large field experiments run with major US retailers and financial brokerages, covering $2.8 million of advertising. In such campaigns the standard deviation of a customer's spending was often about ten times its average, while a profitable campaign might lift the average by only a few percent.

The result is what they called the unfavorable economics of measuring the returns to advertising. Across their experiments, the median 95% confidence interval on the return on investment, the range of returns consistent with the data at the usual standard of proof, was more than 100 percentage points wide, so a campaign could plausibly have lost half its cost or doubled it. To tell a campaign that pays from one that does not, an advertiser would often need an experiment covering more than ten million person-weeks. Few campaigns are that large, and few advertisers run the test.

Without experiments, advertisers compare people who saw an ad with people who did not, and the comparison misleads. Ads are aimed at people likely to buy anyway, such as those who searched for the product, visited the site, or bought before, so the gap between the exposed and the unexposed mixes the ad's effect with the targeting's. Sales also rise when ads run because advertisers spend most in the seasons when demand is highest. The fix is a randomized experiment: hold back a random group from the ads and compare it with everyone else. The difference is the lift, and the worked example at the end of the module shows how to read one.

Economists summarize ads' effect on sales with an elasticity: the percentage change in sales that follows a 1% change in advertising. An elasticity of 0.1 means that 10% more advertising raises sales by 1%, and whether that pays depends on how large the ad budget is relative to sales and on the profit margin.

5b. The Experiments: eBay, Facebook, and Ghost Ads

In the most cited test, Thomas Blake, Chris Nosko, and Steven Tadelis (2015) first stopped eBay's search ads on its own name on Microsoft's and Yahoo's search engines. About 99.5% of the clicks those ads had brought came back through the free listings below them, because people who type “eBay” are already

looking for eBay. The researchers then switched off eBay's ads on ordinary search terms in 68 of 210 US television markets for 60 days. The short-term return was about −63%, with a 95% range of about −124% to −3%: the ads lost money.

The ads did move some people: new and infrequent eBay users bought more when they saw them, but frequent users, who received most of the spending, did not, since they would have come anyway. The lesson spread quickly: measurements that credit a sale to the last ad clicked before it can make ads that change nothing look like a firm's best investment (Module 13).

Data from Facebook let Brett Gordon, Florian Zettelmeyer, Neha Bhargava, and Dan Chapsky (2019) test the shortcuts directly: they took 15 large advertising experiments on the platform, covering 1.6 billion impressions, and compared their results with what standard observational methods estimated from the same data without the randomization. The observational methods usually overstated the lift, even when they used the detailed data Facebook holds on each user.

Experiments are costly because the held-back group must be shown something else or tracked with care. Garrett Johnson, Randall Lewis, and Elmar Nubbemeyer (2017) proposed ghost ads: the ad system notes each time a member of the control group would have been shown the ad, so the comparison is between equally targeted people, without buying placebo ads. Their test of a retailer's retargeting campaign, ads that follow people who have visited its website, found 17% more site visits and 11% more purchases among those shown the ads. Platforms now sell lift tests built on such designs, though the platform selling the ads also runs the test (Module 13).

5c. Television and the Shrinking Estimates

Television has the longest record of study, and its largest test is recent: Bradley Shapiro, Günter Hitsch, and Anna Tuchman (2021) estimated the effect of TV advertising on the sales of 288 brands of packaged goods, using Nielsen's weekly store sales and advertising data. They compared neighboring areas on either side of the boundaries between TV markets, which receive different amounts of a brand's advertising for reasons unrelated to local demand. For about two-thirds of the brands the effect could not be distinguished from zero, and the typical estimate, in the table, was far below the earlier literature's.

The returns were worse: for more than 80% of the brands, the last dollar of TV advertising returned less than it cost, so they would have earned more by spending less. Even judging all the advertising they actually ran, rather than the last dollar, only about a third of the brands earned a positive return.

Meta-analyses, which pool the results of many studies, show the estimates shrinking over the decades. Korkames, Stanley, and Stremersch (2026) corrected for publication bias, the tendency of studies that find no effect to go unpublished; internet ads showed the largest effects in their data.

Study	What it covers	Elasticity found
Sethuraman, Tellis and Briesch (2011)	Meta-analysis of 56 studies published 1960-2008: 751 short-term and 402 long-term estimates	Short-term 0.12 (against 0.22 in a 1984 meta-analysis); long-term 0.24
Shapiro, Hitsch and Tuchman (2021)	TV advertising for packaged-goods brands, 2010-2014	Median 0.014; mean 0.023
Korkames, Stanley and Stremersch (2026)	Meta-analysis of 538 estimates, corrected for publication bias	Short-term 0.0008; long-term 0.03

Sources: the papers as published (*Journal of Marketing Research*, 2011; *Econometrica*, 2021; *International Journal of Research in Marketing*, 2026). The 1984 meta-analysis is Assmus, Farley and Lehmann's.

The arithmetic shows why such small effects rarely pay: at the table's median TV estimate, a brand with $1 billion of annual sales that raised a $50 million TV budget by 10% would gain 0.14% in sales, or $1.4 million, and at a 30% margin that is $420,000 of profit for $5 million of extra spending.

5d. Where Advertising Does Pay

The same research shows where ads earn their keep. They work best on people who do not already buy, as the eBay experiment showed, which fits the Ehrenberg-Bass finding that growth comes from light buyers (Module 4). New products, which have news to tell and no habits to overcome, are where the informative effect shows most clearly, as with the yogurt ads in Module 4.

Effects also outlast the campaign, when there is an effect to begin with. In 55 in-market tests of extra TV weight for established packaged goods, Leonard Lodish and colleagues (1995) found that where the extra advertising raised sales in its first year, the total gain over three years came to roughly double the first year's. Where it had no effect in the first year, it had none later, so long-run gains, on this evidence, come from campaigns that also work in the short run.

Stopping shows the same thing from the other side, as Ehrenberg-Bass researchers found when they studied 57 cases in which Australian packaged-goods brands cut all mass-media advertising for a year or more (Hartnett and colleagues, 2021). On average, sales were 16% below the last advertised year after one year, 25% after two, and 36% after three, with wide variation: large brands that had been growing kept growing for a year or two. The brands chose to go dark, often for financial or strategic reasons, so some may already have been weakening.

The biggest single advertising event shows effects that are real but partly offsetting. Wesley Hartmann and Daniel Klapper (2018) studied beer and soda sales in the 56 largest US markets from 2006 to 2011, using the fact that a market's Super Bowl audience rises when its local team plays. Brands' Super Bowl ads raised their sales in the following weeks, notably during later sports broadcasts, but when the two leading colas both advertised, their gains largely cancelled out.

Political advertising shows the same pattern of small, short effects. Tests of 49 ads from the 2016 presidential race on about 34,000 people found small average effects on views of the candidates and on vote choice (Coppock, Hill and Vavreck, 2020). Texas Governor Rick Perry's randomized TV ads of 2006 raised his standing in the weeks they ran, and the effect was gone within a week of their stopping (Gerber and colleagues, 2011).

5e. The Practitioners' Evidence: The IPA Databank

Practitioners rely on different evidence, much of it from the award entries that Britain's IPA has collected since 1980 for its advertising effectiveness prizes, each with the entrant's own data on sales, market share, and profit. Les Binet and Peter Field (2013) analyzed 996 of these campaigns, for 700 brands in 83 categories, and drew a distinction that has shaped budgets ever since. Brand-building campaigns raise baseline sales slowly and lastingly by building memories, while activation campaigns, such as price promotions and direct-response ads, produce quick spikes that fade. The most effective campaigns combined the two, with about 60% of spending on brand building and 40% on activation.

Their later work refined the ratio: in *Effectiveness in Context* (2018) the average best split was 62:38, with the brand share varying by category and circumstance up to about 80%. For business-to-business marketing, a study for LinkedIn's B2B Institute (2019) put it at 46:54, with more weight on activation. These

ratios became rules of thumb across the industry, and the argument over how budgets should follow them is in Module 6.

The databank's limits are plain: its campaigns were entered for awards, so they are a self-selected set of successes, and their effects are those the entrants report, usually estimated with econometric models rather than experiments. The ratio is a correlation among such entries, not a test of what would happen if a given brand changed its split.

5f. Does Most Advertising Pay for Itself? The Cases as Their Advocates Make Them

The case that it does not: Economists who run advertising experiments, among them Blake, Nosko and Tadelis and Shapiro, Hitsch and Tuchman, argue that much of the money spent at the margin is wasted, and Lewis and Rao showed why firms rarely notice. They hold that their evidence, randomized or close to it, is the strongest available and points one way: eBay's search ads lost money, most TV advertisers would earn more by spending less, and once publication bias is removed the average measured effect is close to zero. Firms keep spending, on this view, because the measurements they rely on credit ads with sales that would have happened anyway, and because the agencies and platforms that sell advertising also supply much of the measurement.

The case that it does: Practitioners and many marketing scientists, among them Binet and Field and the Ehrenberg-Bass Institute, argue that short-run tests miss most of what advertising does. They capture a few weeks of response to a little extra spending, not the long-run work of keeping a brand in buyers' memories. Long-run effects are at least as large as short-run ones, brands that stop advertising lose sales year after year, and campaigns that win new and light buyers build sales that last. A marginal dollar that does not pay says nothing about the average dollar, and even Shapiro and colleagues found that some brands' advertising paid overall (section 5c).

The evidence: The two sides measure different things, and each measurement is sound as far as it goes. At the margin, for established brands over the weeks that experiments can track, most of the advertising studied does not recover its cost. The long-run effects that defenders cite are real but harder to establish: the strongest evidence, the in-market tests of 1995, shows them only for campaigns that already worked in the short run, and the going-dark and IPA studies are correlations among self-selected cases. Advertising pays most reliably where it reaches new buyers or carries news, and least reliably where large, established brands spend heavily on people who already buy them.

The course's reading, stated once: most advertising probably does not pay for itself at the margin, and many large advertisers would earn more by spending less or spending differently. Advertising as a whole is not wasted: it pays most clearly for new products and new buyers, and it helps sustain established brands over time. The waste persists because measuring returns well is costly and rare, and the measuring is often left to those who sell the ads.

Worked example — Reading a Lift Test

Inputs: an online retailer randomly assigns 2 million customers, 1.6 million to see a four-week campaign and 400,000 to a holdout group that sees none of it. In the holdout, spending averages $7.00 a customer with a standard deviation of $75, a level of noise typical of Lewis and Rao's experiments, and the exposed group averages $7.35. The campaign costs $224,000, or $0.14 per exposed customer, and the retailer keeps 50 cents of gross margin on each dollar of sales. Step 1, the lift: $7.35 − $7.00 = $0.35 a customer, 5% of the holdout's average. Step 2, the return: $0.35 × 1.6 million = $560,000 of extra sales, worth $280,000 in margin; less the $224,000 cost, that leaves $56,000, a return on investment of 25%. Step 3, the uncertainty: the standard error of the difference between the two averages is $75 × √(1 ÷ 1,600,000 + 1 ÷ 400,000) = $0.133, so the 95% range on the lift is $0.35 ± 1.96 × $0.133, or $0.09 to $0.61. Step 4, carry the range through: a lift of $0.09 yields $0.045 of margin per customer against $0.14 of cost, a return of −68%, while a lift of $0.61 yields $0.305, a return of +118%. Step 5, compare with zero and with break-even: the lift is 2.6 standard errors above zero, so the ads almost certainly raised sales, but break-even needs a lift of $0.14 ÷ 0.5 = $0.28, only 0.5 standard error below the estimate, so the test cannot say whether the campaign paid. Step 6, the cost of certainty: narrowing the return's range to plus or minus 25 points with the same 80:20 split would take about 27.6 million customers. Result: a test of 2 million people shows that the ads worked and leaves their profitability anywhere from −68% to +118%, the problem Lewis and Rao found across real campaigns.

Module 5 Key Points

Does Advertising Work? The Evidence

• Across 25 large experiments, the typical 95% range on advertising's return was more than 100 percentage points wide.

• When eBay stopped buying search ads on its own name, 99.5% of the lost clicks came back free.

• For more than 80% of the TV advertisers studied, the last ad dollar returned less than it cost.

• Corrected for publication bias, the average short-term advertising elasticity is close to zero: 0.0008.

• Ads pay best on new and light buyers; long-run gains follow only campaigns that also work in the short run.

• The course's reading: most advertising does not pay at the margin, and waste persists because good measurement is costly and rare.

Module 5 Quiz

Check Your Understanding

1. What happened when eBay stopped buying search ads on its own name?

A) Its sales fell by about 63% until the ads resumed
B) Rival sellers bought the term and drew away a fifth of eBay's visitors
C) About 99.5% of the lost clicks came back through free listings
D) Frequent users stopped visiting, so eBay restored the ads within a few weeks

2. What does an advertising elasticity of 0.12 mean?

A) Each ad dollar returns 12 cents of profit
B) Advertising causes 12% of a brand's sales
C) Sales must rise 12% to cover a doubling of the ad budget
D) A 10% rise in advertising raises sales by 1.2%

3. Why did the 2026 meta-analysis by Korkames, Stanley and Stremersch find much smaller effects than earlier ones?

A) It studied only internet advertising
B) It corrected for publication bias, the tendency of studies that find no effect to go unpublished
C) It measured profit rather than sales
D) It relied on one platform's experiments, which run for only a few weeks

4. In a ghost-ads experiment, what happens to the control group?

A) It sees its usual ads while the system logs would-be exposures
B) It is shown a charity ad that the advertiser pays for, not the real one
C) It sees the same ad, but less often, so the effect of frequency can be measured
D) It is kept from seeing any advertising on any site during the test

5. What can the retailer conclude from the module's worked lift test?

A) The campaign returned 25%, with 95% confidence
B) The ads lost money, since the range of returns includes losses
C) Little, because a holdout of 400,000 is too small to detect any lift
D) The ads raised sales, but the test cannot tell if they paid

Module 6 — Who Buys Advertising, and How a Budget Is Set

The Advertisers, the Rules of Thumb, and a Launch Budget

Recommended Reading / Reference

Book: Les Binet and Peter Field, *Media in Focus: Marketing Effectiveness in the Digital Era* (2017), the IPA's follow-up study of how budgets, share of voice, and the balance between brand and activation spending work once most media are digital

Report: Ad Age Datacenter, *Brand Report 2026* (October 2026), the annual rankings of the 50 biggest advertisers in the United States and worldwide, with estimated spending for 2025

6a. The Biggest Advertisers

The largest buyers of advertising are increasingly the online marketplaces that also sell it. Ad Age's *Brand Report 2026*, reported by Meio & Mensagem and AdIndex on October 14, 2026, ranked Alibaba the world's largest advertiser in 2025, passing Amazon. The top 50 advertisers worldwide spent $326.2 billion, up 12.7%, against 6.6% growth the year before; all the figures are Ad Age's estimates.

Rank	Advertiser	Main business	2025 spending	Change
1	Alibaba	Online marketplaces, cloud computing	$30.2 billion	+91.7%
2	Amazon	Online retail, cloud computing, streaming	$23.5 billion	+9.8%
3	PDD Holdings	Online marketplaces (Pinduoduo, Temu)	$16.3 billion	+12.5%
4	L'Oréal	Cosmetics	$16.0 billion	+5.6%
5	Procter & Gamble	Household and personal-care goods	$13.5 billion	+8.9%
6	LVMH	Luxury goods, wine and spirits	$10.4 billion	−1.5%
7	JD.com	Online retail and delivery	$10.1 billion	+89.3%
8	Samsung Electronics	Electronics	$10.0 billion	+7.2%
9	Alphabet	Search, video, devices, cloud computing	$9.9 billion	+13.8%
10	Nestlé	Food and drink	$9.3 billion	+9.8%

Ad Age, *Brand Report 2026* (estimated worldwide spending in 2025), as reported by Meio & Mensagem and AdIndex, October 14, 2026. PDD's growth is 12.5% in AdIndex's report and 15.5% in Meio & Mensagem's.

Below the top ten came Unilever, Comcast, Booking Holdings, Disney, and American Express. Among the top 25, the fastest climbers after the Chinese platforms were a bank and an insurer, Capital One, up 29%, and Progressive, up 28.2%, while Unilever's spending fell 9.7%. The list also shows a circularity in the business: Amazon, Alibaba, and Alphabet are among the world's largest sellers of advertising as well as buyers of it (Modules 8 and 9).

In the United States, the top 50 advertisers spent $144.5 billion in 2025, up 7.2%. Amazon was the largest, at an estimated $16.1 billion, followed by Comcast, Procter & Gamble, Capital One, and Progressive; only Amazon, P&G, and Alphabet ranked in both the US and the world top ten. Counting single brands rather than companies, in measured media and paid social, Amazon led at $4.8 billion, ahead of Progressive at $2.5 billion, Walmart at $1.9 billion, Verizon at $1.5 billion, and Geico at $1.4 billion, and twelve US brands passed $1 billion.

6b. Millions of Small Advertisers

Behind the giants stand millions of small buyers, a market the self-serve platforms created. A dentist, a bakery, or a one-person online shop can open an account with Meta, Google, Amazon, or TikTok, set a budget of a few dollars a day, and enter the same automated auctions as Procter & Gamble (Module 10). Facebook said it had 10 million active advertisers in the third quarter of 2020, the last figure in Statista's series of the company's counts; Meta's more recent figures count narrower groups, such as advertisers using its AI tools (Module 17).

Small advertisers buy almost entirely through the auctions, mostly ads judged by clicks, calls, and sales rather than campaigns built for memory. A firm with $2,000 a month to spend cannot afford a television campaign, but it can afford clicks, which is one reason the share of spending on performance formats has grown.

6c. Which Industries Advertise

Industries differ widely in how much they advertise. In eMarketer's estimates for the United States, as reported by Inside Radio in 2025, retail was the largest spender, with more than a quarter of all advertising, having passed $100 billion in 2024, more than 85% of it digital. Consumer packaged goods, the food, drink, and household products sold in supermarkets, passed $60 billion in 2024 but were among the slowest-growing. Financial services were expected to pass $50 billion in 2025, growing 11.9%, while telecom companies' digital spending was forecast to grow fastest, by 16.3%, and carmakers' slowest, by 2.2%.

Ad Age's global list shows a similar mix at the top. Of the 25 largest advertisers in 2025, six were online platforms or retailers, eight made consumer goods, four were banks, card issuers, or insurers, two were media and telecom groups (Comcast and Disney), two were carmakers (Volkswagen and Toyota), two sold technology (Samsung and Alphabet), and one sold hamburgers (McDonald's). Carmakers put a larger share of their traditional-media budgets into television than any other industry (eMarketer, 2025), and drug makers are covered with the rules on drug advertising (Module 15).

6d. How Budgets Are Set: The Rules of Thumb

Many budgets start from last year's figure and change at the margin, but five methods shape the change. The oldest is a percentage of sales: spend a fixed share of revenue, as packaged-goods and beauty companies long have (Module 1 gives the averages and two leading examples). It is simple and keeps spending affordable, but it runs backward: when sales fall the budget falls, just when a brand may most need support, and a new product with no sales gets no budget at all.

The second is competitive parity: spend what rivals spend, or the same share of sales. It guards against being outshouted but assumes the rivals know the right amount, and it can lock a whole category into overspending. The third method, share of voice, links a brand's advertising to its rivals' and to its market share. John Philip Jones, writing in the *Harvard Business Review* in 1990, found that brands whose share of voice exceeded their share of the market tended to grow, while those below it tended to shrink, and that small brands usually spent above their market share and big brands below it.

The gap between a brand's share of voice and its market share is its excess share of voice (ESOV), and Binet and Field later measured its payoff in the IPA databank. As the LinkedIn B2B Institute summarized it in January 2020, each 10 points of excess share of voice goes with about 0.6 point of market-share growth a

year for consumer brands, and about 0.7 for business-to-business brands. A brand with 10% of its market and 20% of its category's advertising has 10 points of excess voice and, on the rule, would reach about 10.6% after a year. The rule is an average across award entries, not a promise.

The fourth method is objective and task: decide what the advertising must achieve, such as reaching most of the target audience several times in a launch month, cost each task, and add them up. It is the most logical method and the most demanding, since it requires knowing how much reach and frequency produce how many sales, and the worked example applies it. The fifth is the media mix model, fitted to several years of weekly sales and spending by channel, which estimates each channel's contribution and the point at which more spending adds little (Module 13 covers the tools). Its estimates are only as good as the variation in past spending, so experiments are increasingly used to check them.

6e. Brand or Performance: The Budget Choice

The largest allocation decision is now between brand building and performance advertising, the two kinds of spending defined in the Terms at the front, and the IPA's evidence on how to split them is in Module 5. What companies actually do is measured by surveys. Gartner's 2026 survey of marketing leaders found that awareness and conversion, the two ends of the funnel, took 62.6% of media spending (Module 1 gives the survey's figures on overall budgets).

The tilt toward performance has causes beyond the media plan. Performance ads produce numbers within days, while brand advertising pays over years, and chief marketing officers (CMOs) may not stay long enough to collect: Spencer Stuart's figures, as reported in August 2026, put the average CMO's tenure at S&P 500 companies at 4.1 years, the shortest of any core executive role. Finance chiefs ask for returns they can audit, and attribution systems, which credit sales to the ads people clicked last, supply such returns most readily for search, social, and retail media (Module 13).

The case for brand spending rests on the buyers that performance ads miss: most people are not shopping for a given product in a given month, and performance ads harvest demand that something else created. Whether the balance has gone too far is the contested question.

6f. Has Marketing Swung Too Far Toward Short-Term Performance Ads? The Cases as Their Advocates Make Them

The case that it has: Binet and Field, the IPA, and many brand strategists argue that the measurable has crowded out the effective. Attribution credits performance ads with sales that brand advertising created or that would have happened anyway, as eBay's experiment showed (Module 5), so budgets drift toward them. The example most often cited is Adidas. At an industry conference on effectiveness in October 2019, its global media director, Simon Peel, said the company had put about three-quarters of its advertising budget into performance ads because its attribution tools “encourage you to invest more and more in digital.” Econometric modeling then found that brand activity drove more than 65% of its sales, in shops and online, and Adidas moved toward a 60:40 split in favor of brand.

Airbnb went further. Its brand and performance marketing fell 58% in 2020, to $482 million, with performance ads accounting for most of the cut, and its chief executive, Brian Chesky, said the company could “take marketing down to zero and still have 95% of the same traffic.” Airbnb said it would not return to its 2019 spending and would favor brand marketing and public relations (Campaign, reporting the company's

2020 results). James Hurman of the brand-tracking firm Tracksuit summed up the case in 2022: “marketers tend to spend too much of their budgets on performance and too little on brand.”

The case that it has not: Critics answer that the rule behind the alarm is weak. Byron Sharp of the Ehrenberg-Bass Institute, no advocate of narrow targeting, called the 60:40 rule “terrible” and “very misleading” in August 2022: it came, he said, from “a very weird data set, which is award submissions,” and activation has no agreed definition, so no one can check a split. Performance marketers add that their ads are increasingly judged by holdout experiments rather than last clicks, while much brand advertising is never tested at all, and that the largest study of television, the classic brand medium, found most brands overspending on it (Module 5).

For small firms and online sellers, they argue, performance ads are the sales channel itself, and a split drawn from big brands' award entries does not fit them. Binet and Field's own later work found that the best split varies widely by category and situation (Module 5).

The evidence: The case reports are suggestive, not decisive. Adidas's finding came from its own econometric model, and Airbnb cut its marketing in 2020, when the pandemic had collapsed travel, so its cut was not a controlled test. The ratio itself is a correlation among award entries. What experiments do establish is narrower: last-click attribution overstates the return on performance ads, especially those aimed at people already searching for the brand, and the overstatement pushes budgets toward them. Experiments also find that much television advertising fails to pay at the margin, so neither kind of spending has a monopoly on waste.

The course's reading, stated once: the alarm is right about the mechanism and overstated about the remedy. Budgets have tilted toward performance ads partly because last-click measurement flatters them, and firms that rely on it probably underinvest in reaching future buyers. A fixed ratio is no cure, since the evidence for any particular split is weak; the better guide is to test both kinds of spending for their incremental effect and set the split by the results.

Worked example — A $25 Million Launch

Inputs: a new consumer product launched nationally, aimed at 100 million US adults, with the first-year plan in the table below. Reach media are priced per thousand impressions at sourced 2026 rates, while search and retail media are priced per click, because they capture people already looking rather than reaching new ones. Step 1, impressions: the six reach lines buy 1,062.5 million impressions for $19 million, a blended $17.88 per thousand. Step 2, gross rating points: 1,062.5 million ÷ 100 million = 10.6 impressions per target adult, or about 1,063 GRPs. Step 3, reach: if impressions fell on people at random, the standard formula, 1 − e^(−10.6), where e is the mathematical constant 2.718, would put reach at essentially 100%. Real schedules never achieve that, because heavy viewers and scrollers absorb many of the impressions. A formula that allows for the skew, reach = GRPs ÷ (GRPs + 100), reproduces a published planning example in which about 1,000 GRPs reached 91% of the audience about 11 times (Matt Prentis, 2022). Here it gives 1,063 ÷ 1,163 = 91.4%, about 91 million adults, who see the ads 11.6 times each on average. Step 4, frequency: on the same model 76% see the ads three or more times, but over a year that is less than one exposure a month. Step 5, payback: assume each unit nets the maker $3.00 after the retailer's margin and trade discounts and costs $1.50 to make, leaving $1.50, and that first-year buyers buy four units, worth $6.00 each. Breaking even in the first year takes $25 million ÷ $6.00 = 4.2 million buyers created by the campaign itself, 4.2% of the target. Step 6, repeat buying: if 40% of them buy six units a year for two more years, each first-year buyer is worth $6.00 + 0.4 × 6 × $1.50 × 2 = $13.20, and the break-even falls to 1.9 million buyers. Result: $25 million reaches most of the target but less than once a month, and it pays back

only if the product earns repeat purchases. Only a holdout test can show how many of the buyers the ads caused (Module 5).

Line	Budget	Price used	Basis of the price	What it buys
Linear TV	$7.0 million	$30 per thousand	Between the broadcast and cable prime-time prices in Module 11	233.3 million impressions
Streaming and CTV	$4.0 million	$30 per thousand	Within Stillwater Media's programmatic ranges for the large services, $20-$40 (August 2026)	133.3 million impressions
Online video	$2.5 million	$10 per thousand	YouTube in-stream ads, within Stackmatix's range of $6.20-$12.40 (2026)	250.0 million impressions
Social	$4.0 million	$12 per thousand	Within Stackmatix's ranges for US Facebook ($10.50-$16.80) and North American TikTok ($6.80-$12.40), first quarter of 2026	333.3 million impressions
Creators	$1.0 million	$20 per thousand views	Mid-range creator pricing (Module 12)	50.0 million views
Outdoor	$0.5 million	$8 per thousand	Near the programmatic digital out-of-home average in Module 11	62.5 million impressions
Search	$1.5 million	$5.42 per click	Average across US search campaigns (WordStream/LocaliQ, 2026 benchmarks)	About 277,000 clicks
Retail media	$2.0 million	$1.07 per click	The median Amazon cost per click in Module 9	About 1.87 million clicks
Production	$1.75 million		Assumption: TV and video spots, shorter cuts, images	
Agency fees and measurement	$0.75 million		Assumption: 3% of the budget (Module 7 covers agency pay)	
Total	$25.0 million			

Prices are planning assumptions anchored to the sources named; actual prices vary with targeting, season, and negotiation. Impressions = budget ÷ price × 1,000.

Module 6 Key Points

Who Buys Advertising, and How a Budget Is Set

• Alibaba passed Amazon as the world's largest advertiser in 2025, at an estimated $30.2 billion (Ad Age).

• Beside the 50 giants, millions of small firms buy ads through self-serve auctions, mostly performance ads.

• The share-of-voice rule: each 10 points of excess voice goes with about 0.6 point of market-share growth a year.

• Awareness and conversion, the two ends of the funnel, took 62.6% of media spending in Gartner's 2026 survey.

• A $25 million national launch reaches most of its target but less than once a month; repeat buying decides payback.

• The course's reading: last-click measurement tilts budgets toward performance ads, but no fixed brand-to-performance ratio is well supported.

Module 6 Quiz

Check Your Understanding

1. Which company did Ad Age's Brand Report 2026 rank as the world's largest advertiser in 2025?

A) Alibaba
B) Amazon, for another year
C) Procter & Gamble, the largest consumer-goods advertiser
D) Alphabet

2. Under the excess-share-of-voice rule, what should a brand with 10% of its market and 20% of its category's advertising expect?

A) A doubling of its share within two years
B) Stable share, since voice above share only defends a position
C) About 0.6 point of market-share growth a year
D) A gain of 10 share points in one year

3. What did Adidas's global media director say in October 2019?

A) Its TV advertising had failed, so it moved the budget to search
B) Attribution tools had led it to put most of its budget into performance ads, though brand activity drove most sales
C) It had stopped all brand advertising after an econometric study
D) Performance ads drove 65% of its online sales

4. Why does the module's launch example use a reach formula that allows for skew?

A) Industry rules require it for national campaigns that mix television and digital media
B) Digital media cannot reach light viewers
C) Random exposure would put reach near 100%, which real schedules never achieve
D) TV ratings count only heavy viewers, so the random formula would put reach near 50%

5. In the launch example, which assumption cuts the number of buyers needed to break even by more than half?

A) Buyers keep buying for two more years
B) Moving money from linear TV to streaming
C) Dropping the agency fee
D) Doubling the outdoor budget

Module 7 — The Agencies and Media Buying

The Holding Companies, How They Are Paid, and the Squeeze

Recommended Reading / Reference

Book: Ken Auletta, *Frenemies: The Epic Disruption of the Ad Business (and Everything Else)* (2018), a *New Yorker* writer's reporting on the agency groups as the platforms, the consultancies, and doubtful clients closed in on them

Report: Association of National Advertisers (ANA), *The Continued Acceleration of Principal Media* (March 2026), its survey of marketers on agencies that resell media they have bought for themselves; and K2 Intelligence, *An Independent Study of Media Transparency in the U.S. Advertising Industry* (June 2016), the investigation for the ANA that documented rebates

7a. What an Agency Does

An advertising agency sells two things, judgment and buying power, and its work falls into five trades. Creative agencies devise campaigns and write and design the ads, while media agencies plan where the ads should run, for whom, and how often, then negotiate the prices and place them. Public relations (PR) firms seek coverage the advertiser does not pay for. Data and technology units build the audience lists, dashboards, and models behind targeting and measurement, and production studios make the films, photographs, and thousands of digital variants a campaign now needs.

A large advertiser rarely buys all of this from one firm: it may keep a creative agency for each brand or region, one media agency for the world, and specialists for PR, health care, or marketing in stores. The media assignment carries the most money, since a media agency can handle billions of dollars of one client's spending, and it is paid a small share of what passes through its hands.

Most of these firms belong to a holding company, a listed parent that owns many agency networks, shares their back offices, and pools their clients' media buying. Pooling is the point: a group that buys television time and digital inventory for hundreds of clients can demand lower prices and better positions than any one client could. How the groups were assembled, from Interpublic's founding to WPP's takeovers, is told in Module 2.

The holding companies face four kinds of rival. Consultancies such as Accenture sell marketing alongside the technology projects they already run for the same clients, the platforms let advertisers buy directly through self-serve tools (Module 8), and advertisers build agencies of their own (7g). Independent agencies, owned by their founders or partners, compete for creative work and for the clients the big groups lose.

7b. The Holding Companies After the Merger

The industry's map changed on November 26, 2025, when Omnicom completed its purchase of the Interpublic Group (IPG), announced on December 9, 2024. It paid in shares, 0.344 of an Omnicom share for each IPG share, which left Omnicom's existing shareholders with 60.6% of the combined company (Omnicom's filing). Within weeks it retired three of the best-known names in advertising: DDB (Doyle Dane Bernbach) and MullenLowe were folded into Omnicom's TBWA network, and FCB (Foote, Cone & Belding) went into BBDO. Omnicom also announced more than 4,000 job cuts (Mediaweek, December 2025).

The cuts went further: Mediaweek reported on October 25, 2026 that the group was past most of a planned reduction of about 15,000 roles, from about 120,000 people to about 105,000. Omnicom says it is on track to save $900 million of costs from the combination in 2026, and that media planning and buying made up 53% of the revenue of its core businesses in the third quarter of 2026 (Omnicom, October 22, 2026).

The table ranks the groups by their 2025 revenue, as Ad Age's annual Agency Report counted it in June 2026, and gives each group's latest organic growth. Organic growth strips out acquisitions, disposals, and currency moves, so it shows whether the business that existed a year earlier is growing. The revenue figures are not quite comparable: some groups count costs they pass straight through to clients, such as production and travel bills, and others leave them out.

Group (base)	Revenue, 2025	Organic growth, Jul–Sep 2026	Note
Omnicom (New York)	$26.48 billion, counted with Interpublic	+7.2%	Largest by revenue since the merger
Accenture Song (part of Accenture)	$20.0 billion	Not reported separately	The marketing arm of a consulting firm
Publicis Groupe (Paris)	$19.66 billion	+5.0%	Largest by stock-market value (Reuters, February 3, 2026)
WPP (London)	$17.87 billion (−5.2%)	−2.1%	In a recovery plan, Elevate28 (7c)
Dentsu (Tokyo)	$9.60 billion	Not reported as of October 31, 2026	Dropped the sale of its international business (7c)
Havas (Paris)	$3.29 billion	+2.4%	Listed on its own since its split from Vivendi in December 2024

Revenue: Ad Age Agency Report (June 2026), as reported by Meio & Mensagem; Ad Age counts $158.8 billion for the 25 largest agency companies. Organic growth: company releases of October 13 (Publicis), October 14 (Havas), October 22 (Omnicom), and October 29, 2026 (WPP, whose measure is like-for-like revenue less pass-through costs).

The six held about 61% of the revenue of Ad Age's top 25 (calculated from the table). Their chief executives are John Wren at Omnicom, who has run it since 1997; Arthur Sadoun at Publicis; Cindy Rose at WPP, since September 1, 2025; Takeshi Sano at Dentsu, since March 27, 2026; and Yannick Bolloré at Havas. Below them sits S4 Capital, the digital group Martin Sorrell founded in 2018 after leaving WPP, with $995 million of revenue in 2025, down 8.2% (Ad Age).

7c. WPP's Fall and Dentsu's Retreat

WPP shows how fast a group can lose ground. Martin Sorrell built it through takeovers (Module 2) and ran it until 2018, and in 2025 it lost two large media accounts, Mars and Coca-Cola's North American media, to Publicis. Its like-for-like revenue less pass-through costs fell 5.9% in the third quarter of 2025, and Cindy Rose, a former Microsoft executive who had just taken over, said: “Our recent performance has not been acceptable” (as reported by Marketing Dive).

In February 2026 WPP set out a plan it calls Elevate28: stabilize the business in 2026, return to growth in 2027, and accelerate in 2028. It aims for £500 million of gross annual cost savings by 2028, £100 million of them in 2026, and it is selling businesses it no longer counts as core, with more than £200 million of proceeds expected in 2026 (WPP, August 6, 2026). The decline has narrowed each quarter, from 6.7% in the first quarter of 2026 to 2.8% in the second and the smaller fall in the table. WPP Media, its media-buying arm, grew 0.3% in the third quarter (WPP, October 29, 2026).

Investors have marked the group down: WPP's stock-market value was about $5.6 billion on October 29, 2026 (Investing.com), under a third of its 2025 revenue. Rose describes the aim as “a simpler, integrated

WPP – powered by WPP Open," the group's own technology platform (WPP's first-quarter update, April 2026).

Dentsu, Japan's largest agency group, has pulled back from its expansion abroad. Reports in August 2025 said it had hired Mitsubishi UFJ Morgan Stanley and Nomura to explore selling all or part of its international business, and it announced 3,400 job cuts the same month. By late December 2025 the private-equity firm Apollo and rival agency groups had dropped out, leaving Bain Capital as the only bidder, and Dentsu's shares fell 11% (Campaign).

On February 13, 2026, Dentsu abandoned the sale, booked ¥396.1 billion (about $2.6 billion) of goodwill impairments for 2025, writing down the premium it had paid for businesses abroad, and suspended its year-end dividend for the first time. It also named Takeshi Sano, the head of Dentsu Japan, to replace Hiroshi Igarashi as chief executive (PPC Land, February 13, 2026).

7d. Publicis and the Data Bet

Publicis has spent the past decade betting that data, not buying power alone, is what clients will pay for. In 2019 it bought Epsilon, a US marketing-data company, for $4.4 billion, and in March 2025 it added Lotame, which sells audience data and identity tools to publishers and marketers. On May 17, 2026 it agreed to buy LiveRamp for $38.50 a share in cash, about $2.5 billion, a premium of 30% over LiveRamp's last share price (LiveRamp's filing with the Securities and Exchange Commission).

LiveRamp's business is identity resolution: matching the records one company holds on a person to those another company holds, across devices, without either handing over names. It runs much of the plumbing behind data clean rooms, shared databases in which two companies can match their customer lists without seeing each other's, and behind what is replacing third-party cookies (Module 14). LiveRamp's shareholders approved the deal on August 17, 2026, with more than 99% of the votes cast in favor, and the companies expected to complete it by December 31, 2026, after antitrust and US national-security reviews (PPC Land, August 18, 2026).

The logic runs through the industry: as browsers and phones stop letting advertisers follow people from site to site, an agency that controls identity data can sell targeting and measurement that a pure buyer of media cannot. Accenture Song makes the mirror-image bet: a consultancy that already builds its clients' customer databases sells them the marketing that runs on those databases.

7e. How Agencies Are Paid

For most of the 20th century an agency kept 15% of the media it bought, whatever the ads achieved (Module 2 gives the commission's origin). The commission tied the agency's income to the size of the client's budget rather than to results, and it rewarded recommending more spending. Advertisers began to break with it early: a New York Times headline of 1960 read "Shell and Ogilvy Try a Fee Instead of 15%."

Fees now dominate: in the ANA's survey of agency compensation published December 13, 2022, 82% of marketers used a fee in at least one agency agreement, up from 68% in 2016. Most fees are built from labor: the agency estimates the hours each grade of staff will spend on the client, multiplies them by rates that cover salaries and overhead, and adds a profit margin. Among advertisers spending $500 million or more, 53% used fixed or output-based fees, paying per campaign or per piece of work, against 5% in 2016. For media agencies alone, 19% of marketers still paid by commission (MediaPost's report of the same study).

Payment for results has become less common, not more: performance incentives, bonuses tied to sales, brand measures, or agreed goals, were used by 41% of marketers in 2022, against 48% in 2016 and 61% in 2013 (ANA). The groups say this will turn: WPP's Cindy Rose has called a move to outcome-based pay “inevitable” (Campaign, August 6, 2026), though the surveys so far show advertisers moving toward fixed fees instead.

Fees put agencies' margins under steady pressure, because clients push rates down at every review and ask for longer payment terms. The ANA's 2024 report on principal media traced the practice's growth to exactly these pressures: shrinking agency margins, slower payment, and cost-cutting (as summarized by AdNews, July 22, 2025). The groups' answer has been to earn more on the media itself, in ways that do not appear on a client's invoice.

7f. Principal Media and Rebates

In principal-based buying, the agency stops acting only as its client's agent. Its trading desk buys television time or digital inventory on its own account, taking the risk that it goes unsold, then resells it to clients at a price it sets, without disclosing what it paid. The client sees a price per thousand impressions or viewers, not the agency's margin.

Agencies say the client gets a better deal: on a July 2025 earnings call Omnicom's John Wren said clients “opt into” principal media and “get a better deal,” while Omnicom earns “incremental revenue with an incremental margin” (AdNews, July 22, 2025). Advertisers are using it more and trusting it less, according to the ANA's survey of 114 marketers, published March 19, 2026: 58% had used principal media in the past year, up from 47% in its previous survey, mostly for television (74%) and open-web digital ads (43%). Yet 90% were unsure it served their interest, and only 57% had guidelines for it.

Rebates are older: a media owner pays the agency a share of the money it brings, in cash, in free inventory, or through “service agreements,” and the agency may keep it. On June 7, 2016 K2 Intelligence, an investigations firm hired by the ANA, published a report built on 143 interviews with 150 sources. It found “numerous non-transparent business practices, including cash rebates to media agencies” pervasive in US media buying, and said senior executives across the agency business had been aware of some of these practices and had mandated them.

The report also described principal buying with markups “between 30% and 90% in some cases” (Campaign's account), and the agency groups denied wrongdoing. Rules differ by country, and France's Sapin law of 1993 requires an agency to buy media as its client's agent and to pass any discount to the advertiser, which rules out rebates of the kind K2 found. In the United States the matter is left to the contract between advertiser and agency. The FTC's orders against the agency groups concern brand safety, not pay (Module 13).

7g. In-House Agencies and the AI Squeeze

The agencies' oldest rival is the client: in 2023, 82% of the ANA's members had an in-house agency, a team on the advertiser's own payroll that makes ads or buys media, up from 78% in 2018, 58% in 2013, and 42% in 2008. Most use it alongside outside agencies rather than instead of them, and in the ANA's 2026 study (June 22, 2026), 53% of marketers described their in-house team as a strategic partner and 9% mainly as a way to save money (Adweek).

Generative AI adds pressure because it attacks the work agencies bill by the hour: writing copy, making versions of images and video, and adapting ads to each market and format. In Gartner's 2025 survey of chief marketing officers (CMOs), 22% said generative AI had reduced their reliance on agencies, and 39% planned to cut agency budgets. The platforms are automating buying and creative too (Module 17), which offers advertisers a route that bypasses the agency altogether.

The groups have answered differently: Omnicom chose scale and cost cuts, Publicis chose data, WPP chose a simpler structure built on its own technology, and Dentsu chose retrenchment; each group's AI platform is described in Module 17. All of them are trying to earn more from technology and trading and less from hours. That shift raises the question the worked example puts in numbers: how much of an agency's income its client can see.

Worked example — What an Agency Earns on a $100 Million Media Account

Inputs: a client spends $100 million a year on media through one agency. The benchmark fee comes from the CMA's market study of July 2020, which found that media agencies took about 5% of advertisers' programmatic spending (range 4% to 14%, on 2019 UK data). The principal-media and rebate figures below are assumptions for illustration. Step 1, the old commission (7e) on $100 million: $15 million, all of it visible to the client. Step 2, a fee model at the CMA's average: 5% × $100 million = $5 million, shown on the invoice. Step 3, principal media: assume a quarter of the budget, $25 million, goes through the agency's trading desk, which bought that inventory for $20 million. The agency earns $25 million − $20 million = $5 million, a markup of 25% on its cost, below the range K2 reported (7f). Step 4, a rebate: assume media owners pay the agency 2% on $50 million of the remaining spending, $1 million, which the contract does not require it to return. Step 5, add: $5 million + $5 million + $1 million = $11 million, or 11% of the media budget. Step 6, what the client sees: the $5 million fee. The other $6 million, 55% of the agency's income, appears on no invoice line. Step 7, the agencies' side: if the same inventory would have cost the client $27 million bought directly, principal buying saved it $2 million even after the agency's margin. Result: the fee no longer tells a client what its agency earns. The answer depends on prices the client cannot see, which is why most of the ANA's respondents could not say whether principal media served them.

Module 7 Key Points

The Agencies and Media Buying

• Omnicom's purchase of Interpublic, completed November 26, 2025, made it the largest agency group, with $26.48 billion of 2025 revenue.

• WPP's revenue fell 5.2% in 2025, and its Elevate28 plan aims to return the group to growth in 2027.

• Dentsu dropped the sale of its international business in February 2026 after writing down ¥396.1 billion of goodwill.

• Fees have replaced the 15% commission: 82% of marketers in the ANA's 2022 survey used a fee.

• Principal media, agencies reselling inventory they bought for themselves, was used by 58% of marketers in the ANA's 2026 survey.

• In-house agencies, at 82% of ANA members in 2023, and generative AI are squeezing the work agencies bill by the hour.

Module 7 Quiz

Check Your Understanding

1. Which agency names did Omnicom retire after it bought Interpublic?

A) TBWA, BBDO and Omnicom Media

B) Ogilvy, J. Walter Thompson and Young & Rubicam, which had been merged into other networks years earlier

C) Leo Burnett and Saatchi & Saatchi

D) DDB, FCB and MullenLowe

2. What is principal-based media buying?

A) A media owner paying an agency a share of the money the agency brings it

B) An agency buying inventory on its own account and reselling it to clients at a price it sets, without disclosing what it paid

C) A client buying directly from publishers without any agency

D) An agency keeping 15% of every media bill it places

3. In the ANA's 2022 survey, what share of marketers used a fee in at least one agency agreement?

A) 82%, up from 68% in 2016

B) 19%

C) 41%, down from 61%

D) 53%

4. What did Dentsu do on February 13, 2026?

A) It agreed to sell its international business to Bain Capital, the only bidder left

B) It dropped the sale of its business abroad and wrote down goodwill

C) It agreed to buy the data company LiveRamp for about $2.5 billion in cash

D) It merged with the French group Havas to form the world's fourth-largest agency group

5. In the worked example, why does the client see only $5 million of the agency's $11 million income?

A) The fee is paid by the media owners, not by the client

B) The agency books most of its income in the following year, so this year's invoices show only the fee

C) Most of the income goes to the holding company as a dividend

D) The principal-media margin and rebate appear on no invoice

Module 8 — Search, Social, and Video: The Platforms

Google, Meta, TikTok, and the Rest

Recommended Reading / Reference

Book: Steven Levy, *Facebook: The Inside Story* (2020), a veteran technology writer's account, built on wide access to the company, of how Facebook turned a student network into the advertising machine that pays for it

Report: Alphabet, *Alphabet Announces Third Quarter 2026 Results* (October 28, 2026), and Meta Platforms, *Meta Reports Third Quarter 2026 Results* (October 21, 2026), the two companies' earnings releases, which give advertising revenue by line and, for Meta, the number of ads shown and their average price

8a. Who Sells the World's Advertising

Advertising is sold by thousands of companies, but most of the money goes to a few. WPP Media's forecast of June 16, 2026 found that five sellers, Google, Meta, ByteDance, Amazon, and Alibaba, would take 58% of the world's advertising revenue in 2026. Six of the ten largest sellers were Chinese, and for the first time no traditional media owner, no broadcaster or publisher, made the top ten.

Other forecasters measure the same concentration in other ways: WARC put the share of Alphabet, Amazon, and Meta at 59.7% of advertising spending outside China in 2026, or $659.6 billion, rising to 61.5% by 2028 (October 8, 2026). This module covers the platforms that sell search, social, and video advertising, while Amazon, whose ads sit inside a store, is the subject of Module 9, and the market's totals by country and channel are in Module 1.

The order at the top depends on how revenue is counted. eMarketer, a forecasting firm, expects Meta to pass Google in worldwide digital advertising revenue in 2026, with 26.8% of the total against Google's 26.4% (April 13, 2026). It counts Google's revenue after the payments Google makes to the partners that send it traffic, so its 2025 figure for Google is $214 billion, well below the gross total Alphabet reports (8b). Counted gross, as the companies report, Google remains the larger seller.

8b. Alphabet: Search, YouTube, and the Network

Google sells advertising in three lines, and Alphabet reports each. Search and other covers the ads beside search results on Google and on its other properties, such as Maps, Gmail, and the Play store. YouTube ads are the video ads on YouTube, and Google Network is the ads Google places on other companies' websites and apps through its AdSense, AdMob, and Ad Manager products, keeping a share and paying the rest to the sites' owners.

Line	2025	Change on 2024	Jul–Sep 2026	Change on a year earlier
Search & other	$224.5 billion	+13.4%	$63.35 billion	+12.0%
YouTube ads	$40.4 billion	+11.7%	$11.21 billion	+9.2%
Google Network	$29.8 billion	−1.9%	$7.46 billion	+1.4%
All Google advertising	$294.7 billion	+11.4%	$82.0 billion	+10.6%

Alphabet's release for the fourth quarter of 2025 (February 4, 2026), its quarterly report for the third quarter of 2025, and its release of October 28, 2026; 2025 totals built from the nine-month report and the fourth-quarter release. Alphabet's total revenue in the third quarter of 2026 was $127.6 billion, up 25%.

Advertising still pays for most of Alphabet, but a shrinking share of it: ads were 64% of its revenue in the third quarter of 2026 (calculated from the release), as cloud computing and subscriptions grew faster. Search, the core of the business, slowed: Philipp Schindler, Google's chief business officer, said on the October 28 call that its growth eased from the second quarter because of currency moves, a second quarter lifted by the World Cup, and Asian retailers hurt when the European Union ended its customs exemption for low-value parcels. How Google's AI answers carry ads, and how much they earn per search so far, is in Module 17.

Google pays for much of its traffic: traffic acquisition costs (TAC) are its payments to Network partners, the sites that show Google's ads, and to distribution partners, such as Apple, Samsung, browser makers, and phone carriers, that make Google the default search engine on their devices. They came to $16.4 billion in the third quarter of 2026, about 20% of Google's advertising revenue (calculated). How the courts have treated the default-search deals is in Module 16.

The Network is the weak line; the open web's advertising and the auction machinery behind it are the subject of Module 10. YouTube is the growth line after search, and its share of television viewing and its sales on TV screens are covered in Module 11.

8c. Meta: More Ads, and Higher Prices

Meta Platforms, the owner of Facebook, Instagram, Messenger, and WhatsApp, is almost entirely an advertising business: ads brought in $196.18 billion of its $200.97 billion of revenue in 2025, up 22%. It counted 3.64 billion daily active people across its apps in September 2026, up 3% on a year earlier, a count of people rather than accounts that Meta estimates itself.

Meta reports two drivers of its ad revenue, the number of ads it shows (impressions) and the average price per ad, and in 2025 impressions rose 12% and prices 9%. In the third quarter of 2026 the balance shifted toward price: advertising revenue rose 26% to $63.06 billion, with impressions up 10% and the average price per ad up 14% (1.10 × 1.14 = 1.25; the rest is rounding). Rising prices mean advertisers are bidding more for each slot, and Meta credits its AI ranking models with making each ad perform better (Module 17).

North America brings in the most money. Of Meta's $64.60 billion of revenue in the quarter, nearly all of it advertising, users in the United States and Canada accounted for $26.2 billion (up 33%), Europe $14.0 billion (up 21%), Asia-Pacific $17.0 billion (up 19%), and the rest of the world $7.4 billion (up 30%), by the user's location (Meta's quarterly report, October 2026). Meta expects fourth-quarter revenue of $73 billion to $76 billion, and capital spending of $135 billion to $145 billion in 2026, mostly on data centers and servers (October 21, 2026); its ads made with creators are covered in Module 12.

8d. ByteDance, TikTok, and the American Joint Venture

ByteDance, the Beijing-founded owner of TikTok and of its Chinese counterpart, Douyin, is private and publishes no accounts, so its size comes from press reports. Bloomberg reported in December 2025 that it was aiming for about $186 billion of revenue in 2025. The Information reported on September 15, 2026 that it had in fact taken in about $200 billion in 2025, up 29%, and about $120 billion in the first half of 2026. The Chinese outlet 36Kr put TikTok's share of ByteDance's revenue above 30%, and the company does not disclose how much of the total is advertising.

TikTok's American business changed hands on January 22, 2026, after a 2024 law required ByteDance to sell it or see the app banned; the Supreme Court had upheld the law in *TikTok v. Garland* (January 17, 2025). The deal, announced on December 18, 2025, created a US joint venture owned 15% each by Oracle, the private-equity firm Silver Lake, and MGX, an Abu Dhabi investment firm; 5% by other new investors; 30.1% by affiliates of existing ByteDance investors; and 19.9% by ByteDance, the largest single stake. A seven-member board with an American majority, which includes TikTok's chief executive, Shou Chew, oversees it, and Adam Presser runs it.

The venture holds less than its name suggests: ByteDance keeps the recommendation algorithm and licenses it to the venture, which retrains an American copy on American users' data in Oracle's cloud. The advertising, e-commerce, and marketing business stays with TikTok companies that ByteDance controls, not with the venture (EqualOcean and Variety, reporting on the deal's terms). WARC estimates that TikTok will sell $14.5 billion of US advertising in 2026, 38% of its ad revenue worldwide (which implies about $38 billion), and $16.7 billion in 2027 (January 28, 2026); eMarketer's estimate for 2026 is over $17 billion.

8e. Microsoft, Apple, and the Smaller Networks

Below the leaders the market is crowded, and its members are growing at very different speeds. Microsoft sells search ads on Bing and, through LinkedIn, ads aimed at professionals; eMarketer gives it 2.1% of the world's digital ad revenue in 2026. Apple sells ads in its App Store's search results, and it announced in January 2026 that from March 3 it would show several sponsored apps there instead of one, starting in the United Kingdom and Japan and reaching all markets by the end of March (MacDailyNews, January 23, 2026). eMarketer gives Apple 1.6%. John Ternus succeeded Tim Cook as Apple's chief executive on September 1, 2026, and Apple's rules on tracking in other companies' apps are covered in Module 14.

Company	What it sells	Latest figure	Change	Period and source
Microsoft	Search ads on Bing; LinkedIn ads	Search and advertising revenue, excluding traffic acquisition costs	+12%	Jul–Sep 2026; Microsoft, October 28, 2026
Pinterest	Ads among saved images and shopping ideas	Revenue, $1.174 billion	+12%	Jul–Sep 2026; Pinterest, October 29, 2026
Snap	Ads in Snapchat	Revenue, $1.599 billion	+19%	Apr–Jun 2026; Snap, August 3, 2026
Reddit	Ads in discussion forums	Advertising revenue, $847 million	+54%	Jul–Sep 2026; Reddit, October 29, 2026
X	Ads on the former Twitter	Ad revenue, $2.26 billion (estimate)	About half of Twitter's 2021 level	2025; eMarketer, as reported in April 2026

Company releases and eMarketer. Pinterest and Snap report revenue without splitting out advertising. Snap had not reported its July–September results as of October 31, 2026. eMarketer expects X's ad revenue to reach $2.46 billion in 2026.

Reddit is growing fastest, and its daily active visitors rose 12% to 130.3 million (Reddit, October 29, 2026). Pinterest reached 661 million monthly users, up 10%, but 76% of its revenue came from the United States and Canada (calculated from its October 29 release). Snap's daily users rose to 493 million in the second quarter, while those in North America fell 7% to 92 million (Snap, August 3, 2026). Elon Musk bought Twitter, now X, in October 2022; the advertisers' withdrawal and the lawsuits that followed are covered in Module 13.

8f. China's Own Giants

China's advertising market, the world's second-largest (Module 1), is served almost entirely by Chinese companies: Facebook has been blocked there since 2009, and Google withdrew its search engine in 2010. ByteDance, which sells Douyin's ads, is private (8d), but the listed sellers report every quarter, in renminbi (RMB), China's currency, and their results for April to June 2026 show a market split between sellers that are growing and sellers that are not.

Company	Advertising line	Apr–Jun 2026	Change	Reported
Tencent (WeChat)	Marketing services	RMB 43.6 billion	+22%	August 12, 2026
Alibaba (Taobao, Tmall)	Customer management revenue	RMB 82.5 billion	−7% (+1% excluding an accounting change)	August 20, 2026
PDD (Pinduoduo, Temu)	Online marketing services and others	RMB 57.6 billion	+3%	August 24, 2026
Kuaishou (short video)	Online marketing services	RMB 20.6 billion	+4.4%	August 19, 2026
Baidu (search)	Online marketing services	RMB 13.1 billion	−19%	August 18, 2026

Company releases (Kuaishou's as reported by Webull). Alibaba's like-for-like figure excludes subsidies to merchants that it now records as a reduction of revenue rather than as a marketing expense. Baidu reports a further RMB 2.6 billion of “AI-native marketing services” inside its AI business.

The split follows the kind of seller: Tencent, which sells ads in its WeChat messaging app, grew fastest. The marketplaces, Alibaba and PDD, earn their advertising from merchants paying for placement in shopping searches, the model Module 9 examines as retail media's precedent. Baidu, China's main search engine, is shrinking. None of these companies had reported results for July to September as of October 31, 2026.

8g. How the Platforms Sell

The platforms sell almost everything through self-serve auctions: an advertiser, from a sole trader to a global brand, opens an account, chooses an objective such as clicks, app installs, leads, or sales, sets a budget, and uploads the ads. Each time a page or feed loads, an auction decides which ad appears and what it costs; the mechanics are explained in Module 10, and millions of small businesses buy this way (Module 6).

Large advertisers get more help: account teams at the platforms, agencies that run their buying through the platforms' tools, and ways to connect their own customer data. The auction is the same, and so is the platforms' hold on measurement, since the results an advertiser sees come mostly from the platform that sold the ads (Module 13).

Increasingly the advertiser hands the decisions to the platform: Google's Performance Max runs one campaign across Search, YouTube, Display, Gmail, Maps, and Discover, and Meta offers Advantage+. What these automated campaigns decide, and how far the platforms' AI now goes, are the subjects of Module 17.

Prices depend on the objective the advertiser picks: one that asks Meta for leads pays more per thousand impressions than one that asks for clicks, because Meta shows its ads to people it predicts will act. Affect Group, an agency, put typical US Meta CPMs at $10 to $15 for traffic campaigns and $30 to $45 for website lead campaigns in November 2025. The worked example uses the benchmarks of WordStream by LocaliQ, a marketing-software firm whose 2026 figures (published September 23, 2026) come mostly from small and midsize businesses' campaigns.

Worked example — A Small Business Buys Leads on Meta

Inputs: a home-services firm sets a budget of $1,500 a month for a Meta lead campaign. LocaliQ's 2026 benchmarks for lead campaigns are a cost per click of $1.80, a click-through rate of 2.70%, and 8.54% of clicks turning into leads; for traffic campaigns, $0.60 per click at a 1.93% click-through rate. Step 1, clicks: $1,500 ÷ $1.80 = 833. Step 2, impressions: 833 ÷ 0.027 = about 30,900. Step 3, the implied CPM, the cost per click times the click-through rate: $1.80 × 2.70% × 1,000 = $48.60. Step 4, leads: 833 × 8.54% = 71. Step 5, cost per lead, the cost per click divided by the share of clicks that become leads: $1.80 ÷ 8.54% = $21.08. LocaliQ's own average cost per lead is higher, $27.39, because it averages twenty industries' figures rather than dividing one average by another. Step 6, the traffic campaign's CPM: $0.60 × 1.93% × 1,000 = $11.58, so the lead campaign pays about four times as much per thousand impressions. Result: the firm buys about 71 leads a month, and if one in five becomes a customer, an assumption, each new customer costs about $105 in advertising.

Module 8 Key Points

Search, Social, and Video: The Platforms

• Five sellers, Google, Meta, ByteDance, Amazon, and Alibaba, take about 58% of the world's advertising revenue (WPP Media, June 2026).

• Alphabet's ad revenue reached $82.0 billion in the third quarter of 2026; about a fifth went to partners as traffic acquisition costs.

• Meta's ad revenue rose 26% in the third quarter of 2026, driven more by higher prices per ad than by more ads.

• TikTok's US joint venture closed on January 22, 2026, but TikTok's American advertising stays with ByteDance-controlled companies.

• Growth is uneven: Reddit's ad revenue rose 54% in the third quarter of 2026, while China's search leader, Baidu, shrank.

• At 2026 benchmark prices, a small Meta lead campaign pays about $49 per thousand impressions and $21 per lead.

Module 8 Quiz

Check Your Understanding

1. What are Google's traffic acquisition costs?

A) What advertisers pay Google for each click on a search ad
B) The cost of the data centers that serve Google's ads, which has grown as it added AI answers to search
C) Payments to sites showing its ads and to partners like Apple that make it the default search engine
D) Fines Google pays to the European Union

2. What drove most of Meta's 26% rise in ad revenue in the third quarter of 2026?

A) A 26% rise in daily active people
B) A 14% rise in the average price per ad, with impressions up 10%
C) New ads in WhatsApp chats, which doubled the number of impressions
D) Advertisers moving budgets from Google after the search ruling

3. Under the TikTok deal that closed on January 22, 2026, who runs TikTok's US advertising business?

A) TikTok companies controlled by ByteDance, which also keeps the algorithm
B) The joint venture that Oracle, Silver Lake, and MGX co-own
C) Oracle, which also hosts the American users' data
D) An independent trustee chosen by the US government

4. Why does eMarketer show Meta passing Google in 2026 while Alphabet still reports more ad revenue than Meta?

A) eMarketer leaves YouTube's ad revenue out of Google's total
B) Alphabet's fiscal year ends in June, so its totals cover different months
C) eMarketer counts only revenue earned in the United States
D) eMarketer counts Google's revenue after partner payments

5. In the worked example, why does the lead campaign pay about four times the traffic campaign's price per thousand impressions?

A) Lead campaigns can run only on Instagram, where prices per thousand impressions are higher
B) Meta adds a fixed fee to every lead form
C) Meta shows lead ads to people it predicts will act, and those impressions cost more
D) Traffic campaigns are charged only per click, so their impressions cost the advertiser nothing

Module 9 — The Stores Become Media: Retail Media

Amazon, Walmart, and the Ads at the Point of Sale

Recommended Reading / Reference

Book: Dana Mattioli, *The Everything War: Amazon's Ruthless Quest to Own the World and Remake Corporate Power* (2024), a *Wall Street Journal* reporter's account of how Amazon deals with the independent sellers on its marketplace and with its rivals

Report: Amazon.com, *Annual Report on Form 10-K* for 2025 (February 2026), which reports “advertising services” as a line of revenue and says what it includes; Amazon's quarterly earnings releases update the figure

9a. What Retail Media Is

Retail media is advertising sold by a company that sells things, using what it knows about what its customers buy. Its main form is the sponsored listing: a product placed at or near the top of a store's search results or category pages, marked as an ad and usually paid for per click. Around it sit display and video ads on the retailer's site and app, ads on screens in its stores, and off-site ads, which use the retailer's shopper data to target people on other websites, on social media, and on streaming television.

What retailers sell above all is closed-loop measurement: because the same company records who saw an ad and who later bought, it can tell a supplier how many sales followed a campaign, product by product, which a television network or a social app can only estimate. The claim has a limit: sales that follow an ad are not all caused by it, since many shoppers would have bought anyway (Module 5 on incrementality, Module 13 on attribution).

The business has spread beyond stores to what the industry calls commerce media: delivery apps, ride-hailing, travel sites, payment firms, and banks, all of which see what their customers buy. WARC expects retail media to bring in $202.1 billion worldwide in 2026, up 14.3% (October 8, 2026); how it overtook television is told in Module 3. For the United States, eMarketer forecast $58.79 billion in 2025 and $69.33 billion in 2026, with Amazon and Walmart taking $9.42 billion of the $10.53 billion added, about 89%.

9b. Amazon: The Store That Became an Ad Giant

eMarketer ranks Amazon third in the world's digital advertising in 2026, behind Meta and Google, with about 9% of the total (April 13, 2026), and nearly all of it grew from the sponsored listing. Amazon Ads is run by Paul Kotas, who started its cost-per-click ad program (his speaker profile for the CES trade show, January 2026). It now sells sponsored listings to sellers and brands, display ads on and off Amazon, and video ads on Prime Video (Module 11), Twitch, and Fire TV. It also runs a demand-side platform that buys ads across other websites and streaming services (Module 10).

Quarter	2025	2026	Change
January–March	$13.9 billion	$17.24 billion	+24%
April–June	$15.7 billion	$19.81 billion	+26%
July–September	$17.7 billion	$19.64 billion	+11% (12% excluding currency moves)

Quarter	2025	2026	Change
October–December	$21.3 billion	Not yet reported	
Full year	$68.6 billion, up 22% from $56.2 billion		

Amazon's advertising services revenue. Amazon's releases, as reported by Marketing Dive and in Amazon's call of October 29, 2026; changes calculated.

The third quarter's slowdown was mostly a matter of the calendar: Amazon moved its Prime Day sale into the second quarter in most countries in 2026, and Andy Jassy, its chief executive, said the shift took 8.5 percentage points off third-quarter advertising growth. He expected growth to return to the high teens in the fourth quarter (October 29, 2026). Advertising was about 10% of Amazon's $201.9 billion of net sales in the third quarter (calculated), and Amazon does not report the profit it earns on it.

Sponsored listings are sold by auction, mostly per click. Eva, an Amazon marketing firm, found a median cost per click of $1.07 across Amazon's ad types in 150 US brand accounts between July 27 and October 24, 2026, a period that included the October Prime Big Deal Days sale, ranging from $0.80 in beauty to $1.28 in home and kitchen (October 26, 2026). Prices have been rising: Triple Whale, which sells analytics software to online brands, measured Amazon's average cost per thousand impressions at $7.82 in 2025, up 47% from $5.30.

The sponsored listings sit where sellers most need to be seen, in Amazon's own search results, so for a seller they are both advertising and a cost of doing business on the site. Amazon's advertising revenue in 2025 came to about 40% of the $172 billion it collected in “third-party seller services,” the commissions and fulfillment fees sellers pay (Modern Retail, April 23, 2026; calculated). Marketplace Pulse counted about 500,000 active sellers on Amazon.com in March 2026, down from about 584,000 in January 2025, a count Amazon disputes. In its 2026 Seller Index, 46% of sellers named advertising costs among their biggest worries about margins, and nearly half named Amazon's fees (Modern Retail, April 23, 2026). How Amazon prices its sponsored-ad auctions is the subject of an FTC lawsuit covered in Module 10.

9c. Walmart and the Followers

Walmart is the other company eMarketer expects to take most of the growth, a distant second to Amazon. Its global advertising revenue rose 27% to $4.4 billion in the fiscal year to January 2025. In December 2024 it bought Vizio, the television maker, for $2.3 billion, which gave it the ads on Vizio's TV screens and a way to sell advertising to brands it does not stock. With Vizio, advertising reached about $6.4 billion in the fiscal year to January 2026, up 46% (PPC Land).

Growth has stayed fast: in the quarter to July 31, 2026, Walmart's global advertising grew 38% including Vizio; Walmart Connect, its US retail media business, grew 43% without Vizio; and international advertising, led by Flipkart in India, grew 20% (Walmart, August 20, 2026, as reported by PPC Land). On June 23, 2026 it announced the purchase of Vibe, a self-serve platform for streaming-TV ads with more than 10,000 advertisers (PPC Land).

The followers are smaller but share the pattern: each owns a moment of purchase, a list of logged-in customers, and a record of what they bought, and each sells suppliers a place in that moment. The table gives each company's latest public figure; a run rate is a recent period's revenue scaled up to a full year.

Company	Ad business	Latest figure	Source and date
Target	Roundel	$915 million of ad revenue in 2025, a year in which store sales fell	Adweek, March 3, 2026

Company	Ad business	Latest figure	Source and date
Uber	Ads in the ride and Uber Eats apps	Passed a $1.5 billion annual run rate, growing more than 60% a year	Dara Khosrowshahi, chief executive, May 2025
DoorDash	Ads in the delivery app	Passed $1 billion a year; about 150,000 advertisers	Shopifreaks, June 12, 2025
Instacart	Ads in the grocery app	“Advertising and other” revenue of $297 million in April–June 2026, up 16%, equal to 2.9% of the value of orders	Instacart, August 6, 2026

Latest figures found as of October 31, 2026. Uber's release for April–June 2026 (August 5, 2026) gave no advertising figure.

Target's advertising grew 55.3% to $295 million in the last quarter of 2025 (Adweek), and DoorDash bought Symbiosys, an advertising-technology firm, for $175 million in June 2025 (Shopifreaks). Grocers, pharmacies, banks, and airlines have built ad businesses on the same model.

9d. Why Retailers Want It

Retailers earn thin margins on the goods they sell and much wider ones on advertising. A sponsored listing uses pages and screens the retailer already has, shown to customers who are already there, so most of what a supplier pays for it is profit. Walmart's finance chief, John David Rainey, has made the effect plain: in February 2026 he was reported as saying that “fully a third of our profit in the most recent quarter was related to advertising and membership income.”

On August 20, 2026 Rainey said that almost half of Walmart's profit growth came from membership, advertising, and its online marketplace, and that incremental margins could “even go higher” as advertising grew (PPC Land). Amazon gives no comparable figure, but it is the same arithmetic on a far larger scale.

Data is the second reason: a retailer knows which households buy which brands, how often, and at what price, and it can sell suppliers both access to those households and a report of what they bought afterward. Off its own site, the same data lets it sell audiences on other companies' screens, which is why Walmart wanted Vizio's televisions. And in its own store, a retailer decides which products shoppers see first, so advertising has become one of the ways suppliers win that position. The data has gained value as browsers and phones limit tracking from site to site (Module 14), because a retailer's records come from its own logged-in customers and do not depend on third-party cookies.

9e. The Suppliers' Complaint

For the brands that pay for it, retail media often arrives as a new charge on an old relationship. Packaged-goods makers already pay retailers heavily in trade spending, the discounts, allowances, and fees for shelf space and promotions (Module 1), and retail media money tends to come out of the same budgets. WARC reported in October 2023 that shopper and trade marketing budgets were the most common source of funding for commerce media, and that only about one advertiser in ten had a budget set aside for it.

Small suppliers say the advertising is not optional. In confidential submissions to the Australian Competition and Consumer Commission's inquiry into supermarkets, reported by the Australian trade site B&T on September 30, 2024, smaller suppliers said they felt pressed to advertise on Woolworths' Cartology and Coles 360 to keep their commercial relationships, and one described “a new charge for advertising through retail media.” The commission had formed no view at that stage, and neither retailer commented in that report.

Measurement is the second complaint: the retailer sells the ads, runs the auction, and reports the results, and its count of sales that followed an ad includes purchases shoppers would have made anyway; how the industry checks such counts is the subject of Module 13. Bradley Keefer of Keen Decision Systems, which tracks brands' marketing spending, argues that the return on ad spending retailers report ignores “execution costs and margin,” and so flatters the ads (as reported by MediaPost in 2026). The retailers' answer is the closed loop itself: they argue that they can show what an advertised product sold, which trade promotions and television rarely could, and that suppliers keep buying because the ads reach shoppers at the moment of choice.

9f. China's Precedent

China showed earlier how far the model can go: two of its largest online marketplaces earn most of their revenue from merchants paying to be seen and to sell. In the quarter to June 2026, customer management revenue, the fees merchants pay for marketing services such as search placement plus commissions on sales, made up most of Alibaba's China e-commerce revenue. PDD's online marketing services were 51% of its revenue (calculated from their releases of August 20 and 24, 2026; Module 8 gives both lines' growth).

On those sites, and increasingly on Amazon, a search for a product returns a ranked list in which payment is one of the inputs, and the price of being seen rises as more sellers compete for the same searches. The result is a rising share of each sale going to the platform rather than the seller. Marketplace Pulse estimated that Amazon's average cut of a seller's sale, counting commissions, storage, packing and delivery, and advertising, first passed 50% in 2022 (Bloomberg, February 13, 2023), and the worked example measures the cut for one item in 2026.

Worked example — An Amazon Seller's $30 Item

Inputs: a seller lists a $30 kitchen gadget, stored and shipped by Amazon's fulfillment service, that weighs just under a pound in its box. Amazon's referral fee is 15% in most categories, and its 2026 fulfillment fee for a large standard-size item of 12 to 16 ounces priced $10 to $50 is $4.60 (Amazon's rate card, as summarized by AMZ Prep, October 2026). A 3.5% fuel and logistics surcharge has applied to fulfillment fees since April 17, 2026 (Retail Dive). The advertising benchmarks come from the same Eva study (9b), which found a median advertising cost of sale (ACoS, ad spending divided by the sales the ads produced) of 27.7% for Sponsored Products, and total ad spending equal to 10.6% of all sales. The seller's cost of making the item and shipping it to Amazon, $7.50, is an assumption. Step 1, referral fee: 15% × $30 = $4.50. Step 2, fulfillment: $4.60 × 1.035 = $4.76; the surcharge adds 16 cents, close to the 17-cent average per unit that Retail Dive reported. Step 3, Amazon's fees before ads: $9.26, or 30.9% of the price. Step 4, ads averaged over all sales: 10.6% × $30 = $3.18, so Amazon takes $12.44, or 41.5%. Step 5, a sale that came through an ad at the median ACoS: 27.7% × $30 = $8.31, so Amazon takes $17.57, or 58.6%. Step 6, what the seller keeps after its $7.50 cost: $10.06 on the average sale and $4.93 on the ad-driven one, before storage, returns, and overhead. Result: Amazon's share runs from about 41% to 59% of the price, around the Marketplace Pulse estimate above.

Module 9 Key Points

The Stores Become Media: Retail Media

• Retail media, ads sold by stores using their shoppers' purchase data, will bring in $202.1 billion worldwide in 2026 (WARC).

• Amazon sold $68.6 billion of advertising in 2025; its third-quarter 2026 slowdown came mostly from moving Prime Day.

• Walmart's advertising reached about $6.4 billion in the fiscal year to January 2026, helped by buying the TV maker Vizio.

• Retailers want ad income for its margins: Walmart tied about a third of one quarter's profit to advertising and membership.

• Suppliers say retail media is trade spending in a new form, and small ones say it is not optional.

• On a $30 Amazon sale, fees and advertising take about 41% to 59%, around Marketplace Pulse's “roughly half.”

Module 9 Quiz

Check Your Understanding

1. What is closed-loop measurement in retail media?

A) The retailer links who saw an ad to who bought, using its own records
B) An auditor checks a retailer's ad counts against a panel of households
C) Ads are shown only to shoppers who are already inside the store, on screens near the shelves
D) Suppliers pay only when an ad's sales exceed its cost

2. By Amazon's own account, why did its advertising growth slow to 11% in the third quarter of 2026?

A) The FTC ordered it to suspend sponsored ads while its lawsuit proceeds
B) Advertisers moved budgets to Walmart after it bought Vizio
C) Prime Video dropped the ads it had shown to Prime members since 2024
D) Prime Day moved to the second quarter in most countries

3. What did Walmart's finance chief say in February 2026 about advertising and membership income?

A) It had stopped growing once Vizio was counted
B) Fully a third of profit in the most recent quarter was related to it
C) It would be moved into a separately listed company
D) It was still too small to report separately, at under 1% of Walmart's profit

4. What do suppliers mean when they call retail media trade spending in a new form?

A) Retailers record it as a discount on the supplier's invoice
B) It must be bought through the supplier's trade association
C) It is paid from the same budgets as discounts and shelf fees, to keep a place with the retailer
D) Shoppers pay for it through a charge added at the till

5. In the worked example, what share of a $30 sale that came through an ad goes to Amazon in fees and advertising?

A) About 59%, or $17.57
B) About 15%, the referral fee
C) About 31%, before any advertising
D) About 41%

Module 10 — How Ad Tech Works: The Auction

From the Bid to the Screen, and Where the Dollar Goes

Recommended Reading / Reference

Book: Ari Paparo, *Yield: How Google Bought, Built, and Bullied Its Way to Advertising Dominance* (2025), a history of ad servers, exchanges, and header bidding by an ad tech veteran who worked at DoubleClick and Google before founding a company of his own

Report: ANA, *Programmatic Media Supply Chain Transparency Study* (complete report, December 2023), which followed advertisers' open-web spending from the buying software to the impression, and the quarterly benchmark the ANA has published since

10a. One Impression at a Time

Almost every digital ad is sold on its own. When a web page, an app, a search result, or a streaming show has room for an ad, that space is offered for sale and usually goes to the winner of an automated auction that ends before the space is filled. The buyer is not buying a newspaper page or a minute of airtime but an impression: one ad, shown once, on one screen, to a person the bidders may or may not be able to identify.

Three kinds of auction do most of the work. Search engines auction the ads beside their results, and social platforms auction the slots in their feeds. In both, the company that owns the space writes the rules, runs the auction, and keeps the money. The third kind, programmatic buying, sells space on millions of other websites, apps, and streaming services through a chain of separate companies, each taking a fee. Module 1 gives the US total bought this way.

The chain runs as follows. The advertiser, often through its agency, sets goals and budgets in a demand-side platform (DSP), software that bids on its behalf, and the DSP may buy data about audiences from data sellers. Bids go to an ad exchange or supply-side platform (SSP), which offers publishers' space to many DSPs at once. The publisher's ad server, the software that decides which ad fills each slot, compares the exchanges' best offers with deals the publisher has sold directly. The advertiser's own ad server records each delivery, so that each side can check the other's count.

10b. The Search Auction: Bids Times Quality

The first search-ad auctions, run by GoTo.com (Module 3) and later by Overture, as it was renamed, ranked advertisers by bid alone and charged each its full bid for every click. Under that rule, bidders' software raised bids a cent at a time to take the top slot until it was no longer worth the price, then dropped back, and the cycle began again; economists documented the pattern in Overture's data (Edelman and Ostrovsky, 2007).

Google changed both the ranking and the price in AdWords. It ranked ads by the bid multiplied by the predicted chance of a click, so that a relevant ad could beat a richer one, and it charged each winner only what it needed to stay ahead of the ad below. Overture adopted the same pricing. Economists call the design the generalized second-price (GSP) auction, and in 2007 two studies showed that it has a catch. In the single-item auction designed by the economist William Vickrey, bidding one's true value is always the best strategy; under GSP, a bidder can sometimes gain by bidding less (Edelman, Ostrovsky, and Schwarz; Varian).

Google's own description is now more elaborate. It computes an Ad Rank for each ad from six things: the bid, the quality of the ad and its landing page, minimum thresholds, the competitiveness of the auction, the context of the search, and the expected effect of extras such as links and phone numbers. An advertiser pays, Google says, only what is “minimally required to clear the Ad Rank thresholds” and “beat the Ad Rank of the competitor immediately below you”, or a reserve price if no ad is below. The thresholds are reserve prices that Google sets and adjusts. Its help pages also say that an ad winning by a wider margin may pay more per click, which a pure second-price rule would not do, and how Google tuned such settings became evidence in the search monopoly case (Module 16).

Microsoft runs the same kind of auction for Bing's ads, and Amazon for the sponsored listings in its store, where the bidders are mostly the sellers and brands of Module 9. How Amazon set its prices is the subject of the FTC's suit in section 10h.

10c. The Social Auction: Bidding on Predicted Actions

A feed has no query to match, so a social platform must predict who will respond. Meta says the ad that wins its auction is the one with the highest “total value”, which combines three inputs: the advertiser's bid, the estimated action rate, and ad quality. The estimated action rate is a machine-learning prediction of how likely this particular person is to do what the advertiser is paying for, such as click, install an app, or buy. Ad quality is judged from signals such as people hiding the ad. Meta's help pages have long written the combination as bid times estimated action rate plus quality, and they say the highest bid does not always win.

The arithmetic rewards relevance. An advertiser offering $10 for a purchase that the system expects one viewer in 200 to make is worth $0.05 an impression to Meta. One offering $20 for a purchase it expects from one viewer in 500 is worth $0.04, and loses despite bidding twice as much. Other social platforms sell their feeds through auctions built on the same idea.

Most advertisers do not bid impression by impression. They set a budget, a goal, and sometimes a target cost per result, and Meta's software bids for them in each auction and paces the budget through the day; its automated campaigns go further and choose the audience and the ad as well. The prediction is what Meta sells. The more a platform knows about its logged-in users, the better it can estimate each person's chance of acting, and the more an advertiser will pay for that impression, which is one reason the largest logged-in platforms take most of the money (Module 8).

10d. The Programmatic Chain: From Page Load to Ad

Programmatic auctions run on a shared technical language, OpenRTB, the standard for real-time bidding that the IAB Tech Lab says began in 2010. Version 2.5, released for comment in November 2016, added a signal for header bidding. When a page or app with ad space loads, an exchange sends a bid request to the DSPs it works with. The request describes the opportunity: the site or app, the slot's size and position, the device and its rough location, any user identifier the exchange holds, and a deadline.

The deadline is a fraction of a second. In that time each DSP must look up what it knows about the user, decide whether any of its advertisers wants this impression, choose a price, and reply with the price per thousand and the ad. The exchange picks a winner, the publisher's ad server confirms it, and the browser loads the ad, usually in well under a second.

Until the mid-2010s most publishers sold their space through a waterfall: their ad server, usually Google's DoubleClick for Publishers (DFP), offered each impression to one exchange or ad network after another, in order of their past average prices, and the first to meet a set floor took it. Google's own exchange, AdX, was connected to DFP and could bid in real time against the others' fixed prices. Rivals saw that position as an advantage, known as “last look”, which became part of the Justice Department's case against Google (Module 16).

Header bidding, which spread in 2014 and 2015, was the publishers' answer. A few lines of code in the page's header ask several exchanges for bids at once, before the ad server is called, and pass the best offer to it to compete with AdX. AppNexus released Prebid.js, a free version of that code, in 2015, and Prebid.org, an independent body founded in 2017, now maintains it; its guidance is to close the whole auction within 1,000 milliseconds. A Google email of June 2015, later a trial exhibit, called the header bidders “basically remarketers, trading desks and other exchanges who want to break DFP”. Google answered in 2018 with a server-side version of its own, Open Bidding.

Header bidding also ended the second-price auction in display advertising. When each exchange ran its own second-price auction and the exchanges' winners then competed on price in the page, no bidder could be sure what rule set its price. Exchanges moved to first-price auctions, in which the winner pays what it bid, in the late 2010s; Google announced the change for its exchange in March 2019. Buyers responded with bid shading: their software estimates the lowest price likely to win and bids just above it, so a first-price bid no longer reveals what the buyer thinks the impression is worth.

10e. Identity: How a Bidder Knows Who Is There

A bidder pays more for an impression when it can recognize the person. On the web the oldest tool is the third-party cookie, a small file set by an ad tech company's domain while someone visits a publisher's site, which lets that company recognize the same browser on other sites. Because every company sets its own cookie, DSPs and exchanges must match one company's identifier (ID) to another's, a process called cookie syncing. Each failed match leaves impressions that a buyer cannot recognize and will bid less for.

In apps the equivalent is the mobile advertising ID that Apple and Google assign to each phone, which an iPhone app must now ask permission to use. On television sets, buyers rely on the household's internet address and on data that smart TVs collect about what is on the screen. Where no identifier is available, a bid rests on the page's content and the device's rough location.

The strongest identifier is a login. Google, Meta, and Amazon know which of their users is signed in, run their auctions on that knowledge inside their own systems, and let little of the data out; the industry calls them walled gardens, and outsiders buy their space through the platforms' own tools. Outside the gardens, firms have built shared IDs from logins. The Trade Desk's Unified ID 2.0 turns an email address that a person gives a publisher into an encrypted token that bidders can recognize, and identity companies such as LiveRamp match records across firms. How much each signal is worth, and what the law allows, are Module 14's subjects.

10f. Where the Advertiser's Dollar Goes

Every link in the programmatic chain takes a fee, and for years no one could say how much of an advertiser's money reached the publisher. ISBA, the British advertisers' association, and the accounting firm PwC published the first study to follow impressions from advertiser to publisher, and called the money it

could not trace to any party the “unknown delta”. ISBA cautioned that its second study's small sample was not representative.

Study	Sample	What it counts	Finding
ISBA and PwC, May 2020	15 advertisers, 12 publishers	Share of spending that reaches publishers	51%; 15% could not be traced; 12% of impressions matched end to end
UK Competition and Markets Authority, July 2020	UK market, 2019 data	Share of spending that reaches publishers	About 65%; intermediaries kept about 35%
ISBA and PwC, January 2023	11 advertisers, 1.3 billion impressions	Share of spending that reaches publishers	65%; 3% could not be traced; 58% of impressions matched
ANA, December 2023	Members' open-web campaigns	Share of DSP spending that buys a measurable, viewable, fraud-free impression on a site not made for advertising	36 cents of each dollar
ANA with TAG, an industry anti-fraud body, August 6, 2026	Second quarter of 2026	The same share, now called TrueAdSpend	45.1%; transaction costs took 27.2%

ISBA/PwC, May 7, 2020 and January 18, 2023; CMA, online platforms and digital advertising market study, Appendix R (July 2020); ANA, complete report (December 2023); ANA benchmark for the second quarter of 2026 (August 6, 2026). The ANA's samples and definitions have changed from quarter to quarter, so its figures are not a clean series.

The ANA, the American advertisers' association, asked a different question in 2023: not how much reached the publisher, but how much bought an impression worth having. It started from the money entering DSPs and subtracted fees and spending on impressions that were fraudulent, could not be measured or seen, or ran on made-for-advertising (MFA) sites, pages built to buy cheap traffic and fill it with ads. In its first look, published June 19, 2023, MFA sites took 21% of impressions and 15% of spending, and the average campaign ran on 44,000 websites. Its complete report estimated $22 billion of possible savings for buyers.

The two kinds of measure answer different questions. A publisher can receive most of the price of an impression that no one sees, on a page no reader chose, and the ANA counts that money as lost while ISBA counts it as paid. Since 2024 the ANA has published this benchmark quarterly with TAG TrustNet; its report for the third quarter of 2026 is due on November 18, 2026. Buyers have also cut the number of sites they use and bought through shorter routes, a practice called supply-path optimization, connecting directly to publishers or to fewer exchanges.

10g. Who Runs the Pipes: Google, The Trade Desk, and Amazon

In the open web's chain, Google sits on both sides. It runs the most-used publisher ad server (DFP, now part of Google Ad Manager), the largest exchange (AdX), and buying tools for advertisers, Google Ads and Display & Video 360 (DV360). In the Justice Department's ad tech case, Judge Leonie Brinkema found that Google's ad server held 91% of the worldwide market for publisher ad servers, and that AdX held more than 54% of the exchange market, about nine times the next exchange. She found that AdX had kept a 20% fee, above competitive levels, for more than a decade, about twice what rival exchanges commonly charged. What the court ordered is Module 16's subject.

The largest independent DSP, The Trade Desk, owns no media and sells only buying software. Its revenue was $2.896 billion in 2025, up 18%, with net income of $443 million (February 25, 2026). Its interim finance chief said clients spent about $13.4 billion through it, so its revenue equaled about 21.6% of client spending, a rough measure of its fee that analysts call the take rate (the course's calculation). Buyers report effective

fees of 15% to 20%, with cuts of a point or two in return for spending commitments (eMarketer, citing Digiday).

Then growth stalled: revenue rose 3% in the second quarter of 2026, to $715 million, and the company forecast at least $650 million for the third, about 12% below a year earlier; its shares fell 24% in after-hours trading on August 6, 2026. “This quarter did not meet the standard we set for ourselves,” said Jeff Green, its chief executive. In March 2026 Publicis had dropped The Trade Desk from the platforms it recommends to clients after an audit that, as reported, questioned its fees; the two said on June 12, 2026 that they had resolved the dispute, without giving terms (Digiday). The Trade Desk is due to report its third-quarter results on November 9, 2026.

Its sharpest rival is Amazon, which sells its DSP on price and on the shopping data behind it. A Needham analysis reported by the Wall Street Journal in August 2025 put Amazon's DSP fee at about 1%. Amazon does not publish its fees; Digiday reported in April 2025 that it offered new video advertisers a 10% spending bonus of up to $250,000. Netflix's ads became buyable through Amazon's DSP in September 2025, and when Microsoft closed its own DSP, Microsoft Invest, on February 28, 2026, it named Amazon's as the recommended destination for its clients.

The Trade Desk replies that low fees can hide high prices; on its February 2026 call, Green cited one appliance maker's test in which it reached 70% more households on connected TV than Amazon's DSP at 30% lower total cost, a claim from a single advertiser. Guideline, which tracks agencies' billings, estimates that between 2022 and the first quarter of 2026 Google's DV360 rose from 34% to 41% of DSP spending worldwide, The Trade Desk from 22% to 31%, and Amazon from 14% to 19% (via PPC Land). On those estimates, the three are taking share from every other DSP.

10h. The FTC's Suit Over Amazon's Ad Auctions

Auction rules are a pricing lever, and on August 31, 2026 the FTC and 22 states sued Amazon in federal court in Seattle, alleging that it had pulled that lever without telling advertisers. The suit concerns Sponsored Products, Sponsored Brands, and display ads in Amazon's own store, not its DSP. According to the complaint, Amazon told advertisers it ran a second-price auction, in which the winner pays just over the runner-up's bid. From 2019, the FTC says, Amazon added an undisclosed “soft reserve price”, its own estimate of what a slot was worth, which an internal document called an “invented auction participant”, and charged winners up to their full bids.

The FTC says winners in Sponsored Products auctions paid their own full bid in 30% to 40% of cases in 2021, about 70% in 2022, and about 80% in 2024. It says the practice reached more than a million brands and sellers over seven years, “likely extracting tens of billions of dollars”. Amazon called the suit “misguided” and said the complaint “fundamentally misunderstands how advertisers operate”. It said that auctions weighing billions of bids produce varying prices, that advertisers are properly informed, and that no advertiser ever pays more than its own bid. The suit is separate from the FTC's antitrust case against Amazon (Module 16).

The dispute turns on the gap between bid and price. An advertiser promised a price just over the runner-up's bid can safely bid its full value, while one charged its own bid must shade it, and the worked example below measures how large that gap can be.

Worked example — An Impression Sold, and a Search Auction

Inputs: a reader opens a news article whose page has one ad slot, and the publisher runs header bidding with three exchanges, each passing on one DSP's bid. The UK's Competition and Markets Authority found average DSP fees of 14% (range 5% to 42%) and SSP fees of 16% (8% to 25%) in 2019 data (July 2020). A US practitioner put both at 10% to 20%, with audience data often charged at about $1 per thousand (The Drum, April 29, 2025). The advertiser's agency buys this audience at $10.00 per thousand impressions, all fees included. Step 1, the request: the page's Prebid code asks the three exchanges for bids at once, and each sends an OpenRTB request to its DSP with a deadline of, say, 120 milliseconds. Step 2, the bids: DSP 1 offers $6.20 per thousand, DSP 2 $7.50, and DSP 3 $5.80. Step 3, the clearing: under first-price rules DSP 2 wins and pays its bid, $7.50 per thousand, which is $0.0075 for this one impression. Under the old second-price rule it would have paid $6.21, 17.2% less, which is why its software tries to shade its bids toward the lowest price likely to win. Step 4, to the screen: the best header bid goes to the publisher's ad server, which finds nothing higher among its direct deals or Google's exchange and tells the page to load DSP 2's ad, and the advertiser's ad server counts the impression, all within about a second. Step 5, the split of the advertiser's $10.00: the DSP's fee, at 15% of spending, is $1.50 (or 20% of the $7.50 media cost, depending on how the fee is quoted); the audience data cost $1.00; the remaining $7.50 is the winning bid; the exchange keeps 16% of it, $1.20; and the publisher receives $6.30, or 63%. Step 6, the range: with fees at the low end (10% to the DSP, no data, 10% to the exchange) the publisher gets $8.10, or 81%; at the high end (20%, $1.00 of data, 20%) it gets $5.60, or 56%. Step 7, a search auction: three advertisers bid on the query “running shoes”. A bids $4.00 a click with a quality score of 6, B bids $3.00 with a score of 9, and C bids $5.00 with a score of 4; ad rank is bid times quality, and the platform's minimum ad rank to show is 10. The ranks are A 24, B 27, and C 20, so B takes the top slot, A the second, and C the third. Step 8, the prices: each pays the ad rank below it, divided by its own quality score, plus a cent. B pays 24 ÷ 9 + $0.01 = $2.68; A pays 20 ÷ 6 + $0.01 = $3.34; and C, with no one below, pays the minimum, 10 ÷ 4 + $0.01 = $2.51. Step 9, the stakes of the rule: if the platform charged each winner its own bid, the three would pay $3.00, $4.00, and $5.00, which is 12%, 20%, and 99% more for advertisers who had bid expecting a second price. Result: in the display sale the publisher kept 63 cents of the advertiser's dollar, and the fee range alone moves that between 56 and 81 cents. In the search sale the lowest bidder won the top slot and paid less per click than the advertiser below it, because quality multiplies the bid. In both, the rules of the auction move the price as much as the bids do.

Module 10 Key Points

How Ad Tech Works: The Auction

• Search ads are ranked by bid times quality, and each winner pays just enough to beat the ad ranked below it.

• Meta ranks ads by total value: the bid times the predicted chance of the paid-for action, plus ad quality.

• Open-web ads are sold in auctions lasting well under a second; first-price rules became standard by 2019, so buyers shade their bids.

• The ANA found only 36 cents of each programmatic dollar bought a measurable, viewable, fraud-free impression on a site not made for advertising.

• A court found Google's publisher ad server held 91% of its market, and its exchange charged about twice rivals' fees.

• The FTC alleges Amazon described second-price auctions but charged sponsored-ad winners their own bid about 80% of the time.

Module 10 Quiz

Check Your Understanding

1. Under the rule Google describes for its search auction, what does a winning advertiser pay per click?

A) Its maximum bid, as in any first-price auction
B) The average of all bids in the auction, adjusted for quality
C) A fixed reserve price that Google sets for each keyword in advance
D) Just enough to clear the thresholds and beat the ad ranked below

2. Besides the advertiser's bid, what does Meta's auction weigh when it picks the winning ad?

A) How recently the advertiser last won an auction, so that buyers share impressions
B) The size of the advertiser's monthly budget
C) The predicted chance that this person takes the paid-for action, and the ad's quality
D) The advertiser's total past spending with Meta

3. Why did header bidding spread among publishers in 2014 and 2015?

A) It let publishers stop paying for ad servers and sell their space directly
B) It let several exchanges bid at once, ahead of the ad server
C) A court ordered Google to let rival exchanges bid on equal terms in 2014
D) It replaced auctions with fixed prices agreed in advance, which publishers preferred

4. What does the ANA's TrueAdSpend measure that the ISBA/PwC studies do not?

A) The share of spending that reaches the publisher
B) The share of spending that buys a measurable, viewable, fraud-free impression on a site not made for advertising
C) The fees that media agencies charge on top of DSP fees
D) The number of websites on which an average campaign appears

5. What does the FTC's suit of August 31, 2026 allege about Amazon's sponsored-ads auctions?

A) That Amazon's DSP charged hidden fees of 20% on ads bought for other websites
B) That Amazon refused to sell ads to sellers who also used other marketplaces
C) That Amazon openly ran first-price auctions but set its reserve prices too high under federal antitrust law
D) That Amazon described second-price auctions but, via a hidden reserve, often charged winners their own bid

Module 11 — Television, Streaming, Audio, and Outdoor

The Upfronts, the Super Bowl, and the Streamers' Ad Tiers

Recommended Reading / Reference

Book: Amanda D. Lotz, *We Now Disrupt This Broadcast: How Cable Transformed Television and the Internet Revolutionized It All* (2018), a media scholar's account of how cable bundles, and then internet distribution, changed the way American television is paid for, advertising included

Reference: Nielsen, *The Gauge* (monthly), the share of US television viewing taken by broadcast, cable, and streaming, and by each large streaming service: the plainest public record of where the audience has gone

11a. Television's Money

Television is still one of advertising's largest channels, and among its slowest-growing. WARC expects worldwide TV advertising of $176.2 billion in 2026, up 1.2%, and a fall of 1.8% in 2027 (October 8, 2026). In the United States the traditional business, linear TV, meaning programs watched as they air on broadcast stations and cable networks, is shrinking faster. MoffettNathanson estimates that national and local linear TV took $55.2 billion of advertising in 2025, 7% less than in 2024 on its “core” measure, and forecasts $51.6 billion in 2026, down 6%, and $47.9 billion in 2027 (December 4, 2025).

Seller	2025	Change from 2024
National cable networks	$20.3 billion	−10%
Local TV stations	$16.9 billion	−22%
National broadcast networks	$11.6 billion	−8%
Local cable	$4.2 billion	−20%
Syndicated programs	$2.2 billion	−8%
Ad-supported streaming (outside the linear total)	$16.1 billion	+15%

MoffettNathanson, December 4, 2025, via Inside Radio. The first five rows make up the linear total. Local stations fell against a 2024 lifted by the presidential race and the Olympics. Against that 2024, the five rows imply a fall of about 14% in the total, twice the 7% the source gives for its core measure; it does not reconcile the two. The streaming figure uses MoffettNathanson's own definition.

The table shows both the decline and its cause. Viewers have moved to streaming, and advertisers have followed them. Buyers are less gloomy than the sellers' analysts: the IAB's survey of about 200 US ad buyers, published in January 2026, expected linear TV spending to fall 1.7% in 2026 and connected-TV (CTV) spending to rise 13.8%. The two sources measure different things: sellers' revenue in one case, and buyers' plans in the other.

11b. The Upfront and Scatter Markets

American networks sell most of their prime-time advertising in advance. Each May, at presentations in New York known as the upfronts, they show the coming season's programs, and in the weeks that follow advertisers commit to large shares of the year's inventory. In return they usually get lower prices and an audience guarantee: if a program draws fewer viewers than promised, the network makes up the shortfall with extra spots, called make-goods, and buyers can usually cancel part of a commitment later in the year.

What remains is sold closer to air in the scatter market, at whatever the market will pay, which is above upfront prices when demand is strong and below them when it is weak. The guarantees are written in Nielsen's ratings or, increasingly, in rival measures (Module 13).

Media Dynamics, which estimates the market each year, found that streaming services took more than half of the 2026-27 prime-time upfront for the first time, while broadcast and cable together took $16.6 billion (August 2026), and Netflix said its own upfront commitments nearly doubled. The prime-time upfront is not the whole market: Media Dynamics puts about $12 billion more of linear commitments outside prime time. Prices per viewer fell in every category even as volume rose, because supply grew faster than demand: streaming services added ad-supported plans faster than advertisers added budgets, and the linear networks had smaller audiences to sell.

Seller	Volume, 2026-27	CPM	Change in CPM
Broadcast networks	Down 5.3%	$41.65	−4.2%
Cable networks	Down 7.7%	$17.70	−8.5%
Streaming services	$17.2 billion, up nearly 30%	$25.90	−4.9%
All prime time	$33.8 billion, up 9.1%	Not given	−6.2%

Media Dynamics estimates for the 2026-27 prime-time upfront, via Advanced Television, August 12, 2026; MediaPost gives the total as $33.7 billion. CPMs are per thousand adults for a 30-second ad. The streaming average includes free ad-supported channels, which pull it down.

11c. Sports and the Super Bowl

Sports hold their prices because they are watched live, by large audiences, with the ads. In Ad Age's survey of the 2024-25 season, the latest found, a 30-second spot in NBC's *Sunday Night Football* cost an average of $1,008,746, up 14%. ABC's *Monday Night Football* cost $637,718 and Prime Video's *Thursday Night Football* $562,918, while the costliest entertainment program, *American Idol*, cost $131,943. Spot prices fell for 33 of the 53 returning series. Football's demand has held since: NBC said it had sold 90% of its 2025-26 *Sunday Night Football* inventory before the season began, to more than 150 advertisers (CNBC, September 3, 2025).

The Super Bowl is the extreme case. NBCUniversal sold time in Super Bowl LX, played on February 8, 2026 and carried on NBC, Peacock, and Telemundo, at about $8 million per 30 seconds on average, according to trade-press reports, and it said a handful of spots went for more than $10 million. Nielsen counted an average audience of 125.6 million, revised up from 124.9 million, and a peak of 137.8 million, a US record. The average was 1.6% below the record of 127.7 million set in 2025, when Fox sold more than ten spots at $8 million each (CNBC).

Super Bowl LXI, set for February 14, 2027, goes to Disney's ABC and ESPN. Disney asked $10 million a spot, plus a matching $10 million of spending elsewhere in its portfolio (Variety, April 2026), and sold out in early August 2026, at a reported $8 million to $9 million, in what was reported as the earliest sell-out on record (Sportcal; The Desk). Advertisers pay because no other evening offers an audience of that size watching the ads, and because the ads are discussed as news before and after the game; whether they pay back is Module 5's question.

11d. Streaming's Ad Tiers

Netflix added a cheaper plan with ads in November 2022. Its ad revenue passed $1.5 billion in 2025, more than two and a half times 2024's (January 2026). It had aimed for about $3 billion in 2026, but on October 20, 2026 it cut the aim to about $2.7 billion, citing a “slower-than-expected rollout of certain advertising product capabilities”, while still expecting growth of more than 80%. Netflix says the plan has more than 250 million monthly active viewers worldwide, by its own measure: members who watched at least a minute of ad-supported programming, multiplied by an estimate of the people in each household (May 2026).

Netflix plans to bring the ad plan to 15 more countries in 2027: nine in Europe, from Austria to Switzerland, on March 1, with Colombia, Peru, and New Zealand the same month, and Indonesia, the Philippines, and Thailand later in the year (Netflix, September 24 and October 20, 2026). It sells its ads both directly and through the DSPs described in Module 10; in July 2026 it dropped its minimum spending requirements for buyers using The Trade Desk (PPC Land, September 2026).

Amazon took the opposite route: in January 2024 it put ads into Prime Video for all US members by default, charging $2.99 a month to remove them, and so created one of the largest ad-supported audiences in a single day. Amazon says the average ad-supported audience reached 315 million worldwide in the fourth quarter of 2025, from 200 million in April 2024, and that its *Thursday Night Football* games averaged more than 15 million viewers, up 16%.

Peacock, NBCUniversal's service, had 48 million paid subscribers in mid-2026. Its advertising revenue rose 68% in the second quarter of 2026, to $717 million, helped by the NBA playoffs and the World Cup, and the service made its first quarterly profit on Comcast's adjusted measure (July 23, 2026). Disney+, Hulu, HBO Max, and Paramount+ also sell cheaper plans with ads.

11e. Why Streaming Prices Fell

Streaming's prices per thousand viewers have fallen fast, and Netflix's show how far.

When	Netflix's price per thousand	Source
November 2022, at launch	About $65	Adweek
2023	$39 to $45	Wall Street Journal
June 2024, in negotiations	$29 to $35	Wall Street Journal
July 2024, some inventory	$20 to $30	Three buyers, to Adweek
2026, one agency's estimate	$20 to $30 through auctions; $45 to $65 bought directly	Stillwater Media, August 24, 2026

Adweek, August 6, 2024, citing the Wall Street Journal for 2023 and June 2024. Stillwater Media says its ranges come from its own 2026 negotiations and campaign reports.

Amazon's entry was the turning point: in the six months after Prime Video's ads began, eMarketer measured CPMs falling 27.6% on Disney+, 26.3% on Netflix, 20.5% on Prime Video itself, and 13.2% on Peacock (December 2024). Media Dynamics' streaming average fell 17.6% in the 2023-24 upfront and 7.6% in 2025-26, before the smaller fall in the table in section 11b. Most streaming time now sells well below broadcast prices. Open-exchange CTV, bought through the auctions of Module 10, runs about $14 to $25 per thousand on one agency's 2026 estimate, and free services such as Tubi, along with free ad-supported streaming TV (FAST) channels, about $12 to $22 (Stillwater Media).

The most-watched streaming service on American televisions is YouTube, which in Nielsen's Gauge for August 2026 took more than 15% of all US TV viewing for the first time, and TheWrap reported the figure as

15.8%. Alphabet said on its October 28, 2026 call that living-room screens produced more than a third of YouTube's ad revenue in the third quarter.

Estimates of the whole US connected-TV market vary with the definition. MoffettNathanson forecasts $19 billion of CTV advertising in 2026, while the IAB's video report puts US CTV spending at $29.3 billion (May 2026). Both have it growing while linear TV shrinks, and the IAB expects digital channels to take 61% of US spending on TV and video advertising in 2026, up from 58% a year earlier.

11f. Audio: Radio's Slide, Podcasts' Rise

WARC expects worldwide audio advertising of $41.9 billion in 2026, up 2.5%: $25.3 billion for radio, down 1.9%, and $16.6 billion for online audio, up 9.9% (October 8, 2026). In the United States, BIA Advisory Services forecasts over-the-air radio advertising of about $10.08 billion in 2026, up 1.8% after modest declines in 2025, with radio companies' digital sales adding $2.38 billion (April 9, 2026). Spotify reported €482 million of ad-supported revenue in the third quarter of 2026, up 10%, from 499 million monthly users of its free tier (October 22, 2026).

Podcasts are audio's growth: the IAB and PwC counted $2.9 billion of US podcast advertising in 2025, up 17.6% (April 16, 2026). Host-read ads, in which the presenter reads the message and often vouches for the product, cost about $25 to $40 per thousand downloads, and pre-recorded spots of up to 60 seconds about $15 to $30 (Castos, September 2026). The Libsyn and AdvertiseCast marketplace's average for a 60-second spot was $21.37 in September 2024, the last month it published.

Castos cites Nielsen's finding of 71% brand recall for host-read ads, against 62% for announcer-read ones. Most podcast ads are no longer recorded into the episode but inserted into each download as it is requested, and such dynamic insertion accounted for more than 90% of podcast ad revenue in the IAB's study of October 2023.

11g. Outdoor and Print

Out-of-home advertising, on billboards, buses, and screens in public places, is the old medium still growing. WARC expects $63.3 billion worldwide in 2026, up 8.1%, of which digital screens take $21.7 billion. The Out of Home Advertising Association of America (OAAA) counted a record $9.46 billion in the United States in 2025, up 3.6%, with digital at 36.3% of the total. In the second quarter of 2026, quarterly revenue passed $3 billion for the first time, at $3.16 billion, up 10.7%, as digital rose to 38.4% of the total and transit advertising grew 23.9% (OAAA, August 2026).

Outdoor's sellers argue that a billboard cannot be skipped or blocked and reaches people where screens at home do not, and digital screens can now be bought programmatically by time of day and place. Programmatic digital out-of-home cost an average of $7.62 per thousand in the second half of 2024 (Place Exchange, February 2025). The industry also promotes a comparison by Solomon Partners that puts US out-of-home CPMs at $2 to $16, against $13 to $49 for TV. A billboard's audience is estimated from traffic counts rather than measured like a TV audience, so the comparison is approximate.

Print keeps shrinking. WARC expects $44.4 billion of print advertising worldwide in 2026, down 2.4%, made up of $31.3 billion for news brands, down 2.7%, and $13.2 billion for magazines, down 1.5%. Much of the money newspapers and magazines lost went to the platforms of Module 8 and the auctions of Module 10.

Worked example — A Super Bowl Spot Against a Streaming Buy

Inputs: Super Bowl LX's reported average price and Nielsen's average audience (11c), and Media Dynamics' 2026-27 prime-time CPMs for streaming and broadcast (11b). Step 1, the Super Bowl's price per thousand viewers: $8,000,000 ÷ 125,600 = $63.69. Across the range of reported prices, $7 million to $10 million, it runs from $55.73 to $79.62, and measured against the 137.8 million peak it is $58.06. Step 2, the same $8 million in streaming buys $8,000,000 ÷ $25.90 × 1,000 = 308.9 million impressions; in broadcast prime time it buys 192.1 million. Step 3, impressions are not people. A streaming plan delivers its impressions over weeks and repeats them to the same viewers. If it shows the ad to each person four times on average (an assumption), 308.9 million impressions reach 77.2 million people, at $103.60 per thousand people reached. Step 4, the break-even: $63.69 ÷ $25.90 = 2.46, so a streaming plan whose average frequency is above about 2.5 costs more per person reached than the Super Bowl. Step 5, the caveats: Nielsen's count includes viewers aged 2 and over, while Media Dynamics' CPMs are per thousand adults, which flatters the Super Bowl a little. Production is extra in both cases, and a streaming plan can choose its viewers while the Super Bowl takes everyone. Result: per impression, the Super Bowl costs about two and a half times the average streaming price. Per person reached, it can be the cheaper buy, because it puts one ad in front of about 125 million people in one evening, each counted once.

Module 11 Key Points

Television, Streaming, Audio, and Outdoor

• US linear TV advertising is forecast to fall 6% in 2026, to $51.6 billion (MoffettNathanson).

• Streaming took more than half of the 2026-27 prime-time upfront for the first time, while prices per viewer fell.

• A Super Bowl LX spot cost about $8 million, roughly $64 per thousand viewers.

• Netflix cut its 2026 ad revenue aim to about $2.7 billion on October 20, 2026.

• YouTube passed 15% of all US TV viewing in August 2026, the most of any streaming service.

• Out-of-home advertising keeps growing while radio and print shrink; WARC expects it to rise 8.1% worldwide in 2026.

Module 11 Quiz

Check Your Understanding

1. What is the scatter market in American television?

A) Ad time sold closer to air at market prices, after the upfront
B) Free spots that networks give when a program misses its audience guarantee
C) Ads placed across many cable networks at once by a single sales firm
D) Unsold streaming inventory, auctioned through exchanges

2. Which statement matches Media Dynamics' estimates for the 2026-27 prime-time upfront?

A) Linear TV's volume rose while streaming's fell, as advertisers went back to live sports
B) Prices per viewer rose in every category, because advertisers' demand outran supply
C) Streaming's share passed half for the first time, as prices per viewer fell
D) Broadcast networks had the lowest CPM of the three kinds of seller, despite rising volume

3. In the worked example, why can a Super Bowl spot be the cheaper buy per person reached, despite its higher price per thousand?

A) Networks give Super Bowl advertisers make-goods worth half the $8 million price
B) Streaming services charge separately for every impression served to children, unlike networks
C) It reaches about 125 million people once each, while streaming repeats its ads
D) Super Bowl prices include production, while streaming prices do not

4. How did Amazon's approach to ads on Prime Video differ from Netflix's?

A) Amazon sold its ads only through programmatic exchanges
B) Amazon launched its ads in Europe first and the US two years later
C) Amazon added a cheaper ad plan that viewers had to opt into, as Netflix did
D) Amazon put ads into Prime Video for all US members by default and charged extra to remove them

5. Which of these channels does WARC expect to grow fastest worldwide in 2026?

A) Out-of-home
B) Broadcast and cable TV
C) Over-the-air radio
D) Newspapers and magazines

Module 12 — Influencers, Creators, and Sponsorship

Paid Voices, Disclosure, and the Names on Stadiums

Recommended Reading / Reference

Book: Taylor Lorenz, *Extremely Online: The Untold Story of Fame, Influence, and Power on the Internet* (2023), a technology reporter's history of how bloggers, YouTubers, and influencers turned audiences into businesses, and how the platforms courted and squeezed them

Reference: Federal Trade Commission, *Guides Concerning the Use of Endorsements and Testimonials in Advertising* (revised June 2023), with the agency's plain-language answers to common questions and its 2024 rule on consumer reviews: the US rules for paid voices

12a. Creators and What Brands Spend on Them

A creator is someone who builds an audience by publishing their own work online, and an influencer is a creator whom brands pay to reach that audience; the words overlap, and the business uses both. Goldman Sachs estimated in April 2023 that the creator economy was worth about $250 billion and could reach about $480 billion by 2027. It counted about 50 million creators worldwide, of whom about 4% earned more than $100,000 a year.

What brands spend on creators depends on what is counted, and the three best-known estimates differ by more than three times.

Estimate	Scope	Figure	Source
Influencer marketing	Worldwide, 2025	$32.55 billion, from $24 billion in 2024	Influencer Marketing Hub, annual benchmark report
Influencer marketing	United States, 2025	$10.52 billion, up 15.0%; about $12.2 billion forecast for 2026	eMarketer, March 13, 2025
Creator advertising	United States, 2025	$37.1 billion, up 26%; $44 billion forecast for 2026	IAB, 2026

Influencer Marketing Hub, 2025 report; eMarketer; IAB, as reported by eMarketer and with the IAB/PwC internet advertising report of April 2026.

Influencer marketing, in the narrow sense, is what brands pay creators to make and post sponsored content. The IAB's count of creator advertising rests on a broader definition, so the three figures can be neither added nor compared directly. Each is useful for its own question: what brands pay creators, or how much advertising money now runs through creators' work.

The platforms pay creators too, and sell creators' work to advertisers. YouTube said in September 2025 that it had paid more than $100 billion to creators, artists, and media companies over the previous four years, much of it a share of the advertising sold against their videos. Meta said on its October 21, 2026 call that partnership ads, in which a brand pays to run a creator's post as an ad, had passed an annual run rate of $15 billion, up 97% in a year. In a study published October 7, 2026, Meta also claims that such ads cut the cost of each additional purchase by 45%; it is the seller measuring its own product, a problem Module 13 takes up.

TikTok Shop turns creators into salespeople, paying them commissions on goods sold through their videos and live streams. Momentum Works estimates TikTok Shop's US gross merchandise value (GMV), the

value of goods sold, at $15.1 billion in 2025, up 68%, and its worldwide GMV at $64.3 billion (February 2026).

12b. What Creators Charge

Creators are priced per post, by platform and by tier of followers, from nano to mega, and each source draws the lines differently. Published rates within a tier differ by 5 to 20 times. Some sources report long-standing ranges with no stated sample, and others, such as impact.com, report what creators ask on their rate cards.

Tier	Influencer Marketing Hub	impact.com
Nano	500 to 10,000 followers: $20 to $100	Under 15,000 followers: $500 to $2,000
Micro	10,000 to 100,000: $100 to $5,000	15,000 to 75,000: $2,000 to $8,000
Mid-tier	Included in micro	75,000 to 250,000: $8,000 to $20,000
Macro	100,000 to 500,000: $5,000 to $10,000	250,000 to 1 million: $20,000 to $45,000
Mega	Over 500,000: $10,000 or more	Over 1 million: $45,000 or more

Price per Instagram post (impact.com: per Instagram post or TikTok video). Influencer Marketing Hub's ranges as republished by Meltwater (June 2026); impact.com's US averages, compiled from creators' rate cards (May 2025).

Because follower counts say little about how many people see a post, buyers increasingly compare creators by cost per thousand views, the creator's version of a CPM. Influencer Marketing Hub puts typical creator CPMs at $5 to $15 on Instagram, $3 to $10 on TikTok, and $10 to $30 on YouTube (August 2026). Modash, an influencer-search company, puts them higher from its own experience: $15 to $45 for Instagram Reels, $7 to $25 on TikTok, and $50 to $75 or more for a dedicated YouTube video. It notes that paid ads cost far less per thousand because they carry no production costs and borrow no one's credibility. Format matters too: a 60-to-90-second mention inside a YouTube video, called an integration, costs about $15 to $25 per thousand views, against $30 or more for a video devoted to the brand (CreatorDB, March 2025).

Extras raise the price. Later's 2026 guide says that permission to run a creator's post as a paid ad, called whitelisting, adds 30% to 50% to the fee, and that agencies add 10% to 20% for macro and mega creators (July 7, 2026). Many deals also pay a commission of 10% to 20% on sales (impact.com). Brands that hire an influencer agency pay for its share: in an ANA survey of late 2025, agencies kept about 30% of influencer spending and creators received about 70%, and 61% of marketers said their agency's pay was not transparent or was unknown to them (MediaPost, February 2026).

12c. The Creator as Company

The largest creators have become consumer-goods companies. Jimmy Donaldson, known as MrBeast, runs the most-subscribed channel on YouTube, and his company, Beast Industries, had sales of more than $400 million in 2024 while seeking investment at a valuation of about $5 billion (Bloomberg, via Fortune, February 27, 2025). Its investor presentation, as reported by Fortune, showed Feastables chocolate selling about $250 million and earning more than $20 million, while the media business lost about $80 million and the company as a whole about $60 million (March 11, 2025).

The pattern is the creator economy's in small: the videos build the audience, and the products sold to that audience make the money. Beast Industries has kept expanding, taking a $200 million investment from the crypto company Bitmine in January 2026 and buying Step, a banking app for young people, on February 9, 2026 (Fast Company).

New kinds of sponsor follow the audience. In the 90 days to February 3, 2026, the prediction markets Polymarket and Kalshi sponsored at least 101 and 140 YouTube videos respectively (Tubefilter); how gambling ads are regulated is Module 15's subject.

12d. Disclosure: The Rules for Paid Voices

The law's starting point is that an audience should know when a recommendation is paid for. In the United States, the FTC's Endorsement Guides call for a clear disclosure of any “material connection” between an endorser and a brand, such as payment, free products, or a family or business tie. The FTC revised them in June 2023, adding guidance on fake and incentivized reviews. The guides explain how the FTC reads the law against deception, but they are guidance, not a rule.

In November 2023 the FTC sent warning letters to two trade groups, the American Beverage Association and the Canadian Sugar Institute, and to 12 influencers they had paid to post about aspartame and sugar without clear disclosure. The letters noted that civil penalties could reach $50,120 per violation.

Penalties became easier to seek with the Consumer Reviews and Testimonials Rule, in force since October 21, 2024. It bans fake reviews and testimonials, including ones attributed to people who do not exist or generated by AI, and paying for reviews on condition that they say something good or bad. It also bans undisclosed reviews by a company's own managers, company-run review sites posing as independent, suppressing negative reviews by threats, and buying or selling fake signs of social-media influence, such as followers or views.

Because it is a rule, the FTC can seek civil penalties from a first violation, now up to $53,088 each. It sent warning letters to ten companies on December 22, 2025, and brought its first cases in 2026. In April it charged TruHeight, a seller of height supplements for children, with offering customers incentives for five-star reviews, having employees post fake ones, and running fake social-media profiles, and announced a proposed settlement. On May 11 the Justice Department, acting for the FTC, and the state of Illinois sued Premium Home Service, accusing it of inventing local repair businesses and posting fake five-star reviews for them.

Britain moved in October 2026, following guidance for creators that the CMA issued on September 3, 2025. A guide issued jointly by three bodies makes “#ad” the expected label: the ASA, which regulates advertising; the Committee of Advertising Practice (CAP), which writes the advertising codes; and the CMA. It says “Sponsored”, “Spon”, and “Sp” are not enough, and it covers AI-generated content (Lewis Silkin, October 23, 2026). Under the Digital Markets, Competition and Consumers Act 2024, the CMA can now fine breaches of consumer law directly, up to 10% of worldwide turnover.

France has a law written for influencers, Law 2023-451 of June 9, 2023. It requires paid posts to be labeled “publicité” or “collaboration commerciale”, requires retouched and AI-generated images to be labeled, and bars influencers from promoting cosmetic surgery. The EU has no influencer law of its own. Its general ban on hidden advertising applies, and the Commission's planned Digital Fairness Act is expected to address influencer marketing (Module 15).

12e. Sponsorship: Names on Shirts, Stadiums, and Games

Sponsorship buys association rather than attention: a brand pays to attach its name to a team, an event, or a building, and to use the connection in its own advertising. Europe's sponsorship market was €34.45 billion

in 2025, up 4.7%. Sport took 72% of it, and football 54% of the sport total, about €13 billion (European Sponsorship Association and Nielsen Sports, March 2026).

American leagues have sold space they once kept bare. The NBA approved jersey patches in April 2016 as a three-year pilot from the 2017-18 season, a patch about 2.5 inches square on the front of the jersey. The average deal is now worth more than $9 million a year, and a few exceed $20 million (SponsorUnited, June 24, 2025). Naming rights run longer: Crypto.com agreed in November 2021 to pay a reported $700 million over 20 years to rename Los Angeles's Staples Center, described at the time as the largest such deal in American sport.

The largest events sell worldwide rights. The top tier of Olympic sponsorship, The Olympic Partner (TOP) programme, counted AB InBev, Airbnb, Alibaba, Allianz, Coca-Cola, Deloitte, Omega, P&G, Samsung, TCL, and Visa among its partners in April 2026, when JPMorganChase joined through the 2030 Winter Games (Sportcal, April 28, 2026).

FIFA's 2026 World Cup, held in the United States, Canada, and Mexico in June and July, closed a four-year cycle in which FIFA expected about $13 billion of revenue. About $4.3 billion came from television rights, more than $3 billion from tickets, and about $2.8 billion from sponsorship, $1.8 billion of it in 2026 (estimates and FIFA projections reported by Milano Finanza, July 20, 2026). FIFA said in March 2026 that it had sold every global sponsorship package for the tournament, calling the result the “highest sponsorship revenue ever recorded by a stand-alone sporting event”, without giving a figure. Its top-tier partners were Adidas, Aramco, Coca-Cola, Hyundai-Kia, Lenovo, Qatar Airways, and Visa, and it forecasts more than $14 billion of revenue for 2027-30 (Sportcal).

Product placement, paying to put a product inside a film, a show, or a video, is sponsorship's quieter relative. PQ Media estimated worldwide spending on it at $29.6 billion in 2023, up 12.3% (July 2024).

Worked example — Pricing a Creator Campaign Against Paid Social

Inputs: a brand books ten creators for one post each, at rates within the ranges of Later's 2026 guide (July 7, 2026). Six nano creators with 8,000 Instagram followers each charge $350 a Reel (Later's range, $200 to $500); three micro creators with 60,000 TikTok followers each charge $1,600 a video ($800 to $2,500); and one macro creator with 400,000 Instagram followers charges $7,500 a Reel ($2,500 to $10,000 or more). For paid social, Influencer Marketing Hub cites about $8.96 per thousand impressions on Meta. Step 1, the fees: 6 × $350 + 3 × $1,600 + $7,500 = $2,100 + $4,800 + $7,500 = $14,400. Step 2, the audience: 48,000 + 180,000 + 400,000 = 628,000 followers, so the campaign costs $22.93 per thousand followers. Step 3, the views, which no one knows until the posts run. If each post is viewed by a number equal to 25% of the creator's followers, the campaign gets 157,000 views, at $91.72 per thousand; at 50%, it gets 314,000 views at $45.86; at 100%, 628,000 views at $22.93. Step 4, paid social: the same $14,400 at $8.96 per thousand buys about 1.61 million impressions, 2.56 times the creators' combined followers. Step 5, the extras: through an agency that keeps the ANA's average share (section 12b), the brand pays $14,400 ÷ 0.70 = $20,571 for the same posts, and running the macro creator's Reel as a paid ad adds $2,250 to $3,750 to that creator's fee. Result: per view, these creators cost several times what paid social does, and only a view rate of about half their followers brings the campaign near the top of Modash's range for Reels. The case for creators rests on what each view does, through trust, the creator's own production, and content the brand can reuse as ads. A rate card cannot show that, but a lift test can (Module 5).

Module 12 Key Points

Influencers, Creators, and Sponsorship

• Estimates of creator spending differ by definition; the IAB counted $37.1 billion of US creator advertising in 2025.

• Meta's partnership ads, which run creators' posts as paid ads, passed a $15 billion annual run rate in 2026.

• Published creator rates differ 5 to 20 times within a tier, so buyers compare cost per thousand views.

• The FTC's reviews rule, in force since October 2024, allows fines for fake reviews and bought followers; its first cases came in 2026.

• Britain's October 2026 guide makes #ad the expected label; “Sponsored” or “Spon” alone is not enough.

• Europe's sponsorship market reached €34.45 billion in 2025, most of it in sport.

Module 12 Quiz

Check Your Understanding

1. Why can the IAB's estimate of US creator advertising exceed Influencer Marketing Hub's worldwide figure?

A) It is measured in euros and converted at a high exchange rate
B) It rests on a broader definition of the spending tied to creators
C) It adds several years of spending together
D) Influencer Marketing Hub leaves the United States out of its worldwide total

2. In the worked example, how do ten creators costing $14,400 compare with the same money spent on paid ads on Meta?

A) They deliver more impressions, because every one of the 628,000 followers sees every post
B) Their cost per thousand followers equals Meta's CPM of $8.96
C) Hiring an agency lowers their total cost by a third
D) Per view, they cost several times as much unless their posts are very widely viewed

3. What changed when the FTC's Consumer Reviews and Testimonials Rule took effect on October 21, 2024?

A) Influencers had to register with the FTC before they could post sponsored content
B) The Endorsement Guides were withdrawn and replaced by the rule
C) The FTC could fine first violations for fake reviews and bought followers
D) Review sites and platforms became liable for every fake review their users post

4. Under Britain's October 2026 guide for influencers, which label is expected on a paid post?

A) “Sponsored”
B) “#ad”
C) “Spon”
D) “Paid partner”

5. What did Beast Industries' investor presentation show, as Fortune reported it in 2025?

A) Its Feastables chocolate made a profit while its media business lost money
B) Its YouTube videos made most of its profit; its products broke even
C) Every division was profitable, including media
D) It had no outside investors and wanted none

Module 13 — Counting: Ratings, Attribution, and Fraud

How the Industry Measures What It Buys

Recommended Reading / Reference

Book: Tim Hwang, *Subprime Attention Crisis: Advertising and the Time Bomb at the Heart of the Internet* (2020), a former Google policy lead's argument that online advertising is priced on metrics few buyers can check, and that they overstate what the buyers get

Report: Interactive Advertising Bureau and Media Rating Council, *Attention Measurement Guidelines* (November 2025), the industry's first standard for how attention metrics must be built and disclosed, published on the IAB's website

13a. What Gets Counted, and by Whom

Every advertising deal rests on a count the buyer cannot make for itself. A television network sells an audience that a ratings firm measures, a website sells impressions its own servers record, and a platform reports the sales its ads produced. Money changes hands on those numbers, so disputes about measurement are disputes about price.

Measurement answers three different questions, each with its own tools. The first is exposure: how many people, and which ones, had the chance to see an ad, the question that ratings, impression counts, and viewability answer. The second is attention, whether the ad registered at all. The third is effect: whether the ad changed what anyone did, the territory of attribution and experiments. The industry is most confident about the first question and least about the third, though budgets are set on the third.

Since the 1960s the industry has policed its counts through the Media Rating Council (MRC), which it set up after congressional hearings into the accuracy of broadcast ratings. The MRC measures nothing itself; it audits measurement services against its standards and grants or withholds accreditation, a seal that buyers and sellers use to decide which numbers to trust. Ratings firms, platforms, and verification companies submit to its audits voluntarily, and some withdraw.

13b. Television's Currency and Its Rivals

In television the count is called the currency, because networks and advertisers write their deals on it. Nielsen has supplied it for decades from a panel of homes whose meters record who is watching. A panel tells buyers which people watched, but it is small, and audiences for niche channels and streaming shows are too thin to measure precisely. Big data from set-top boxes and from smart TVs that recognize what is on screen has scale, but it sees households and devices rather than people.

Nielsen's answer is Big Data + Panel, which blends the two. The MRC accredited it on January 22, 2025, the first time it had accredited a hybrid of panel and big data that estimates individual viewers. Nielsen says the product draws on 45 million households and 75 million devices, and it served as the currency for that year's upfront negotiations (Module 11). Nielsen has lost the seal before: the MRC suspended its national TV accreditation in 2021 over undercounting during the pandemic and later restored it.

The new accreditation has been tested too: in March 2026 the MRC reported “unusual changes in measurement results” and problems with how well Nielsen's panel represents some households, and asked for corrective action. In May 2026 it said the accreditation remained unchanged (eMarketer; The Desk).

The networks have tried to loosen Nielsen's hold. In 2023 several of the largest network groups formed a US joint industry committee (JIC), run through OpenAP, a joint venture of TV networks, to certify rival currencies. It certified Comscore and VideoAmp in April 2024 and iSpot in August 2024. Certification by sellers is not accreditation by an independent auditor, and the rivals' progress with the MRC has been uneven. In June 2026 the MRC dropped VideoAmp from its audit process because the company did not intend to continue, and Nielsen withdrew its cross-media product, Nielsen One, from the process (MRC, July 13, 2026). MediaPost reported in July 2026 that iSpot's accreditation covers only counting when ads aired, not who saw them.

The stakes are concrete: when a network guarantees an audience and the currency shows it fell short, the network owes the advertiser extra spots, called make-goods. A change in method that lowers ratings therefore costs sellers inventory, and one that raises them costs buyers money. Each side prefers the count that favors it, which is why both want an auditor neither controls.

13c. Did Anyone See It? Viewability and Attention

Online, the basic unit is the impression: one ad served to one screen. A served ad may never be seen, because it loads below the part of the page the reader views, in a background tab, or after the reader has scrolled past. The MRC and the Interactive Advertising Bureau set the industry's answer in 2014. A display ad counts as viewable if at least half its pixels are on screen for one continuous second, and a video ad if half its pixels play for two continuous seconds.

Verification firms such as DoubleVerify and Integral Ad Science measure viewability for advertisers, checking what publishers and platforms report. The standard is a floor, not proof of attention: an ad can meet it while the reader looks elsewhere. Many buyers pay only for viewable impressions, and sellers price their inventory accordingly.

Attention metrics try to go further. Vendors score how much attention an ad received, using eye-tracking panels, signals such as time on screen, scrolling, and sound, or both, and advertisers can buy against the scores. On November 12, 2025 the MRC and the IAB finalized guidelines for such metrics. They mandate no single method; instead they set principles every method must meet and disclose, and they say attention should not be used to measure a campaign's outcomes (PPC Land, November 25, 2025).

13d. Who Gets the Credit? Attribution and Experiments

Attribution assigns the credit for a sale among the ads a buyer encountered before it. The oldest rule, last click, gives all the credit to the final ad clicked. It was long the default in web analytics because it is simple, but it favors ads that appear just before a purchase. Searches for the brand's own name and retargeting ads, shown to people who have already visited a site, collect much of the credit, while ads that are seen but not clicked, such as video and television, collect none; the table compares the main methods.

Method	What it uses	What it gets right	Blind spot
Last click	The final ad clicked before a sale	Simple; ties credit to an action the advertiser can see	Credits ads shown to people already about to buy; ignores ads seen but not clicked

Method	What it uses	What it gets right	Blind spot
Multi-touch attribution	Every tracked ad in a buyer's path, weighted by fixed rules or a statistical model	Gives earlier ads some of the credit	Needs cross-site tracking that browsers and Apple now limit (Module 14); still measures correlation, not cause
Media mix model	Weekly sales against spending by channel, prices, and seasons, usually over two to three years	Covers television and other offline media; needs no personal data	Needs years of varied spending; struggles to separate channels whose budgets rise and fall together
Lift test	A randomly chosen group of people, or set of regions, that sees no ads	Measures what the ads caused	Costly and noisy; answers for one campaign at a time

Lift tests, also called incrementality tests, are designed and read as in Module 5's worked example.

Mix models have returned as tracking has declined, and the platforms now supply them. Google released Meridian, an open-source mix model, for general use on January 29, 2025, and Meta offers its own, Robyn. Both are free, and both come from companies whose share of budgets the models help decide; Forrester observed that Meridian is most useful to brands that spend heavily on Google. Experiments remain the reference against which the other methods are judged, but holding customers out forgoes sales, and the effects are small enough to need large samples.

Worked example — Last Click Against a Lift Test

Inputs: a retailer spends $50,000 over four weeks on ads aimed at 1,000,000 people who have visited its site. It holds out a random 20%, or 200,000 people, who see none of the ads. On last click, its analytics credit the ads with 2,500 orders averaging $80, and its gross margin is 30%. Step 1, the last-click return: 2,500 × $80 = $200,000 of sales, a ROAS of 4.0. Gross profit is 30% × $200,000 = $60,000, so the campaign appears to make $10,000 after its cost. Step 2, the experiment: the 800,000 people who could see the ads placed 22,400 orders, or 2.80%, and the 200,000 held out placed 5,400, or 2.70%. Step 3, the lift: 0.10 percentage point × 800,000 = 800 orders caused by the ads, a relative lift of 3.7%. Step 4, the incremental return: 800 × $80 = $64,000 of sales, an incremental ROAS of 1.28. Gross profit is $19,200, a loss of $30,800 after the $50,000 cost. Step 5, compare the two: only 800 of the 2,500 credited orders, 32%, were caused by the ads, and the other 68% would have happened anyway. To break even the ads needed $166,667 of extra sales, about 2,083 orders, a lift of 0.26 point. Step 6, the uncertainty: with groups this size the 95% range on the lift runs from about 0.02 to 0.18 point (Module 5 shows the method). Gross profit then lies between about $3,900 and $34,500, short of the cost even at the top. Result: last click reported a profitable campaign, and the experiment shows a loss. The error runs this way for ads shown to people already close to buying, and the other way for ads that work without a click.

13e. Platforms Grading Their Own Homework

The largest sellers of advertising also measure it. Google, Meta, and Amazon run the auctions, count the impressions, and report the sales their ads produced, using data no outside party sees in full; the trade calls them walled gardens. Their reports are detailed and often the only data available, but the seller's interest is plain, and the best-known error overstated what advertisers received.

That error was Facebook's: in September 2016 it disclosed that its figure for the average time people spent watching video ads had counted only views longer than three seconds, which inflated the average. The Wall Street Journal reported the overstatement at 60% to 80%, and advertisers who sued alleged inflation of 150% to 900%. Facebook agreed in October 2019 to pay $40 million to settle the suit, while calling it without merit (Fortune).

The remedies are audits and shared tools. Platforms submit some of their metrics to the MRC; Netflix's ad system, for example, is accredited for counting US in-stream video impressions (TheWrap, August 10, 2026). Data clean rooms (Module 14) let advertisers check a platform's claims against their own sales records. Platforms also run lift tests for their advertisers, but in those the platform designs the test, chooses the control group, and reports the result.

13f. Fraud: Bots, Spoofed Sites, and Whose Numbers to Believe

Advertising fraud is the sale of impressions or clicks that no person produced. Bots load pages and click ads, malware runs hidden ads on infected computers, and spoofing passes off cheap or fake sites as premium ones. The MRC's broader term is invalid traffic (IVT), which covers routine non-human traffic such as search-engine crawlers as well as deliberate fraud.

Estimates of the loss differ more than 200-fold. Juniper Research put it at $84 billion in 2023, 22% of online ad spending, rising to $170 billion to $172 billion by 2028. Its report, *The Global Cost of Ad Fraud: 2023-2028*, is distributed by Fraud Blocker, which sells protection against click fraud and calls it “our in-depth study.” The ANA's quarterly benchmark of large advertisers' programmatic buying found invalid traffic at 0.1% of spending in the second quarter of 2026.

The two figures measure different things: Juniper defines fraud as the deliberate, repeated clicking of pay-per-click ads and drew on data from 78,000 sources in 45 countries (Marketing Brew, October 5, 2023). The ANA's benchmark covers large advertisers whose buying passes through verification filters, and it counts what got through. Each source also has an interest: a seller of protection gains from a large number, and an industry body whose members buy the media gains from a small one.

The largest proven cases were industrial: Methbot ran from rented data centers, faking video views on more than 6,000 spoofed publisher domains between September 2014 and December 2016 and taking more than $7 million. A second scheme, 3ve, used 1.7 million infected computers and took more than $29 million. US prosecutors unsealed an indictment of eight people in November 2018, and Methbot's organizer, Aleksandr Zhukov, was sentenced to ten years in prison on November 10, 2021 (Department of Justice).

13g. Brand Safety: From GARM to the FTC's Orders

Brand safety means keeping a brand's ads away from content that could embarrass it. The modern push dates from March 2017, when advertisers including AT&T and the British government paused spending on YouTube after The Times of London reported their ads running beside extremist videos. In June 2019 the WFA founded the Global Alliance for Responsible Media (GARM), after further incidents including the live-streaming of a terrorist attack in Christchurch, New Zealand.

GARM gave its members, who included advertisers, agencies, and platforms, common definitions of harmful content. Its brand safety floor listed content no ad should fund, such as terrorism and child sexual abuse, and a wider framework graded risk across categories; in 2022 it added misinformation. Agencies applied such standards for their clients, often through lists of sites and channels to avoid.

The backlash came from Congress and from X. On July 10, 2024 the House Judiciary Committee published *GARM's Harm: How the World's Biggest Brands Seek to Control Online Speech*, alleging that GARM coordinated its members to withhold ad spending from conservative content. In August 2024 X sued the WFA and four advertisers, including Mars and Unilever, alleging an illegal boycott. Days later the WFA closed GARM, citing the cost of the suit.

Judge Jane Boyle dismissed X's case with prejudice in March 2026, writing that “the very nature of the alleged conspiracy does not state an antitrust claim.” X appealed in April, then settled with the WFA on July 29, 2026 on undisclosed terms; the WFA said it would not revive GARM.

The FTC took up the charge: as a condition of Omnicom's purchase of Interpublic (Module 7), it issued an order, final in September 2025, barring the merged firm from steering clients' money by political viewpoint unless a client directs it, and from using broad shared exclusion lists. On April 15, 2026 the FTC and eight states filed a complaint against WPP, Publicis, and Dentsu, with consent orders that a federal court in Texas approved. It alleged that since 2018 the three had colluded through GARM and the Advertiser Protection Bureau of the American Association of Advertising Agencies on a “brand safety floor” aimed at “misinformation”, using ratings from NewsGuard and the Global Disinformation Index (GDI). Havas, the last of the five holding companies, accepted a similar order in early July 2026.

The FTC also investigated Media Matters for America, a media watchdog whose report on ads beside extremist posts on X, which X said had been manufactured, preceded advertisers' departures from the platform in November 2023. A federal judge blocked the FTC's demand for documents on August 15, 2025, finding it likely retaliation for speech the First Amendment protects, and the FTC withdrew the demand in April 2026.

13h. Is “Brand Safety” Advertisers Protecting Themselves, or a Coordinated Boycott? The Cases as Their Advocates Make Them

The case that it is advertisers protecting themselves: the WFA, the advertisers X sued, and media watchdogs such as Media Matters argue that brand safety is ordinary prudence. A brand that pays for an ad beside a terrorist video or child abuse suffers for it, as the YouTube episode of 2017 showed, and each advertiser has the right to decide what its money funds. GARM, in this view, was a standards body that gave buyers a common vocabulary for risk, while each member chose its own lists and budgets. Declining to buy, they add, is not censorship, because the content stays up, and the WFA said X's suit misconstrued GARM's purposes.

The case that it is a coordinated boycott: the House Judiciary Committee under Jim Jordan, X, and the FTC under Chairman Andrew Ferguson argue that brand safety became a cartel. On their account, competing agencies and advertisers agreed on common standards instead of competing on them. The misinformation category, they say, let the members treat lawful conservative news and opinion as unsafe, cutting outlets and platforms off from advertising revenue. They cite *FTC v. Superior Court Trial Lawyers Association* (1990), in which the Supreme Court held that a boycott by competitors seeking economic advantage breaks antitrust law even when it carries a political message.

The evidence: the record is thinner than either side's claims. The one court to rule on the antitrust theory, in X's case, dismissed it, and X then settled. The FTC's charges were settled rather than tried: each of the five holding companies accepted a consent order, which binds it without a court deciding the facts. The House report draws on subpoenaed internal messages about standards and about particular platforms, but it does not measure how much spending moved because of GARM. X's advertising revenue fell after its sale (Module 8), yet no court has separated the effect of shared standards from individual advertisers' decisions. The finding in the Media Matters case concerned the FTC's motive, not the advertisers'.

The course's reading, stated once: each advertiser choosing where its own ads run is lawful and common, and the FTC's orders themselves preserve that choice. The contested practice is narrower: competing buyers agreeing on shared standards and lists, including a misinformation category that calls for judgments about news and opinion. Whether that coordination moved money by viewpoint was alleged and settled, never proven or disproven in court. The orders draw the line where the law now stands, allowing the client's choice and barring joint exclusion, and both descriptions fit part of the record.

Module 13 Key Points

Counting: Ratings, Attribution, and Fraud

• Nielsen's Big Data + Panel won MRC accreditation in January 2025; rival TV currencies rely mainly on certification by the networks themselves.

• A display ad counts as viewable when half its pixels are on screen for one second, which proves exposure, not attention.

• Last-click attribution credits ads shown to people already about to buy; lift tests often find most of those sales would have happened anyway.

• Fraud estimates range from a vendor-distributed $84 billion a year to a sliver of large advertisers' spending; who counts shapes the number.

• The FTC's orders bar the five agency holding companies from joint brand-safety standards and viewpoint steering, while each client may still choose.

• The course's reading: advertisers' own choices are lawful; coordination among competing buyers is the contested part, alleged and settled but never tried.

Module 13 Quiz

Check Your Understanding

1. What did the MRC do on January 22, 2025?

A) Suspended Nielsen's national TV ratings over undercounting during the pandemic
B) Accredited Nielsen's Big Data + Panel, a hybrid of panel and big-data measurement
C) Certified VideoAmp and Comscore as currencies for the upfront
D) Issued its first guidelines for attention metrics, which mandate eye tracking

2. Under the MRC's 2014 standard, when does a display ad count as viewable?

A) When the page that carries it finishes loading, whether or not the reader scrolls to it
B) When the whole ad stays on screen for two continuous seconds
C) When the reader moves the cursor over it
D) When at least half its pixels are on screen for one continuous second

3. In the module's worked example, why did last click show a profit while the lift test showed a loss?

A) Most orders credited to the ads came from people who would have bought anyway.
B) The lift test counted only orders placed on mobile phones, missing those from computers.
C) Last click counted each order twice, once for the search ad and once for the retargeting ad.
D) The holdout group saw more ads than the exposed group.

4. Why do Juniper Research's and the ANA's fraud figures differ so widely?

A) The ANA counts only fraud that has been prosecuted in court
B) Juniper's figure includes fraud in television ratings
C) They measure fraud in different kinds of traffic
D) The ANA's benchmark was taken before programmatic buying became common

5. What do the FTC's consent orders with the agency groups (2025-26) still allow?

A) Sharing exclusion lists among agencies, provided every client approves
B) Following GARM's brand safety floor until a new standard is written
C) A client directing its own agency to avoid particular sites of its choosing
D) Placing ads only on sites rated reliable by NewsGuard

Module 14 — Personal Data, Privacy, and the End of Tracking

Cookies, Device IDs, and the Laws That Limit Them

Recommended Reading / Reference

Book: Joseph Turow, *The Daily You: How the New Advertising Industry Is Defining Your Identity and Your Worth* (2011), a University of Pennsylvania scholar's account of how ad networks and data brokers learned to profile readers and put a price on them, written as the tracking business took shape

Reference: Alessandro Acquisti, Curtis Taylor, and Liad Wagman, “The Economics of Privacy,” *Journal of Economic Literature* (June 2016), a survey of what economists had learned about who gains and who loses when personal data are collected and traded

14a. How Targeting Works

Targeting means choosing who sees an ad on the basis of data about them. First-party data are what a company learns from its own customers: purchases, logins, and visits to its own site or app. Third-party data come from someone else. Data brokers assemble them from public records, loyalty cards, apps, and websites into profiles sold by the segment, such as new parents or people likely to buy a car.

The platforms add tools of their own. An advertiser can upload its customers' email addresses, scrambled so they cannot be read, and the platform shows ads to the matching accounts, or to a lookalike audience of users who resemble them. Retargeting follows people who visited a site with its ads elsewhere. A data clean room lets two companies, often an advertiser and a platform or retailer, match their records under agreed rules and see only totals. Contextual targeting needs no personal data at all: it places ads by the content of the page or video, as print and television always have.

Following a person across sites and apps requires an identifier, long the third-party cookie on the web and the phone's mobile advertising identifier (ID) in apps. Logged-in platforms can recognize their own users without either, so limits on cookies have weighed more on the open web than on Google, Meta, and Amazon, though Apple's limit on app tracking hit Meta hard (Module 10 explains how identity enters the auction).

14b. The Browsers Close the Door

Apple's Safari and Mozilla's Firefox blocked third-party cookies by default years before Google moved, as the table shows. Chrome, used by most of the world's web users, announced a plan to make them obsolete. Google proposed to replace them with the Privacy Sandbox, a set of browser tools meant to let ads be targeted and measured without following individuals from site to site. Advertisers and publishers objected that the tools would give Google more control, and the UK's CMA, concerned about the effect on competition, took binding commitments from Google over the project in 2022.

Gatekeeper	Change	Date
Safari (Apple)	Intelligent Tracking Prevention begins limiting cross-site cookies	2017
Firefox (Mozilla)	Enhanced Tracking Protection switched on by default for all users	September 3, 2019
Chrome (Google)	Plan announced to make third-party cookies obsolete within two years	January 14, 2020

Gatekeeper	Change	Date
Safari (Apple)	All third-party cookies blocked by default (Safari 13.1)	March 2020
iPhone apps (Apple)	App Tracking Transparency: apps must ask before tracking users across other companies' apps and sites	April 26, 2021
Chrome (Google)	Phase-out pushed back to 2023, then to the second half of 2024	June 2021; July 27, 2022
Firefox (Mozilla)	Total Cookie Protection on by default: each site's cookies kept in a separate jar	June 14, 2022
Chrome (Google)	Reversal: third-party cookies stay, with a user choice promised	July 22, 2024
Chrome (Google)	No standalone choice prompt; the setting stays in Chrome's menus	April 22, 2025
Chrome and Android (Google)	Most Privacy Sandbox tools retired	October 17, 2025

Sources: Chromium blog (January 14, 2020); Google's Privacy Sandbox posts (July 22, 2024; April 22, 2025; October 17, 2025); The Register (July 28, 2022); InfoQ (April 13, 2020); Android Central (September 3, 2019); TechCrunch (June 14, 2022); VideoWeek (May 6, 2021).

The deadline slipped repeatedly, then disappeared. Announcing its reversal, Google said that “instead of deprecating third-party cookies, we would introduce a new experience in Chrome that lets people make an informed choice.” It then retired most of the Sandbox's tools, including Topics and Protected Audience, “in light of their low levels of adoption,” keeping a few meant for security and logins. As of October 31, 2026, Chrome accepts third-party cookies by default.

14c. Apple's Prompt and the Fight Over It

Apple's App Tracking Transparency (ATT) requires every app to ask permission before tracking a user across other companies' apps and websites, which includes reading the phone's advertising ID. Most users said no. In the first weeks the analytics firm Flurry counted about 4% of US daily users agreeing, and 11% worldwide, among those who made a choice that day. AppsFlyer, which measures app-install advertising, reported a global average near 37%; the figures differ because each firm counts a different group of users. AppsFlyer said in April 2025 that half of users now consent, without publishing its method.

The change fell hardest on advertisers that measured sales from app ads and on the platforms that served them. Meta's finance chief, David Wehner, estimated on February 2, 2022 that it would cost the company “on the order of $10 billion” of revenue that year, adding: “We can't be precise on this. It's an estimate.”

Competition authorities took up a different complaint, that Apple's prompt was harsher on other companies' apps than on its own. France's competition authority fined Apple €150 million (about $173 million) on March 31, 2025, finding the prompt “neither necessary for nor proportionate with” its privacy aims and its double consent requests a particular burden on smaller publishers. Italy's authority fined it €98.6 million on December 22, 2025 over the same double consent request. Apple warned in October 2025 that pressure in Europe might force it to withdraw the feature there.

Germany chose commitments over a fine. On August 13, 2026 the Bundeskartellamt made binding Apple's pledges that its own apps and others' will get consistent, neutral prompts, that publishers get more room to explain their requests, and that the prompt may be combined with consent under the GDPR. Apple has four months to carry them out, and they last seven years (VideoWeek, August 17, 2026).

14d. Europe: Consent, “Pay or Consent,” and the Omnibus

The GDPR, in force since May 25, 2018, lets a company process personal data only on one of six legal grounds, of which consent matters most for advertising. Meta first relied on another, the contract users accept when they sign up. Ireland's Data Protection Commission, which oversees Meta for the EU, fined it €390 million (about $449 million) on January 4, 2023, ruling that it could not rely on contract for behavioral advertising. Ireland's larger fine that May, €1.2 billion, was for sending Europeans' data to the United States, not for targeting.

The EU's highest court then shaped what consent must look like. On July 4, 2023, in a German competition case against Meta, the Court of Justice of the European Union (CJEU) held that competition authorities may weigh breaches of the GDPR. It added that users who refuse tracking should be offered, “if necessary for an appropriate fee, an equivalent alternative” without it. Meta answered with a subscription: from November 2023, EU users could pay €9.99 a month on the web or €12.99 on phones to see no ads, or else consent to tracking. In November 2024 it cut the prices to €5.99 and €7.99 and added a free option with less personalized ads.

The European Commission found the model wanting under the DMA. On April 23, 2025 it fined Meta €200 million for failing to give users a service that uses less of their personal data. In December 2025 it accepted Meta's commitment to offer, from January 2026, a choice between sharing all data for fully personalized ads and sharing less for more limited ones, and said it would watch how many users choose each. Britain took its own course: after talks with its Information Commissioner's Office, Meta began offering a no-ads subscription in the UK in September 2025, at £2.99 a month on the web and £3.99 on phones.

The EU is now revisiting its cookie rules. The Commission's Digital Omnibus, proposed on November 19, 2025, would have moved them from the ePrivacy Directive into the GDPR and had websites honor a browser signal of each user's choice. The negotiating mandate the Council of the EU agreed on October 21, 2026 drops both. Cookie rules stay in a rewritten ePrivacy article with six purposes exempt from consent, among them first-party audience measurement and security, and a site must wait at least four months before asking again after a refusal (PPC Land, October 25, 2026). The European Parliament has not yet set its position, so the final text remains open as of October 31, 2026.

14e. The United States: State Laws, Brokers, and Location Data

The United States has no general privacy law. As of October 12, 2026, 24 states had signed comprehensive privacy laws and 20 had them in force, counting Florida's narrow one (privacylawmap.com; Husch Blackwell). Oklahoma, Alabama, Louisiana, and Vermont signed theirs in 2026, taking effect between January 2027 and January 2028. Most such laws give residents the right to opt out of targeted advertising and the sale of their data, and several require companies to honor a browser setting that opts out automatically.

Congress has debated a federal law for years without passing one. The SECURE Data Act (H.R. 8413), introduced on April 22, 2026 by Representative John Joyce with Energy and Commerce Chairman Brett Guthrie, would be enforced by the FTC and state attorneys general, give individuals no right to sue, and largely override state laws. It had a hearing on June 3, 2026 and had not been enacted as of October 31, 2026.

The FTC has used its existing powers against the trade in location data, which can reveal visits to clinics, places of worship, and shelters. In January 2024 it barred X-Mode Social and InMarket from selling sensitive location data. In December 2024 it took similar action against Gravy Analytics and Mobilewalla, with final

orders on January 14, 2025. Kochava, which the FTC sued in August 2022, settled on May 4, 2026: it and a subsidiary may not sell or share sensitive location data without consumers' “affirmative express consent” (FTC).

California has gone furthest against brokers. Its Delete Request and Opt-out Platform (DROP), opened on January 1, 2026 under the state's Delete Act, lets residents ask every registered broker to delete their data in one request. More than 242,000 had used it by late February, more than 575 brokers had registered, and from August 1, 2026 brokers must check the platform for new requests every 45 days (International Association of Privacy Professionals, March 2, 2026).

14f. Targeting That Excludes

Targeting can also exclude. Ad systems that let advertisers choose audiences by age, sex, or neighborhood, or that learn on their own whom to show an ad, can keep housing, job, and credit ads from groups the law protects. In a settlement with the Department of Justice on June 21, 2022, Meta agreed to stop using its “Special Ad Audience” tool, a lookalike feature, for housing ads. It also agreed to build a Variance Reduction System (VRS), which adjusts delivery so that the people who see a housing ad match the eligible audience more closely by sex and estimated race or ethnicity. Meta paid a civil penalty of $115,054, the most the Fair Housing Act then allowed, and an outside firm reported on the system's results through October 2024.

14g. Is Targeted Advertising Worth Its Privacy Cost? The Cases as Their Advocates Make Them

The case that it is: Meta, the IAB, and many publishers and small advertisers argue that tracking pays for free services and lets small businesses compete. When Apple announced its prompt, Facebook ran full-page newspaper ads on December 16, 2020 under the headline “We're standing up to Apple for small businesses everywhere.” Citing its own research, it said small businesses' sales per advertising dollar would fall by about 60% without personalized ads. Targeting, on this view, lets a small seller reach a few thousand likely buyers instead of paying for a mass audience, and lets a publisher charge more for each reader. Rules that make consent hard, they add, mostly strengthen the largest platforms, which can target with the data they already hold.

The case that it is not: Apple, privacy advocates such as the Electronic Frontier Foundation, and many European data-protection authorities argue that the trade is made without real consent, and the advocates add that its gains are overstated. Apple answered Facebook that “this is a simple matter of standing up for our users,” who “should know when their data is being collected and shared across other apps and websites” and “should have the choice to allow that or not.” The data that power targeting also feed a broker market that sells records of where people go, including clinics and places of worship. And independent research, they note, finds that tracking adds little to publishers' revenue, so readers get less in return for their data than they are told.

The evidence: the largest measured gains from tracking go to advertisers and platforms. A study of more than 70,000 advertisers on Meta estimated that losing data from other sites and apps would raise the median cost of a customer the ads actually won from $38.16 to $49.93, by 31%, hurting small advertisers most (Wernerfelt, Tuchman, Shapiro and Moakler, *Marketing Science*, 2025). After ATT, click-through rates on Meta's ads optimized for sales fell 37%, and the revenue of firms that depended on them fell between 8%

and 40% relative to less exposed firms, depending on the method, as spending shifted to Google (Aridor, Che, Hollenbeck, Kaiser and McCarthy, *Management Science*, 2025).

For publishers the studies disagree widely, as the worked example below shows. For apps the harm looks small: the number of apps dipped after ATT and quickly recovered (Cheyre, Leyden, Baviskar and Acquisti, revised June 2026). Kesler (revised July 31, 2026) finds that apps barely changed how they make money, while the share of surviving apps disclosing data linked to their users rose from 10.4% to 24.1% (PPC Land, August 7, 2026). The privacy side resists pricing: in experiments, people demanded far more to give up privacy than they would pay to protect it (Acquisti, John and Loewenstein, 2013).

The course's reading, stated once: targeting is worth most to the platforms that hold logged-in data and to small advertisers that sell online, and least, on the independent evidence, to the open-web publishers in whose name it is often defended. Apple's prompt shows what a clear choice produces: most users declined, advertisers that relied on Meta's sales-optimized ads lost revenue, apps survived, and spending moved toward platforms that can target with their own data. The clearest harms, from the resale of sensitive locations, come from brokers, and that is where US enforcement has acted.

Worked example — What Tracking Is Worth to One Publisher

Inputs: a news site sells 30 million programmatic impressions a month at $2.50 per thousand when buyers can track the reader, an illustrative price. Three studies estimate what the same impressions fetch without tracking. Marotta, Abhishek and Acquisti (2019) found that a cookie raised one large media company's revenue by only about 4%, or $0.00008 per ad, which implies about $2 per thousand without one. Johnson, Shriver and Du (2020) found that the 0.23% of ad impressions coming from users who had opted out of tracking fetched 52% less. Ravichandran and Korula (Google, 2019) removed cookies from a random slice of traffic for 96 days and found revenue fell 52% on average for the top 500 publishers, 64% at the median, and 62% on average for news publishers. Step 1, revenue with tracking: 30,000,000 ÷ 1,000 × $2.50 = $75,000 a month, or $900,000 a year. Step 2, the price without tracking: $2.50 ÷ 1.04 = $2.40 on the first study; $2.50 × 0.48 = $1.20 on the second and on Google's average; $2.50 × 0.38 = $0.95 for Google's news publishers; and $2.50 × 0.36 = $0.90 at Google's median. Step 3, monthly revenue at those prices: $72,115, $36,000, $28,500, and $27,000. Step 4, the yearly loss: $34,615 on the first study, against $468,000 to $576,000 on the others, a 13- to 17-fold spread. Result: the studies measure different things. Google's test removed cookies from part of its exchange's traffic while buyers could still track readers elsewhere, so their bids moved to identifiable impressions. Opted-out users are a small, self-selected group, and the media-company study compares one publisher's impressions with and without a cookie. A publisher that loses tracking alone faces something like Google's numbers. If every publisher lost it at once, advertisers would have nowhere else to send the money, and prices would likely fall by less.

Module 14 Key Points

Personal Data, Privacy, and the End of Tracking

• Safari and Firefox block third-party cookies by default; Chrome abandoned its plan to remove them and still accepts them.

• Most users refused tracking under Apple's prompt; Meta put the cost at about $10 billion of revenue in a single year.

• EU pressure pushed Meta beyond “pay or consent”: since January 2026, EU users can share less data and still use its free service.

• The US has no federal privacy law; 20 state laws are in force, and the FTC polices the trade in location data.

• Estimates of what tracking adds to a publisher's revenue differ more than tenfold, depending on whether one publisher or all lose it.

• The course's reading: targeting is worth most to platforms with logged-in data and small online sellers, least to open-web publishers.

Module 14 Quiz

Check Your Understanding

1. What did Google announce in July 2024 about third-party cookies in Chrome?

A) It would block them for every user by the end of 2024
B) It would replace them with its Topics tool for all Chrome users
C) It would block them only for users in the European Union
D) It would keep them and offer users a choice

2. Why did France and Italy fine Apple over App Tracking Transparency?

A) Its double consent requests burdened other apps more than Apple's own
B) It let apps track users across other companies' sites without asking permission
C) It sold users' advertising identifiers to data brokers without their consent
D) It showed the prompt only in English

3. What did the CJEU say in July 2023 about users who refuse tracking?

A) Platforms may refuse to serve them
B) They should get a tracking-free equivalent, if need be for an appropriate fee
C) They must be shown no advertising of any kind
D) Their refusal lapses after four months unless they renew it, so sites may ask again

4. In the worked example, why does Google's cookie test show a far larger loss than the media-company study?

A) Buyers could still track readers elsewhere, so bids moved to identifiable impressions.
B) Google's test lasted a single day, too short for buyers to adjust their bids to the change.
C) The media company sold only video ads.
D) Google measured clicks rather than revenue.

5. What did Meta's 2022 settlement with the Department of Justice require?

A) A ban on all targeting for housing, job, and credit ads
B) A payment of $115 million to fair-housing groups
C) A system to bring the audience that sees housing ads closer to the eligible audience
D) Moving housing ads into a separate app

Module 15 — The Rules: Truth, Children, Drugs, Politics, and Vice

Who Polices Advertising, and What It Forbids

Recommended Reading / Reference

Book: Ivan L. Preston, *The Great American Blow-Up: Puffery in Advertising and Selling* (1975; revised edition 1996), a University of Wisconsin scholar's history of how American law came to tolerate exaggerated sales talk while punishing false statements of fact

Report: Federal Trade Commission, *Policy Statement on Deception* (October 14, 1983) and *Policy Statement Regarding Advertising Substantiation* (1984), the two short documents that still define what the FTC treats as a misleading ad and what proof an advertiser must hold before making a claim

15a. Truth in Advertising: The Law and Its Enforcers

In the United States the FTC is the main enforcer of truth in advertising. Section 5 of the FTC Act forbids “unfair or deceptive acts or practices,” and Section 12 forbids false advertising of food, drugs, medical devices, services, and cosmetics; Module 2 tells how the commission gained these powers. Its 1983 policy statement on deception sets the test it still applies. An ad is deceptive if it contains a claim or omission likely to mislead a consumer acting reasonably in the circumstances, about something material to the decision to buy.

A second statement, of 1984, requires an advertiser to have a “reasonable basis” for an objective claim before making it, so a claim made without evidence is unlawful even if it later proves true. Vague praise such as “the best coffee in town” is puffery, which no reasonable consumer takes as fact, and the law leaves it alone. The FTC can order a company to stop and, at times, to correct the record. In *Warner-Lambert v. FTC* (1977) a federal appeals court upheld an order requiring about $10 million of Listerine advertising to state that the mouthwash would not help prevent colds or sore throats.

Getting money back is harder. In *AMG Capital Management v. FTC* (April 22, 2021) the Supreme Court held that the FTC cannot win refunds directly in court under the provision it had used for decades. The commission now leans on its trade regulation rules, whose violation carries civil penalties and opens the way to refunds, and on statutes that grant such powers directly, which is one reason rulemaking has become central to its work.

Competitors police one another too. Under Section 43(a) of the Lanham Act a company can sue a rival over false advertising, and the Supreme Court strengthened such suits in 2014. *Lexmark International v. Static Control Components* widened who may bring them, and *POM Wonderful v. Coca-Cola* held that a juice label complying with federal labeling rules could still be challenged as misleading. Many disputes between rivals go instead to the National Advertising Division (NAD), the industry's self-regulatory forum, which refers advertisers that ignore its decisions to the FTC.

The NAD handled 185 cases in 2025, including those still pending. Of the 143 it closed, 106, about 74%, were challenges brought by competitors, 34 came from its own monitoring, and 3 were brought in the public interest (NAD annual report, 2025). State attorneys general enforce their own consumer-protection laws, often called little FTC Acts, and join federal cases. In the UK the self-regulatory ASA rules on most complaints, with the CMA and other regulators as a backstop.

15b. The FTC in 2026

Andrew Ferguson has chaired the FTC since January 20, 2025. Since Melissa Holyoak left in November 2025, only two of its five seats have been filled, both by Republicans, Ferguson and Mark Meador, as of October 31, 2026. On June 29, 2026 the Supreme Court ruled 6-3 in *Trump v. Slaughter* that the president could remove a commissioner, Rebecca Slaughter, despite the FTC Act's limits on removal.

Chief Justice John Roberts wrote, “If anything more is left of *Humphrey's*, we overrule it,” ending the 1935 precedent, *Humphrey's Executor v. United States*, that had protected the commissioners of independent agencies. The same day, in *Trump v. Cook*, the Court left the Federal Reserve's governors protected against removal without cause. Justice Sonia Sotomayor, dissenting for three justices, wrote that the decision “does not just overrule precedent; it all but ignores that precedent exists.” Since the ruling, the president may remove any FTC commissioner at will.

Ferguson's stated priorities mix consumer protection with new targets. He has called the antitrust case against Amazon a top priority and promised many cases under the children's privacy law (MLex). In a podcast interview on September 30, 2026 he said the agency was bringing “a ton” of cases against subscription traps and deceptive pricing, especially by car dealers, and planned a policy statement on surveillance pricing, the use of personal data to set individual prices (Troutman Pepper). He also said the FTC had returned more than $2 billion to consumers in his first year.

The rule against subscription traps shows how fragile such rules can be. Adopted in October 2024, it required cancelling to be as easy as signing up, but the Eighth Circuit vacated it in July 2025 because the FTC had skipped a required preliminary analysis of its costs. On March 11, 2026 the commission began again with an advance notice, asking among other things whether discounts offered to keep cancelling customers are themselves unfair. Comments closed on April 13, 2026, and no replacement rule was in force as of October 31, 2026. Its rule on hidden fees, in effect since May 12, 2025, requires sellers of live-event tickets and short-term lodging to show the total price up front. Rules on endorsements and fake reviews are in Module 12.

15c. Children

Children are the audience US federal law protects most directly. The FTC's amended rule under COPPA, published on April 22, 2025, with compliance required from April 22, 2026, applies to sites and apps aimed at children under 13. It requires separate parental consent before a child's data are disclosed to third parties, including for targeted advertising, which the FTC said is not integral to a service, and written limits on how long data are kept. In a policy statement of February 25, 2026 the FTC eased enforcement against companies that collect children's data only to verify their age.

Other countries go further. Britain banned ads for identifiable less healthy foods and drinks on television and on-demand services between 5:30 a.m. and 9 p.m., and in all paid online media, from January 5, 2026. The industry had applied the rules voluntarily from October 1, 2025, and ads for a brand that show no such product are exempt (UK government; ASA). Australia barred children under 16 from holding social-media accounts from December 10, 2025, with fines of up to A$49.5 million for platforms that fail to take reasonable steps. About 4.7 million accounts had been removed or restricted by mid-December (Mediaweek, March 31, 2026). In effect, the rule also takes under-16s out of the audiences that platforms can target with ads.

In the EU, the DSA bars online platforms from showing ads based on profiling when they are aware with reasonable certainty that the user is a minor, and the Commission set out how platforms should protect minors in guidelines of July 14, 2025.

15d. Prescription Drugs

Only the United States and New Zealand allow drug makers to advertise prescription medicines to the public, known as direct-to-consumer (DTC) advertising. In the US the FDA requires a print ad to carry a brief summary of the drug's risks. A broadcast ad need only state the major risks, provided it makes “adequate provision” for viewers to find the full labeling elsewhere, and an FDA rule issued in November 2023 requires that statement to be clear, conspicuous, and neutral. The FDA's guidance of 1997, made final in 1999, said a website, a toll-free number, a print ad, and doctors would satisfy adequate provision, which made the television drug ad commercially workable (Latham & Watkins, July 14, 2026).

The money followed. MediaRadar counted about $10 billion to $11 billion of prescription-drug advertising in 2024: $10.1 billion as Fierce Pharma reported it, $10.8 billion as TheWrap did. AbbVie alone spent $1.2 billion on Skyrizi and Rinvoq, and Novo Nordisk $530 million on Ozempic and Wegovy (Fierce Pharma). iSpot measured $5.96 billion on national television in 2025, up 16% (eMarketer). Drug makers are among the largest advertisers on American television, so the rules matter to broadcasters as well.

On September 9, 2025 a presidential memorandum directed the Department of Health and Human Services (HHS) and the FDA to enforce the rules more strictly and to close what HHS calls the adequate-provision “loophole.” The FDA said it had sent thousands of letters asking companies to remove misleading ads, and about 100 cease-and-desist letters (Fierce Pharma). The law firm Covington counted more than 50 warning letters and more than 50 untitled letters in the third quarter of 2025, most of them after the memorandum.

The FDA's regulatory agenda of July 2026 lists a proposed rule, “Transparency in Direct-to-Consumer Advertising,” for December 2026. It would remove adequate provision, so a broadcast ad would have to include a full brief summary of side effects, contraindications, and effectiveness, a standard Latham & Watkins describes as long deemed “commercially unworkable in most cases.” Spending has shifted rather than stopped. National TV drug advertising fell 5.3% to $2.82 billion in the first half of 2026, and spending on GLP-1 diabetes and weight-loss drugs fell 71% (iSpot via eMarketer). Companies have said they are not radically changing course, and some have moved spending online (Endpoints, February 2026).

15e. Should Prescription Drugs Be Advertised on TV? The Cases as Their Advocates Make Them

The case that they should be: PhRMA, the drug makers' trade group, and the advertising industry argue that the ads send people to doctors. “DTC advertising provides patients with important fact-based, useful and accessible information,” its chief executive, Stephen Ubl, said when the memorandum was signed. Conditions from depression to high cholesterol often go undiagnosed or untreated, and an ad that prompts a conversation can be the first step to care. A doctor still decides what to prescribe, and the FDA already polices the ads' content. Truthful commercial speech has First Amendment protection, which the Supreme Court applied to drug marketing in *Sorrell v. IMS Health* (2011). Ending adequate provision would, in their view, be a ban adopted without Congress, and industry lawyers have called it a de facto ban.

The case that they should not be: Health and Human Services Secretary Robert F. Kennedy Jr. and FDA Commissioner Marty Makary, with many doctors and health researchers, argue that the ads drive demand

for costly drugs. “Pharmaceutical ads hooked this country on prescription drugs,” Kennedy said, and Makary said that “for far too long, these ads have distorted the doctor/patient relationship.” The American Medical Association called for a ban in 2015. The broadcast format, they argue, lets cheerful images compete with a rapid list of risks while the full facts sit on a website few viewers visit.

The evidence: both sides find support. Of the 73 most-advertised drugs of September 2015 to August 2021 that health agencies in Canada, France, or Germany had rated, only 20 were judged to offer high added benefit over existing treatments. The rest took nearly $16 billion of about $22 billion spent advertising them on television (Patel and others, *JAMA Network Open*, 2023). On the other side, Bradley Shapiro compared areas on either side of television-market borders and found that 10% more antidepressant advertising raised new prescriptions by 0.3%. It brought no rise in patients abandoning treatment, and a fall in missed workdays worth about $770 million (*American Economic Journal: Microeconomics*, 2022). Michael Sinkinson and Amanda Starc found that statin ads raised rival brands' sales as well as the advertiser's own, a sign that the ads drew new patients into treatment rather than only moving them between brands (*Review of Economic Studies*, 2019).

The course's reading, stated once: the evidence supports neither side's strongest claim. The ads do bring untreated patients to doctors, with measured benefits in the best-studied class, and they also concentrate on costly drugs that regulators abroad rate as adding little. The live question in 2026 is disclosure rather than a ban, which would face a First Amendment challenge: whether requiring the full risk summary on air is a fair condition of access to television or a ban by other means. The rule the FDA expects to propose, and any court challenge to it, will decide which.

15f. Political Advertising

Political advertising is the least restricted kind in the United States, because the First Amendment protects political speech most strongly. AdImpact counted $11.21 billion of spending and reservations for the 2025-26 election cycle as of October 27, 2026. That is a record for a midterm cycle and narrowly passes the $11.20 billion of the 2023-24 presidential cycle (TheWrap). Module 5 notes how small and short-lived the measured effects of such ads are.

Federal rules govern who pays rather than what is said. Ads must carry “paid for by” disclaimers under FEC rules, and candidates' own ads include their statement of approval. Broadcast stations must sell candidates time at their lowest rates in the 45 days before a primary and the 60 days before a general election, and may not censor a candidate's ad, even a false one. National cable networks and online platforms have no such duties and set their own policies.

Synthetic media has brought the first new rules. In September 2024 the FEC declined to write a rule on artificial intelligence, issuing instead an interpretation that its ban on fraudulently misrepresenting a candidate applies whatever the technology. States moved faster: Public Citizen counted 33 state laws on election deepfakes, of which 31 were in force after Hawaii's was struck down and Massachusetts' expired (October 21, 2026). Google has required disclosure of synthetic content in election ads since November 2023, and Meta since January 2024.

The EU went further and lost its largest sellers. Its regulation on political advertising, applied from October 10, 2025, requires labels showing who paid and how an ad was targeted, allows targeting with personal data only with explicit consent and never with sensitive data, and sets fines of up to 6% of turnover. Google had announced in November 2024 that it would stop selling political ads in the EU when the rules applied, and Meta, which announced its exit in July 2025, calling the legal uncertainty “untenable,” stopped in October 2025.

15g. Tobacco, Alcohol, and Gambling

Tobacco is the most restricted product. Cigarette ads have been off US radio and television since the federal ban described in Module 2. The Master Settlement Agreement of November 23, 1998, between 46 states and the largest cigarette makers, banned cartoon characters, billboards, transit ads, branded merchandise, and paid product placement. It also requires payments estimated at about $206 billion over its first 25 years. Since 2009 the FDA has regulated tobacco marketing under the Family Smoking Prevention and Tobacco Control Act.

Alcohol advertising is governed mostly by the industry's own codes. The beer industry's code limits ads to media whose audience is expected to be at least 73.6% of legal drinking age; the Beer Institute and the Brewers Association raised the threshold from 71.6% in July 2022 to track the census. Distillers kept spirits ads off television by voluntary agreement until 1996.

Gambling rules are changing fastest, as the table shows. The United States has no federal law on betting ads, so sportsbooks follow the American Gaming Association (AGA) code. A bill from Senator Richard Blumenthal and Representative Paul Tonko, reintroduced in March 2025, would ban sports-betting ads from 8 a.m. to 10 p.m. and during live sport; it had not been enacted as of October 31, 2026.

Place	Main rule on gambling ads	In force
United States	The AGA's voluntary code requires an expected audience at least 73.6% aged 21 or older and bars marketing aimed at college students; several states ban offers described as “risk-free”	Code updated 2023
United Kingdom	Betting firms' voluntary “whistle-to-whistle” ban on ads during live televised sport before 9 p.m.; Premier League clubs' ban on gambling sponsors on the front of match shirts	2019; from the 2026-27 season
Italy	Near-total ban on gambling advertising and sponsorship under the “Dignity Decree”	2018, with existing contracts allowed to run into 2019
Australia	No gambling ads during live sport, from 15 minutes before to 5 minutes after, between 5 a.m. and 8:30 p.m.; at most three an hour on free-to-air TV; none at venues or on uniforms; no celebrities; online ads only to logged-in, verified adults	January 1, 2027

Sources: American Gaming Association; Industry Group for Responsible Gambling; Premier League (April 2023); Italian Decree-Law 87/2018; DeucesCracked (August 2026) on the Australian law passed that month. A parliamentary inquiry chaired by Peta Murphy had recommended a full ban in 2023; critics called the new Australian rules timid, and the industry warned that bettors would move offshore.

Prediction markets, which let users trade contracts on the outcomes of events including games, have become large advertisers. Kalshi spent $41.1 million on national television in 2026 through early September, against $57.9 million by FanDuel and $111 million by DraftKings (CasinoBeats, September 4, 2026). The National Football League (NFL) barred prediction-market ads from Super Bowl LX. Whether the markets are gambling is being fought in court: in August 2026 the Ninth Circuit let Nevada enforce its gambling law against Kalshi. Module 12 covers their sponsorships of creators.

15h. Europe's Platform Rules and What Comes Next

The EU's DSA sets advertising rules for online platforms. They must label ads and show who paid for each and why it reached a user, may not target ads using sensitive data such as religion or sexual orientation, and, if very large, must keep public repositories of the ads they run. The Commission's first fine under the DSA, €120 million (about $138 million) against X on December 5, 2025, was partly for an ad repository that

lacked ads' content, topic, and payer and was hard to use. The same day it accepted TikTok's commitments to update its repository within 24 hours and include targeting data.

Two more EU measures reach advertising. The AI Act's transparency duties have applied since August 2, 2026: people must be told when they are dealing with an AI system, deepfakes must be disclosed, and systems that generate synthetic content must mark it in a machine-readable form. Providers of generative systems already on the market before that date have until December 2, 2026 to add the marking (Morgan Lewis, August 2026). Module 17 covers AI-made advertising itself.

The Digital Fairness Act, the Commission's planned law on dark patterns (designs that steer users into choices they did not intend), addictive design, influencer marketing, and unfair personalization, is expected in the fourth quarter of 2026, tentatively on November 11 (MLex, September 29, 2026). Commissioner Michael McGrath leads the work.

Module 15 Key Points

The Rules: Truth, Children, Drugs, Politics, and Vice

• The FTC's test since 1983: a claim likely to mislead a reasonable consumer about something material is deceptive; puffery is not.

• *Trump v. Slaughter* let the president remove FTC commissioners at will, overruling the 1935 precedent that had protected them.

• Only the US and New Zealand allow prescription-drug ads; a rule expected in December 2026 would require full risk summaries on air.

• Political ads reached a record $11.21 billion in the 2025-26 US cycle, with few limits on what they may say.

• Gambling ad rules are tightening abroad; Australia bans betting ads around live sport from January 2027.

• The course's reading: drug ads both inform and steer toward costly drugs; the live question is disclosure, not a ban.

Module 15 Quiz

Check Your Understanding

1. Under the FTC's deception policy, which claim is the least likely to be treated as deceptive?

A) “Clinically proven to cut colds by half,” with no clinical study behind it
B) “The best coffee in town,” painted on a café's window
C) A ticket price shown without the mandatory fees
D) A product review written by the seller's own staff and posted as a customer's

2. What did the Supreme Court decide in Trump v. Slaughter (June 29, 2026)?

A) The FTC may not seek refunds for consumers in federal court
B) The president may remove Federal Reserve governors at will
C) FTC commissioners serve fixed terms the president cannot cut short, as Humphrey's Executor held
D) The president may remove FTC commissioners at will

3. What would the FDA's expected rule on “adequate provision” change?

A) Drug ads would be banned on television before 9 p.m., when children are watching
B) Broadcast drug ads would have to carry a full brief summary of risks
C) Drug makers could no longer advertise on social media or pay influencers
D) Ads would have to state each drug's list price

4. What did Bradley Shapiro find about antidepressant advertising?

A) Ads moved patients between rival brands without changing total prescriptions, so only the advertisers gained
B) Ads raised the prices of the advertised antidepressants
C) Ads increased the side effects patients reported to the FDA
D) More ads brought more new prescriptions and fewer missed workdays, without more patients quitting treatment

5. Why did Google and Meta stop selling political ads in the EU in 2025?

A) The EU's new political-ad rules seemed legally uncertain to them
B) The Digital Services Act banned all paid political advertising online
C) EU elections were postponed that year, so parties had little reason to spend
D) Advertisers had moved their political spending to television and radio

Module 16 — Antitrust: The Cases Against the Platforms

Google, Meta, Amazon, and the Regulators Abroad

Recommended Reading / Reference

Book: David S. Evans and Richard Schmalensee, *Matchmakers: The New Economics of Multisided Platforms* (2016), two economists' plain account of why businesses that serve two groups at once, such as users and advertisers, set prices and compete differently from ordinary firms: the economics behind every case in this module

Reference: The two liability opinions against Google: Judge Amit Mehta's on search in *United States v. Google LLC* (Washington, D.C., August 5, 2024) and Judge Leonie Brinkema's on ad tech in *United States v. Google LLC* (Alexandria, Virginia, April 17, 2025), long but readable, posted with the remedies decisions on the Justice Department's case pages

16a. What Antitrust Asks of a Platform

American antitrust law rests on the Sherman Act of 1890, whose second section makes it illegal to monopolize a market, though having a monopoly is not itself illegal. The government must prove two things: that a firm has monopoly power in a market the court accepts, and that it kept or gained that power by excluding rivals rather than by offering a better product. The template for technology cases is *United States v. Microsoft*, in which the federal appeals court in Washington held in 2001 that Microsoft had protected its Windows monopoly with exclusive contracts that shut out a rival browser.

Advertising platforms make each step harder, because they are two-sided markets, a term the economists Jean-Charles Rochet and Jean Tirole formalized in 2003: a search engine or social network serves users, usually at no charge, and sells their attention to advertisers, setting the price on each side with the other in mind. A zero price to users is therefore no proof of competition, and a high price to advertisers no proof of monopoly. Tirole received the 2014 Nobel memorial prize in economics for his analysis of market power and regulation.

Two features tend to entrench a leader. Network effects make a social network more useful to each user as others join, and in search, scale feeds itself: every query shows an engine which results people choose, so the largest engine has the most data with which to improve. How high a barrier this creates is disputed: Google argued at trial that it won its users with a better product, and the government that it bought the scale that kept its product ahead.

The third charge is self-preferencing: a platform that runs a marketplace and also sells in it, or runs an auction and also bids in it, can tilt the rules toward its own business. The European Commission fined Google €2.42 billion in 2017 for favoring its own shopping service in search results, and the Court of Justice of the European Union, the EU's highest court, upheld the decision in September 2024.

Behind every case lies an argument about the standard: since Robert Bork's *The Antitrust Paradox* (1978), American courts have asked mainly whether conduct harms consumers, usually through higher prices or lower output. In 2017 Lina Khan, later chair of the FTC, argued in “Amazon's Antitrust Paradox” that this test misses the harm done by firms that keep prices low while taking control of a market. The cases below test both views, and their remedies test whether courts will break up firms or only regulate their conduct.

16b. Google Search: A Monopoly Kept by Default

The Justice Department and a group of states sued Google on October 20, 2020, over the payments that made Google the default search engine in Apple's Safari browser, on Android phones, and in browsers such as Mozilla's Firefox. After a ten-week trial in 2023, Judge Amit Mehta of the federal court in Washington ruled on August 5, 2024: “Google is a monopolist, and it has acted as one to maintain its monopoly.” He found that Google handled about 90% of general searches, nearly 95% on mobile devices, and that it paid $26.3 billion in 2021 for default placement.

The ruling reached advertisers through a second monopoly, in general search text advertising: the text ads shown with search results. Google's monopoly power, maintained by the default deals, Mehta wrote, “has enabled Google to increase text ads prices without any meaningful competitive constraint.” He rejected two other claims: Google did not, he found, monopolize a broader market for all search advertising, which would include ads on Amazon and social media, and the states had not shown that Google used Search Ads 360, its tool for managing campaigns across search engines, to hold back Microsoft's Bing.

Remedies followed on September 2, 2025, and they were milder than the government had asked. Mehta refused to order the sale of the Chrome browser and declined to ban payments for default placement, on which browser makers and phone companies depend. He barred exclusive contracts for Search, Chrome, the Google Assistant, and the Gemini app. He ordered Google to share parts of its search index and data on how users interact with results with qualified competitors, and to offer them search results and search text ads to resell, for six years. Generative artificial intelligence (AI) shaped the outcome: Mehta wrote that “the emergence of GenAI changed the course of this case,” because chatbots had become potential rivals since the trial.

The final judgment was entered on December 5, 2025, and a technical committee to help enforce it was named in January 2026. Both sides appealed to the federal appeals court in Washington, Google asking it to reverse the finding of monopoly. The Justice Department and the states ask it to uphold liability and to send back Mehta's refusal to ban default payments; they did not challenge his rejection of a Chrome sale (Courthouse News, July 2026). Google's opening brief was filed on May 22, 2026, the government's on July 28, and the last brief was due on October 29; the court had not heard argument as of October 31, 2026.

16c. Google Ad Tech: Broker, Exchange, and Auctioneer

The second federal case, filed on January 24, 2023, concerned the machinery that sells display ads on websites (Module 10 explains how it works). Google owned the leading ad server for publishers, DoubleClick for Publishers (DFP); the leading ad exchange, AdX; and the leading buying tools for advertisers. In the same auction it could act for the seller, run the marketplace, and act for the buyer. The conduct at issue included “Last Look,” which let AdX see rival exchanges' bids before making its own, and Unified Pricing Rules, which stopped publishers from setting a higher minimum price for AdX than for other exchanges.

After a three-week trial in 2024, Judge Leonie Brinkema of the federal court in Alexandria, Virginia, ruled on April 17, 2025 that Google had illegally monopolized two markets, publisher ad servers and ad exchanges for open-web display ads, and had illegally tied DFP to AdX. She rejected a third claim, that Google monopolized ad networks serving advertisers, because the government had not proved that such a market existed. Her findings on market shares and fees belong with the mechanics, in Module 10.

The government asked for a breakup: the sale of AdX, the release of the ad server's final auction code as open source, and a sale of DFP if those measures failed. On September 2, 2026 Brinkema refused all three

in a two-page order, and her 106-page opinion was unsealed on September 16. Structural remedies, she wrote, were “neither realistic nor needed.” The government had identified no viable buyer for AdX, a forced sale had no precedent as a cure for tying, and Google's estimates of the time needed to pull the systems apart ran far longer than the government's.

She chose conduct rules instead, for six years rather than the 15 the government sought, applying worldwide and taking effect 60 days after the final judgment is entered. Prebid, the open-source software publishers use to run header bidding, must be able to request real-time bids from AdX, and AdX must send real-time bids to rival publisher ad servers as it does to DFP. Google may not bring back First Look, which gave AdX the first chance to bid, or Last Look; it must retire Unified Pricing Rules and let publishers export their data. A court-appointed monitor with access to Google's documents and code, a three-member technical committee, and an internal compliance officer will police the rules.

As of October 31, 2026, the final judgment had not been entered. The two sides filed competing drafts on October 2, disagreeing over how quickly AdX must connect to Prebid (six months for the government, twelve for Google), whether Google may phase in the rules outside the United States, and how much power the monitor should have (AdExchanger, October 5, 2026). Google says it will appeal the finding of liability, while the Justice Department called the remedies a “significant victory” and said it was reviewing its legal options (Reuters, September 16, 2026). A separate ad tech suit by Texas and other states, in federal court in Texas, was postponed in July 2025 until the Virginia court enters its final judgment.

16d. Meta and Amazon

The FTC sued Facebook, now Meta, in December 2020, arguing that it bought Instagram in 2012 for $1 billion and WhatsApp in 2014 for about $19 billion to neutralize rising rivals, and so kept a monopoly in “personal social networking.” After a trial in April and May 2025, Judge James Boasberg of the federal court in Washington ruled for Meta on November 18, 2025. He framed the question in the present tense: “Whether or not Meta enjoyed monopoly power in the past, though, the agency must show that it continues to hold such power now. The Court's verdict today determines that the FTC has not done so.”

The market decided the case: Boasberg included TikTok and YouTube, finding that the apps now compete for the same time, much of it spent watching videos the apps recommend, so Meta's share of that market fell short of monopoly. The FTC announced on January 20, 2026 that it would appeal, and as of October 31, 2026 the appeal was being briefed in the appeals court in Washington, where the Washington Legal Foundation filed a brief supporting Meta on August 27, 2026.

The FTC and 17 states sued Amazon on September 26, 2023, in federal court in Seattle. The complaint says Amazon keeps monopolies in online superstores and in services for marketplace sellers by punishing sellers who offer lower prices elsewhere and by making the Prime badge depend on using Amazon's own fulfillment service. It also says Amazon degrades its search results by replacing relevant listings with paid ads, raising sellers' costs. Judge John Chun has set trial for March 29, 2027, and on October 9, 2026 he refused Amazon's request to move it to June 2027 (MLex). The FTC's separate suit over Amazon's sponsored-ads auctions is covered in Module 10.

A different FTC case, over how Amazon enrolled shoppers in Prime and made it hard to cancel, ended on September 25, 2025 with a $2.5 billion settlement: a $1 billion civil penalty and $1.5 billion in refunds.

16e. Europe, Britain, and Elsewhere

Europe's rule, Article 102 of the Treaty on the Functioning of the European Union, forbids a dominant firm to abuse its position. The European Commission investigates and decides, and it can fine a company up to 10% of its worldwide turnover before any court rules; the company may then appeal to the EU's courts. On September 5, 2025, the Commission fined Google €2.95 billion (about $3.4 billion) for favoring AdX since at least 2014, through both its publisher ad server and its buying tools for advertisers. Its preliminary view was that only a divestment of part of Google's ad tech business would end Google's “inherent conflicts of interest,” and it gave Google 60 days to propose a remedy.

Google's answer of November 14, 2025 kept the business whole: publishers could set different minimum prices for different bidders, and Google's tools would work more readily with rivals'. Teresa Ribera, the Commission's executive vice-president for competition, had said that Google “must now come forward with a serious remedy.” The Commission was still assessing the offer when Brinkema chose conduct rules over a sale (The Next Web, September 2, 2026), and it had announced no remedy decision as of October 31, 2026; Google said it would appeal the fine.

Other governments have acted too, and Britain has gone furthest toward writing rules in advance instead of proving a violation afterward: its Digital Markets, Competition and Consumers Act 2024 lets the CMA give a firm “strategic market status” and then impose conduct requirements on it without a finding of abuse.

Authority	Step	Date	What it requires or seeks
Britain's CMA	Strategic market status for Google in general search and search advertising	October 2025	The first designation under the 2024 act
Britain's CMA	First conduct requirements imposed	June 3 and June 17, 2026	Publishers may keep their content out of Google's AI features and get attribution, with nine months to comply; fair ranking; data portability. Choice screens went back to consultation on September 23, 2026
Japan Fair Trade Commission	Cease-and-desist order	April 15, 2025	Google must stop requiring Android phone makers to preinstall its search and Chrome in favored positions; five years of monitoring
Canada's Competition Bureau	Case before the Competition Tribunal	Filed November 28, 2024	Seeks the sale of two Google ad tech tools and a penalty; the Tribunal dismissed Google's constitutional challenge on March 4, 2026

Sources: CMA decisions and consultations, 2025-26; Japan Fair Trade Commission, April 15, 2025; Competition Bureau of Canada, November 28, 2024; the Competition Tribunal's ruling of March 4, 2026, as reported. Status as of October 31, 2026.

These rules aim at search and phones rather than at ad prices, but they change the terms on which Google's ad business runs. The British publisher requirement, for example, bears directly on the traffic that publishers lose to AI answers (Module 17).

16f. Who Enforces in Washington

The cases have outlived the administrations that brought them: the search and Meta suits were filed under President Trump's first administration, and the ad tech and Amazon suits under President Biden. At the Justice Department, Gail Slater, who took charge of the Antitrust Division in March 2025, resigned on February 12, 2026. Omeed Assefi led the division in an acting capacity until June, with Associate Attorney General Stanley Woodward overseeing its cases.

President Trump nominated Adam Candeub, the FCC's general counsel, to lead the division. The Senate Judiciary Committee advanced the nomination 12-10 along party lines on September 24, 2026 (MLex), and no vote by the full Senate had been reported as of October 31, 2026. The department said he could advise on antitrust matters before confirmation (Global Competition Review, September 17, 2026). The FTC's own leadership, and the Supreme Court's 2026 ruling on the removal of its commissioners, are covered in Module 15.

16g. Monopolies That Overcharge, or the Cheapest Way to Reach a Customer? The Cases as Their Advocates Make Them

The case that they are monopolies that overcharge: The Justice Department and the state attorneys general make it, with the European Commission, the FTC under chairs of both parties, scholars such as Lina Khan and Tim Wu, advocacy groups such as the American Economic Liberties Project, and many news publishers. Two federal judges, after full trials, have found that Google held monopolies that bear directly on advertisers, including the power to raise text-ad prices. The default payments are their exhibit: a firm pays billions a year to keep a position only when the position earns more.

In ad tech, they argue, Google acted for buyer and seller while owning the exchange between them, a conflict that the former ad tech executive Dina Srinivasan, writing in 2020, compared with conflicts that securities law forbids in financial markets. Amazon, on the FTC's account, has turned its search results into a toll that sellers must pay to be seen. Network effects and data keep rivals from entering, so high prices persist, and advertisers pass them on in the prices of what they sell.

The case that they are the cheapest way yet built to reach a customer: The companies make it, with economists and lawyers in the consumer-welfare tradition, such as those at the International Center for Law & Economics, and with business groups. Advertisers pay what an auction sets, often only when someone clicks or buys, and can measure what they get; a bakery can start with a few dollars a day, which no newspaper or broadcaster ever offered. Targeting has cut the cost of reaching a likely buyer, and the platforms keep adding impressions and advertisers, the opposite of the restricted output that monopolies produce.

Rivals are many and growing: TikTok, YouTube, the stores' ad businesses (Module 9), and now chatbots (Module 17). The courts, they note, agreed in part: Boasberg found that Meta competes with TikTok and YouTube, Mehta rejected a monopoly of search advertising as a whole, and Brinkema found no monopoly on the advertisers' side of display. Google said a forced sale of AdX “would have made it harder for small businesses” to reach customers (Reuters, September 16, 2026), and the Washington Legal Foundation told the appeals court that “the social-media market is a thriving and dynamic one and Meta is no monopolist.”

The evidence: The trial records support a narrow version of the first case and much of the second: two judges found monopoly power at specific chokepoints, and the Commission found abuse at one of them, while every claim of a wider monopoly failed. The best direct price comparison is older: Britain's CMA reported in its 2020 market study that Google's prices for like-for-like search terms ran 30% to 40% above Bing's, a gap that could reflect a higher value per click as well as market power.

Rising prices per ad, as at Meta (Module 8), can reflect rising demand as well as power. And the experiments in Module 5, which find that many campaigns earn less than advertisers believe, show that advertisers can overpay without any monopoly, so overpayment alone does not settle the question.

The course's reading, stated once: The record supports chokepoint monopolies, not a monopolized advertising market. Where the courts found monopoly, Google's conduct raised prices or shut out rivals; elsewhere advertisers have real choices, and the record does not show them overcharged by a monopoly; both can be true at once. Whether conduct rules, rather than breakups, lower what advertisers and publishers pay is the question the decrees and the appeals now under way will begin to answer.

Module 16 Key Points

Antitrust: The Cases Against the Platforms

• Judge Mehta found in 2024 that Google monopolized search text ads and could raise their prices, but he ordered conduct rules, not a breakup.

• Judge Brinkema let Google keep AdX, imposing six years of worldwide conduct rules policed by a court-appointed monitor; the final judgment awaited entry.

• The FTC lost its Meta case in November 2025 because the court counted TikTok and YouTube as competitors; the FTC has appealed.

• Europe fined Google €2.95 billion and prefers a divestiture; Britain now sets conduct rules for Google's search business in advance.

• *FTC v. Amazon*, over sellers' prices, the Prime badge, and ads in search results, is set for trial on March 29, 2027.

• The course's reading: Google held chokepoint monopolies, but no court has found advertising as a whole monopolized.

Module 16 Quiz

Check Your Understanding

1. What did Judge Mehta find about general search text ads in August 2024?

A) They belong in one market with Amazon's and social ads, so Google had no monopoly
B) Bing's text ads had become cheaper and more effective than Google's on mobile devices
C) Google's default deals let it raise their prices without any meaningful competitive constraint
D) Their prices were set by auctions Google could not influence

2. Why did Judge Brinkema refuse to order the sale of AdX in September 2026?

A) No viable buyer had been identified, and a forced sale had no precedent as a cure for tying
B) She concluded that Google had not monopolized ad exchanges after all, reversing her 2025 ruling
C) The European Commission had already ordered the same sale
D) Google agreed to sell DFP instead

3. Why did the FTC lose its monopoly case against Meta in November 2025?

A) The court ruled that deals the FTC had cleared in 2012 and 2014 could not be challenged later
B) Meta sold WhatsApp before the ruling
C) A Supreme Court ruling on the FTC's commissioners ended the case
D) The court counted TikTok and YouTube as Meta's competitors

4. In its decision of September 5, 2025, what did the European Commission say was the only way to end Google's conflicts of interest in ad tech?

A) Different minimum prices for different bidders
B) A fine of up to 10% of worldwide turnover
C) A divestment of part of Google's ad tech business
D) Six years of conduct rules and a monitor

5. What does Britain's strategic market status regime allow the CMA to do?

A) Order the sale of a business without first holding a hearing
B) Impose conduct rules on a firm without proving abuse
C) Approve in advance the prices Google charges advertisers
D) Fine Google for past abuse, as the European Commission did in 2025

Module 17 — AI and the Next Shift

Machine-Made Ads, Ads Inside Answers, and Shopping Agents

Recommended Reading / Reference

Book: Ajay Agrawal, Joshua Gans, and Avi Goldfarb, *Prediction Machines: The Simple Economics of Artificial Intelligence* (2018), three Toronto economists' short account of AI as a fall in the cost of prediction, which explains why targeting and bidding were the first parts of advertising that machines took over

Report: McKinsey & Company, *The Agentic Commerce Opportunity: How AI Agents Are Ushering in a New Era for Consumers and Merchants* (October 17, 2025), the most-cited forecast of shopping by AI agents, with its warning about what agents may do to advertising and retail media

17a. Machines That Buy the Ads

Artificial intelligence (AI) entered advertising through prediction long before it could write or draw. The auctions in Module 10 already rest on machine-learned estimates of how likely each person is to click or buy, and the next step automated the advertiser's own decisions. Google's Performance Max, opened to all advertisers in 2021, and Meta's Advantage+ campaigns, introduced in 2022, let an advertiser set a goal and a budget and leave the system to choose the audience, the placements, the bids, and which combination of images and text to show.

Google carried the approach into search keywords with AI Max for Search, announced on May 6, 2025, which matches ads to queries the advertiser never listed and rewrites the ad text to fit them. Google says that advertisers who turn it on typically get 14% more conversions or conversion value at a similar CPA or ROAS, and 27% more for campaigns that had relied mostly on exact and phrase keyword matching, figures from Google's own data.

Meta set the boldest goal. The Wall Street Journal reported on June 2, 2025 that Meta aimed, by the end of 2026, to let an advertiser supply a product image and a budget target and have its AI do the rest. The AI would produce the entire ad, including imagery, video, and text, choose whom to show it to on Facebook and Instagram, and suggest how to spend the budget. Different viewers would see different versions of the same ad, varied in real time by factors such as location.

On its call of October 21, 2026, Meta did not claim to have reached that goal. Its finance chief, Susan Li, said that “nearly 10 million advertisers” had used at least one of its generative AI creative tools in the third quarter, and that use of its video-generation features grew 25% from the second quarter. She described an AI business assistant that lets marketers plan, create, and optimize campaigns in conversation, with the advertiser approving each step.

Meta also reported what its models add: a new ad-ranking model, it said, raised conversions by 1.7% for campaigns seeking sales off Meta's apps and by 9.6% for ads promoting apps, and advertisers who took its assistant's recommendations got almost 5% more conversions at the same cost per conversion. These are Meta's own measurements, and the last compares advertisers who chose to adopt the recommendations with similar advertisers who did not, so it may reflect who chose as much as what the tool did.

17b. Machine-Made Ads

Generative AI now makes the ads as well as placing them. In a survey of buyers published in July 2025, the IAB found that about 30% of digital video ad creative was built or enhanced with generative AI, a share expected to reach about 40% in 2026, and that nearly 90% of advertisers used or planned to use it for video. Amazon introduced Creative Agent, which builds display, video, audio, and streaming-TV ads, at its unBoxed conference on November 11, 2025.

Big brands have tested how much of it audiences will accept: in November 2024 Coca-Cola remade its 1995 commercial “Holidays Are Coming” with AI, working with the studios Silverside AI, Secret Level, and The Wild Card, and drew wide criticism. The company ran two more AI-made holiday ads on November 3, 2025, which were criticized again. Coca-Cola called the 2024 film a commercial success and said a team of five specialists generated more than 70,000 video clips in 30 days for the new ones (NBC Bay Area, November 2025).

AI companies have also become large buyers of ordinary advertising: at Super Bowl LX (its date and prices are in Module 11), OpenAI advertised its Codex coding tool, Google its Gemini assistant, Amazon its Alexa+ assistant, and Meta the AI glasses it makes with Oakley. Anthropic's ad carried the line “Ads are coming to AI. But not to Claude,” a pledge to keep its assistant free of advertising and a contrast with OpenAI's plans for ChatGPT (17d).

The agency groups have answered with AI platforms of their own: WPP Open, Publicis's Marcel, Omnicom's Omni, Dentsu.Connect, and Havas's Converged.AI (The Media Leader, April 16, 2026). What cheaper creative and automated buying do to the agencies' business is covered in Module 7, and the rules for labeling AI-made content in Module 15.

17c. Ads Inside the Answers

Search is where AI changes most directly what an ad is: Google began placing ads inside AI Overviews, the generated summaries at the top of its results, on mobile in the United States in October 2024. On December 19, 2025 it extended them to English-language searches in 11 more countries, among them Australia, Canada, India, and Nigeria. AI Mode, its conversational search, carries ads too, and advertisers using Performance Max, broad match, or AI Max became eligible for them automatically.

On its call of October 28, 2026, Alphabet described both the opportunity and the cost. Sundar Pichai said query growth was “the highest it's been in the last five years” and that the Gemini app had reached 1 billion users. Philipp Schindler, Google's chief business officer, said Google was scaling ads in AI Overviews and AI Mode to more users and queries, taking AI Mode ads to other countries, and adding multi-turn ads, which follow a conversation, in the United States.

He also named the cost: “the strong growth in AI Mode queries had an impact because of the lower monetization we were seeing in AI Mode.” His case for patience was that many AI Mode queries come early in a purchase, letting Google “capture new high value commercial intent,” and that click-through rates on sponsored and organic product listings inside AI Mode were already similar. How Google's search revenue fared that quarter is in Module 8.

17d. Ads in ChatGPT

OpenAI turned ChatGPT into an ad seller in 2026, in the steps below.

Date, 2026	Step	Source
January 16	Announced a test for logged-in adults on the Free and Go ($8 a month) tiers in the United States; previewed a product-carousel format	OpenAI; SiliconANGLE
February 9	First ads shown, labeled as sponsored and set apart from the answer	OpenAI; TechCrunch
Late March	The US pilot passed $100 million of annualized revenue within six weeks, with more than 600 advertisers	OpenAI
May	Self-serve Ads Manager in beta in the US; cost-per-click bidding added to CPM pricing; a $50,000 minimum dropped; Dentsu, Omnicom, Publicis, and WPP as agency partners	Axios, May 5
August 31	Ads reached a $1 billion annualized run rate; Ads Manager opened in India, Europe, the Middle East, and North Africa	OpenAI; Reuters
September	AI tools for advertisers; tests of agents sponsored by businesses, with some US advertisers	Reuters
October 21	26 more countries, 89 in all; more than 100,000 advertisers since launch	OpenAI; Reuters

Plus, Pro, Business, Enterprise, and Education accounts carry no ads. Revenue run rates are OpenAI's statements, not audited figures.

OpenAI's launch announcement set rules meant to keep ads apart from answers. “Ads do not influence the answers ChatGPT gives you,” it said. Advertisers receive aggregate results such as views and clicks, not chats, chat history, or memories. Ads are withheld from users who say or are predicted to be under 18 and are “not eligible to appear near sensitive or regulated topics like health, mental health or politics.” Free users may opt out of ads in exchange for fewer free messages.

Investor presentations reported by Axios in April 2026 projected OpenAI's ad revenue at $2.5 billion for 2026, rising to $100 billion in 2030. eMarketer's Nate Elliott said the August run rate left OpenAI well short of its target for the year (Reuters, August 31, 2026).

Not every AI company has followed: Perplexity, an AI search engine that began trying ads in 2024, stopped taking new advertisers in October 2025 and has since said it is stepping away from ads in its answers. One of its executives told the Financial Times that with ads, “a user would just start doubting everything” (eMarketer, 2026).

Forecasters expect the format to grow quickly from a small base: WPP Media put spending on ads in generative search at $5.1 billion in 2026, rising to more than $100 billion by 2030 (June 2026), and eMarketer forecasts that ads in AI search will take 13.6% of US search ad spending by 2029. These are projections; the worked example shows how far the format has to go.

Worked example — What an Answer Must Earn

Inputs: Google's statement in March 2025 that it sees “more than 5 trillion searches” a year, and Alphabet's Google Search & other revenue for 2024, $198.1 billion, mostly search ads but also ads in Gmail, Maps, and the Play store. ChatGPT's inputs are the 2.5 billion prompts a day that OpenAI said it received (Axios, July 2025) and the August 31 run rate in the table above. Step 1, Google's revenue per thousand searches: $198.1 billion ÷ 5,000 billion × 1,000 = $39.62, a ceiling, because Google said “more than.” Step 2, ChatGPT's prompts in a year: 2.5 billion × 365 = 912.5 billion. Step 3, ChatGPT's ad revenue per thousand prompts: $1 billion ÷ 912.5 billion × 1,000 = $1.10, also a ceiling, since prompts have grown since mid-2025. Step 4, the gap: $39.62 ÷ $1.10 = about 36. Step 5, what parity would take: 912.5 billion prompts × $39.62

per thousand = about $36 billion a year at mid-2025 volume, more than the $25 billion that OpenAI's presentations projected for 2028, as Axios reported. Result: about seven months after its first ads, ChatGPT earned roughly one thirty-sixth as much per query as Google's search. The comparison counts every query on both sides; because ChatGPT shows ads only on two tiers, in some countries, and away from sensitive topics, it earns more per eligible prompt. The same $39.62 is the bar that Google's own AI Mode must clear for each thousand queries it takes from classic search, and Google says it does not yet (17c).

17e. Fewer Clicks for the Open Web

Answers that satisfy a reader on the results page send fewer readers onward, and the open web's advertising depends on those visits. The Pew Research Center tracked the browsing of 900 American adults in March 2025 and reported on July 22, 2025 that an AI summary appeared on about one Google search in five. When one appeared, users clicked a traditional search result on 8% of visits, against 15% when none did, and clicked a link inside the summary on just 1%. They ended their browsing after 26% of pages with a summary, against 16% without; Google disputed the study's method.

Fewer visits mean fewer impressions for publishers to sell through the exchanges described in Module 10, and less leverage against the platforms that summarize their work. Penske Media, the owner of Rolling Stone and Variety, sued Google in September 2025 over AI Overviews, and Britain's CMA has required Google to let publishers keep their content out of its AI features (Module 16).

17f. Shopping Agents

The largest change would come if software agents, not people, did the shopping. The pieces arrived in 2025: in April Amazon began testing Buy for Me, which lets some US users have the Amazon app buy from other brands' websites. On September 16 Google published the Agent Payments Protocol, with more than 60 partners including American Express, Mastercard, and PayPal, so that agents can pay on a user's signed instructions.

On September 29, 2025 OpenAI launched Instant Checkout, letting ChatGPT users buy from Etsy sellers without leaving the chat, on an Agentic Commerce Protocol written with the payments company Stripe. Merchants pay a fee on sales, and OpenAI said product results would remain “organic and unsponsored.” Walmart agreed on October 14, 2025 to sell through it. Google answered on January 11, 2026 with a Universal Commerce Protocol for checkout inside AI Mode and the Gemini app, with Walmart, Shopify, Etsy, and Target among its early adopters.

Use has lagged the announcements: in March 2026 OpenAI scaled back Instant Checkout, shifting shopping to merchants' own apps inside ChatGPT while keeping the protocol, after The Information reported weak adoption. Forrester counted only about a dozen Shopify merchants live; Shopify put the number at about 30. Schindler said in October 2026 that retailers were adopting Google's protocol rapidly, without giving figures.

Forecasts run far ahead of use: Morgan Stanley estimated in December 2025 that agents could account for up to $385 billion of US online sales by 2030, 10% to 20% of US e-commerce. McKinsey projected in October 2025 that agents could orchestrate $900 billion to $1 trillion of US consumer retail sales, and $3 trillion to $5 trillion worldwide, by 2030. It also warned that agents may “bypass traditional ad channels” and weaken retail media, while impartial “sponsored smart suggestions” could survive.

17g. Will AI Agents Shrink Advertising, or Hand the Platforms Even More of It? The Cases as Their Advocates Make Them

The case that agents will shrink advertising: McKinsey's analysts and many publishers make it, and the AI companies that refuse ads lend it support. Advertising works on people's attention and memory (Module 4), and an agent has neither: it compares price, specifications, delivery times, and reviews against its user's instructions. If agents do the buying, sellers will spend less on persuasion and more on accurate product data, prices, and reviews. The shopping searches that retail media sells (Module 9) would be made by software that can ignore sponsored listings.

Answers that end the search, they add, already cut the visits that fund the open web's ads, as Pew found. And an assistant paid to recommend puts at risk the trust that makes people use it. That was the reason Perplexity gave for leaving ads, and the point of Anthropic's Super Bowl pledge.

The case that agents will hand the platforms even more of it: Google, Meta, and OpenAI make it, and forecasters such as WPP Media and eMarketer project it. Automation lowers the cost of making and placing an ad, which brings in more and smaller advertisers, and better prediction raises what each ad earns. Whoever owns the assistant owns the moment of decision and can sell a place in it, as Google sold space beside its results, and Schindler calls agentic experiences “clearly additive.”

Automated buying also runs on the platforms' logged-in data, which advertisers cannot take elsewhere. So the more the machine decides, the more of each budget, and of the work agencies once did (Module 7), stays with the platform. Meta's adoption figures and OpenAI's first-year advertiser count are, for this side, evidence that advertisers follow the machines.

The evidence: So far the money fits the second case better, since Google's search revenue kept growing through the third quarter of 2026 as AI answers spread (Module 8), and Meta, which set the boldest automation goal, grew its ad revenue faster than Google or Amazon did that quarter (Modules 8 and 9). But the costs that support the first case are visible too: lower earnings per AI answer, the gap per query in the worked example, fewer clicks for publishers, Perplexity's exit, and the retreat of the largest test of agent checkout. The course found no study that measures how agents change what shoppers buy or what advertisers spend, so both cases rest on forecasts.

The course's reading, stated once: Through 2030, AI looks more likely to concentrate advertising than to shrink it. The assistants, the logged-in data, and the automated buying belong to a few platforms, and their ad revenue is still growing. Each AI answer will carry less advertising than a classic search for some time, so the platforms must sell more answers or charge more for each, and publishers and agencies bear the cost of the transition. The scenario in which agents bypass advertising needs agents that buy at scale, and as of October 2026 they did not.

Module 17 Key Points

AI and the Next Shift

• Nearly 10 million advertisers used Meta's generative AI tools in the third quarter of 2026, but fully automated campaigns remained a goal.

• Google says an answer in AI Mode earns less than a classic search, even as AI Mode queries grow fast.

• OpenAI says ChatGPT's ads reached a $1 billion annual run rate within seven months of the first ads.

• Pew found that users clicked through about half as often when Google showed an AI summary, cutting visits to publishers.

• Shopping agents lag their forecasts: OpenAI scaled back Instant Checkout in March 2026 for lack of use.

• The course's reading: AI is more likely to concentrate advertising in a few platforms than to shrink it, at least through 2030.

Module 17 Quiz

Check Your Understanding

1. What did the Wall Street Journal report in June 2025 that Meta aimed to do by the end of 2026?

A) Stop selling ads to Instagram users under 18, starting in the United States
B) Replace its ad auction with fixed monthly prices for small businesses
C) Hand its European ad sales to one of the large agency holding companies
D) Have AI make and target whole ads from a product image and a budget

2. What did Google's Philipp Schindler say about AI Mode on Alphabet's October 2026 call?

A) Strong growth in AI Mode queries weighed on results because AI Mode monetized at a lower rate
B) AI Mode answers earn more than classic searches because they come later in a purchase
C) Google will keep AI Mode free of ads until 2027
D) Ads in AI Mode are limited to the United States

3. Which is one of the rules OpenAI set for ads in ChatGPT?

A) Advertisers may see a user's chat history once the user clicks an ad
B) Ads run on every tier, including Pro and Enterprise
C) Ads may not appear near sensitive topics like health, mental health, or politics
D) Advertisers may pay to have their products named inside the answer, marked as sponsored

4. In the worked example, about how many times more per thousand queries did Google earn from search than ChatGPT earned from ads?

A) About 4 times
B) About 36 times
C) About 360 times
D) Roughly the same, once only ad-eligible prompts are counted

5. What happened to OpenAI's Instant Checkout?

A) It became Walmart's main online sales channel
B) Britain's CMA ordered it closed
C) It was scaled back in March 2026 after weak adoption, though its protocol stayed
D) It was merged into Google's Universal Commerce Protocol

Module 18 — Where Things Stand, and What to Watch

The Argument, the Contradictions, and the Calendar

Recommended Reading / Reference

Book: Michael Schudson, *Advertising, the Uneasy Persuasion: Its Dubious Impact on American Society* (1984), a sociologist's case that advertising is less powerful than either its makers or its critics claim, a counterweight to both

Report: Interactive Advertising Bureau and PwC, *Internet Advertising Revenue Report*, published each April, the yearly count of American online advertising by format, the first place to see whether search, social, and retail media keep their 2025 order

18a. The Argument on One Page

Advertising is the price of most free media, and the reader pays it in attention and personal data, then again at the till. For a century and a half the business sold space and time in bulk: a column inch, a page, a thirty-second spot seen by whoever happened to be watching (Module 2). Since 1994 it has learned to sell a single person's attention, one impression at a time, by automated auction (Modules 3 and 10). That change moved the money to whoever controls a logged-in audience and knows what it buys.

The winners are therefore not the old media owners but four kinds of gatekeeper: the search box, the social feed, the video player, and the store (Modules 8, 9, and 11). Each owns its audience's attention and its own data about that audience, and each runs the auction in which advertisers bid for it. The agencies that once stood between advertisers and media have become smaller middlemen, squeezed by fees, in-house teams, and software (Module 7).

The evidence on whether advertising works is weaker than the industry's growth implies. The best experiments find that much advertising, especially the last dollars of large, established advertisers, returns less than it costs, while ads aimed at new and light buyers, and at people searching for something to buy, can pay well (Module 5). Spending keeps rising anyway, because measurement is costly, because the platforms measure their own results, and because the advertising that does pay is very profitable for those who sell it (Modules 13 and 9).

The rules have tightened, but unevenly. Europe regulates consent, transparency, and the largest platforms' conduct directly; the United States protects children, polices deception, and has stepped up enforcement on drug ads, and many of its states regulate election deepfakes, but it has no federal privacy law (Modules 14 and 15). American courts found Google a monopolist twice and chose rules of conduct over breakups (Module 16). AI now changes both who makes ads and where they appear, moving them from pages people visit into answers and, perhaps, into the choices of software agents that shop (Module 17).

18b. Five Contradictions

No single module states these; each draws on several.

1. Record spending, weak returns at the margin. The world has never spent more on advertising (Module 1), yet the largest studies find that most advertisers would lose little by spending less (Module 5). The two coexist because an advertiser's average return can be positive while its last dollar's return is negative, and

because the sellers' measurements tend to credit ads with sales that would have happened anyway (Module 13). Spending is set by habit, competition, and the fear of losing share, not only by measured returns (Module 6).

2. Privacy rules that helped the biggest trackers. Browsers, Apple, Europe's courts and regulators, and American states have cut the tracking of people across other companies' sites and apps (Module 14). The cuts fell hardest on independent ad-tech firms and the open web, whose targeting relied on third-party data (Module 10). The platforms whose users log in, and the retailers who know what their customers bought, lost the least and grew fastest (Modules 8 and 9). Rules meant partly to restrain the largest data collectors made their own data more valuable.

3. Monopolies found, revenues rising. Two American courts found that Google held unlawful monopolies, one in search text ads and one in the tools that sell display ads (Module 16). Google's advertising revenue kept growing through both cases (Module 8). Meanwhile the FTC lost its case against Meta because the court counted TikTok and YouTube as competitors for users' time. The market is concentrated at its chokepoints and contested in the fight for attention, and the remedies so far have been rules of conduct, which leave the companies' size untouched.

4. An industry built on measurement that cannot agree on its numbers. Advertising's selling point since 1994 has been that it can be counted (Module 3). Yet television has rival currencies, viewability proves only that an ad could have been seen, and fraud estimates differ by orders of magnitude depending on who counts (Module 13). Retail media sells itself on closed-loop measurement, linking ads to sales, but the retailer that sells the ads also does the counting (Module 9). The more precisely ads can be counted, the more depends on who keeps the count.

5. AI makes ads cheaper to make and harder to place. Generative tools cut the cost of making and testing ads, which squeezes the agencies that billed for that work (Modules 7 and 17). The same technology answers questions directly, so fewer people click through to the websites that carry display ads, and the new space for ads sits inside answer boxes run by a handful of companies (Module 17). Cheaper ads and scarcer places to show them both favor the firms that own the answers.

18c. What to Watch

Date	Event	Why it matters
November 9, 2026	The Trade Desk's third-quarter results	Whether the largest independent buyer of open-web ads can hold its business against Amazon's cheaper DSP and the platforms' closed systems (Module 10)
November 11, 2026 (tentative)	The European Commission's proposal for a Digital Fairness Act	The first EU-wide attempt to write rules for influencer marketing, dark patterns, and personalized offers (Modules 12 and 15)
November 18, 2026	The ANA's programmatic benchmark for the third quarter	Whether the share of advertisers' money that buys a real, viewable impression keeps rising under buyers' pressure (Module 10)
December 2, 2026	The EU AI Act's deadline for machine-readable marks on AI-made content from systems already on the market	The first legal test of whether synthetic ads can be labeled at scale (Modules 15 and 17)
December 2026 (expected)	The FDA's proposed rule ending the “adequate provision” standard for broadcast drug ads	Whether full risk summaries make television drug ads too long and costly to run (Module 15)

Date	Event	Why it matters
December 2026 (expected)	WPP Media's *This Year Next Year* forecast	The first full revision of 2026's growth and the first forecast for 2027 (Module 1)
By December 31, 2026	Publicis's purchase of LiveRamp expected to close	Whether owning a data-matching company gives an agency group an edge the platforms cannot copy (Module 7)
January 1, 2027	Australia's limits on gambling ads; Oklahoma's and Louisiana's privacy laws take effect	A test of partial bans against the full ban reformers wanted, and two more states in the American patchwork (Modules 14 and 15)
Late January and early February 2027	Alphabet's, Meta's, and Amazon's fourth-quarter results	Whether AI answers dent search revenue, and whether Amazon's third-quarter ad slowdown came only from moving Prime Day (Modules 8, 9, and 17)
February 14, 2027	Super Bowl LXI, on ABC and ESPN	Whether the year's largest single audience still commands a rising price as viewing fragments (Module 11)
March 1, 2027	Netflix's ad plan reaches nine more European countries, with Colombia, Peru, and New Zealand that month	Whether streaming's ad tiers can scale abroad while prices per viewer fall (Module 11)
March 29, 2027	The FTC's antitrust trial against Amazon	A trial that will test, among other claims, whether the ads in Amazon's search results raise sellers' costs and degrade results for shoppers (Modules 9 and 16)
About April 2027	The IAB and PwC report on 2026 US internet ad revenue	Whether social and retail media keep outgrowing search as AI answers spread (Modules 1 and 3)
No date set	The final judgment and appeals in Google's ad tech case; the appeals in the search case; the EU's remedy decision; the FTC's appeal against Meta; its suit over Amazon's ad auctions	Whether courts move from conduct rules toward structural change, and whether auction design becomes a matter of consumer-protection law (Modules 10 and 16)

Dates as announced or reported by October 31, 2026; the sources are in the modules named.

18d. The Last Word

Every ad a reader sees has passed through a chain of decisions: an advertiser's budget, an agency's plan or a platform's software, an auction lasting a fraction of a second, and a measurement that will be used to justify the next budget. At almost every link, the party doing the measuring is also the party being paid. That is not proof of waste, and the best ads inform buyers, launch products that would otherwise fail, and pay for journalism, video, and services that billions use without paying cash.

The reader's best defense is the one the experiments teach: ask what would have happened without the ad. A brand that charges more for the same product, a feed tuned to hold attention, a sponsored result placed above an organic one, and an answer box that carries an ad are each worth what they deliver beyond the alternative. Advertisers who run that test spend differently, and readers who ask that question buy differently.

Module 18 Key Points

Where Things Stand, and What to Watch

• Advertising moved from selling space and time in bulk to selling single impressions by auction to whoever owns a logged-in audience.

• The evidence says much advertising does not pay at the margin; spending persists because measurement is costly and sellers measure.

• Privacy rules hit independent ad tech and the open web hardest, leaving platforms with logged-in first-party data stronger.

• Courts found Google a monopolist twice but chose conduct rules, so the companies' size and revenues remain untouched.

• AI makes ads cheaper to produce while moving ad space into answer boxes that a few companies control.

• The reader's test, like the advertiser's: what would have happened without the ad?

Module 18 Quiz

Check Your Understanding

1. How does the course reconcile record advertising spending with experiments showing weak returns?

A) The experiments were all run before digital advertising existed, so their findings no longer apply
B) Advertising spending is mostly required by law in regulated industries
C) Record spending reflects inflation alone, since real advertising spending has fallen for a decade
D) Average returns can be positive while the last dollar loses money, and sellers over-credit ads

2. Why did privacy rules end up strengthening the largest platforms, in the course's account?

A) Regulators exempted them from consent rules in exchange for paying fines
B) Their users log in, so they need no third-party tracking
C) They bought most of the independent ad-tech firms that the rules had weakened
D) The rules applied only in Europe, where the platforms have few users

3. What have the remedies in the American cases against Google had in common so far?

A) Each imposed rules on conduct rather than breaking the company up
B) Each ordered Google to sell a major business within two years
C) Each required Google to stop selling advertising in the United States
D) Each was overturned on appeal before taking effect

4. Which event set for March 29, 2027 does the watch list name?

A) Super Bowl LXI
B) The EU's decision on remedies in its adtech case against Google
C) The FTC's antitrust trial against Amazon
D) Australia's new limits on gambling ads

5. What is the course's last word to readers?

A) Ignore advertising, since none of it carries useful information
B) Ask what would have happened without the ad, as the experiments teach
C) Trust the platforms' measurements, since they have the most data
D) Prefer brands over store brands, since brands guarantee quality

Answer Sheet

Answers to every module's Check Your Understanding quiz.

Module	Answers
Module 1 Quiz	1. A 2. B 3. C 4. D 5. C
Module 2 Quiz	1. B 2. D 3. A 4. A 5. D
Module 3 Quiz	1. C 2. B 3. C 4. A 5. D
Module 4 Quiz	1. B 2. B 3. A 4. C 5. D
Module 5 Quiz	1. C 2. D 3. B 4. A 5. D
Module 6 Quiz	1. A 2. C 3. B 4. C 5. A
Module 7 Quiz	1. D 2. B 3. A 4. B 5. D
Module 8 Quiz	1. C 2. B 3. A 4. D 5. C
Module 9 Quiz	1. A 2. D 3. B 4. C 5. A
Module 10 Quiz	1. D 2. C 3. B 4. B 5. D
Module 11 Quiz	1. A 2. C 3. C 4. D 5. A
Module 12 Quiz	1. B 2. D 3. C 4. B 5. A
Module 13 Quiz	1. B 2. D 3. A 4. C 5. C
Module 14 Quiz	1. D 2. A 3. B 4. A 5. C
Module 15 Quiz	1. B 2. D 3. B 4. D 5. A
Module 16 Quiz	1. C 2. A 3. D 4. C 5. B
Module 17 Quiz	1. D 2. A 3. C 4. B 5. C
Module 18 Quiz	1. D 2. B 3. A 4. C 5. B