

"[AI] is a redefinition of the category of advertising. So if you think about what percent of GDP is advertising today, I would expect that that percent will grow." - Mark Zuckerberg, [May 2025](#)

What is Advertising? Why should I care?

Advertising is the business of turning **attention** into **intent**. It allows anyone with a **captive audience** to monetize surplus attention by facilitating **third-party economic activity**. The industry generates, in aggregate, >\$1T in annual revenues, representing a >1% and growing share of global GDP.

The first recorded ads were [posters etched into papyrus scrolls](#) in Ancient Egypt. Since then, **every technology platform throughout history has enabled new, better forms of advertising**. The printing press enabled newspaper ads in the 1800s, federal postage networks enabled direct mail in the 1900s, and radio and cable television networks enabled linear ads in the 2000s. The 21st century has seen the rise of digital channels like email, search, social, e-commerce, streaming, smart TV, influencers, gaming, AI/LLMs, and AR/VR that now collectively represent ~75% of global ad spend. Today, [two of the top seven](#) largest companies in the world - despite their founders' [initial disdain](#) for the underlying business model of advertising - generate the vast majority of their revenues from digital ads. And they're not alone: Netflix, Reddit, Instacart, Amazon, Uber, Disney and many others are on a similar path.




Figure 1: [Generative Value](#) (Eric Flaningam), [Visual Capitalist](#) (Therese Wood)

With Big Tech generating massive profits and crowding out the space, the **number of new adtech companies being started and funded over the past decade has fallen off a cliff**, down 80-90%. The space is ripe for disruption as recent unlocks in **generative AI** and **zero-knowledge cryptography** are collapsing the margins at some layers of the value chain and creating opportunities in others. For non-consensus investors and founders, there's never been a better time to build and invest in adtech.¹




Figure 2: [Tracxn](#) as 05/16/2025

¹ Highly recommend [this talk by Applovin founder Adam Foroughi](#) where he talks about failing to raise a \$1m pre-seed from Silicon Valley VCs in 2012. Applovin later raised \$4m from angels, became profitable in its first year, and is worth \$135B today.

What makes a good Advertising business?

A platform's **advertising value** is a product of four factors:

$$\text{Advertising Value} = \text{Audience} * \text{Engagement} * \text{Proximity} * \text{Attribution}$$

Audience measures a platform's reach in terms of engaged users. **Engagement** measures the amount and depth of engagement. **Proximity** measures how likely it is for users to engage in commerce on the platform. **Attribution** measures the precision with which sales can be attributed to specific ad campaigns.

New advertising channels become dominant when they excel at one or more of these dimensions. Take **engagement** for example: the average American spent 0.5 hours per day reading the daily newspaper in '80s, 3 hours per day listening to the radio in the '90s, 5 hours per day watching TV in the '00s, and today spends 5+ hours per day on their smartphone, half of which is scrolling social media. With every new platform both the **amount** and the **depth** of user engagement grows. This is most easily measured by information time-density, or information transmitted per user second: text is the lowest-dimensionality medium, followed by audio (+tone), images (+visual), videos (+time) and AR/XR (+depth). Technologies like LCDs, 4G/LTE, and now spatial computing raise the bar for what high user engagement looks like.

Proximity has three sub-factors. First, how much **data** does the platform have about its users? Platforms with richer information about users' past behavior can serve more highly-relevant ads with better conversion. Second, how much **commercial intent** do users bring? E-commerce platforms like Amazon deliver [superior advertising results](#) not only because they have data, but also because users go to Amazon with the explicit intent to purchase something; Google, where users go to find information and/or solve a problem, has a lower level of intent; Meta and Netflix, where users go to be entertained, is even lower. The final factor is, how frictionless is the **purchasing experience**? Platforms that can reduce the mental overhead needed to get users from discovery to purchase drive asymmetrically better results.

Attribution is the single most important lever of advertising spend. To see why, imagine you're the CEO of a widget company with a 15% cost of capital: if every \$1 of advertising spend drives >\$1.15 of annual recurring gross profits to your company, the EV-maximizing decision is to spend as much capital as your balance sheet (or ability to raise external financing) allows. In the real world, the attribution of gross profits to a specific campaign is lossy. Your CMO might tell you that - with a 90% confidence interval - a given campaign drove between \$1.14 and \$1.17 in first-year gross profits per dollar spent. Meanwhile, your CFO tells you that - with a 90% confidence interval - you can raise external financing at a cost of capital between 13% and 15%. How much of your balance sheet should you spend on marketing?

Advertising channels with strong attribution deliver results within a **tighter range of gross profit per dollar spent**, allowing CEOs to increase spend with the confidence that it will drive profitable growth with a reasonable margin of safety above their cost of capital. The counterweight to strong attribution is sacrifices in **user privacy**. The advertising industry is heavily regulated under regimes like [GDPR](#) in Europe, [PIPL](#) in China, [LGPD](#) in Brazil, [DPDA](#) in India, and [CCPA](#) in California. These regulations generally mandate user opt-in for data collection, providing users with some level of transparency and control as to how data is used, and making companies liable for storing data securely.

In short, the ideal advertising channel has a 1) large **audience**, with 2) strong **engagement**, and 3) user willingness or **intent to transact**, coupled with 4) the capability to **attribute** sales to specific campaigns without running afoul of consumer protection regulations.

Who are the players in the Advertising value chain?




Figure 3: the advertising supply chain.

Brands spend money on marketing to generate sales. So long as the return on incremental marketing spend is higher than a brand's cost of capital, it is always value-accretive for brands to increase their marketing spend. Brands generally spend up to 10-20% of top-line revenues on sales and marketing.

Agencies help brands determine what to say (creative) and how to say it (media). Agencies are people-intensive service businesses that run on thin 10-15% operating margins. They have largely consolidated into five publicly-traded conglomerates - Publicis, Omnicom, Interpublic, WPP, and Dentsu – with a combined 400k employees and \$70B+ annual revenues, trading at 5-8x EBITDA multiples.

Demand-side platforms (DSPs) help agencies manage campaigns by bidding on advertising inventory across exchanges and publishers. DSPs generate SaaS-like 80% gross and 20-30% EBITDA margins. The largest independent DSP, TheTradeDesk, generates \$2.5B ARR and trades at 40x EBITDA.²

Supply-side platforms (SSPs) help publishers monetize their advertising inventory. Think of them as yield optimizers that maximize revenues by participating in ad auctions across various exchanges. For example, Roku's SSP helps TV manufacturers monetize their inventory by serving programmatic ads.³ Since they're deeper in the value chain, pure-play SSPs Magnite and PubMatic generate \$1B ARR, but with lower margins (60% gross and 25% EBITDA) and lower multiples (15x EBITDA) than DSPs.

Publishers own the rails connecting creators and audiences, e.g. cable companies connect TV stations to their viewers and social media connects influencers to their fans. Publisher economics vary depending on the cost and capital intensity of maintaining the rails: cable and streaming businesses share revenue with TV stations, while social media platforms source UGC at ~zero cost thanks to their network effects.

Adtech infrastructure businesses come in four flavors: targeting, exchanges, servers and attribution.

Targeting businesses use data to improve ad efficiency by serving relevant ads to relevant audiences.

Exchanges host programmatic auctions that match orders between SSPs and DSPs. **Servers** select and show the winning ads to the audience. **Attribution** companies measure the effectiveness of campaigns.

In practice, the industry has seen a meaningful amount of consolidation and vertical-integration, in particular with SSPs, exchanges, servers and attribution commonly integrated into a single platform.

² Like many (non-network effect) software businesses, both DSPs and SSPs have high adjusted EBITDA margins (20-30%), but after stripping out stock-based compensation and intangibles amortization from M&A, operating margins are in the single digits.

³ Roku can be considered a **publisher-SSP hybrid** because they secure long-term contractual rights to the majority of their advertising inventory via OS licensing deals with TV manufacturers.

Where does Advertising value flow?

Let's say an agency charges brands a 15% fee on every dollar of advertising spend. The agency uses a DSP to allocate the remaining amount, after its own 15% fee, across various exchanges. The DSP then bids on various exchanges which, when bundled with ad server and attribution, charges another 10% fee (5% to buyers, 5% to sellers). The net proceeds from the auction are disbursed to the SSP, which charges another 15% fee, and finally remits the remaining amount to the publisher who owns the actual inventory. When it's all said and done, the **ads supply chain earns \$0.45** and leaves only \$0.55 to the publisher.

After operating costs the advertising industry generates \$0.10 in profits, valued by the public markets at a 14x multiple, with the majority of value (excluding publishers) going to the DSP and infrastructure layers. Publishers earn 75% of the profits and 66% of enterprise value created from every dollar of ad spend.

	Take-Rate	Revenue	Margin	Profits	Multiple	Value
Agencies	15%	\$0.15	15%	\$0.02	7x	\$0.16
DSPs	15%	\$0.13	20%	\$0.03	20x	\$0.51
Infra (exchange/server)	10%	\$0.07	40%	\$0.03	15x	\$0.43
SSPs	15%	\$0.10	20%	\$0.02	15x	\$0.29
Total excl. Publisher	--	\$0.45	22%	\$0.10	14x	\$1.39
Publisher	--	\$0.55	40%	\$0.22	10x	\$2.21
Total		\$1.00	--	\$0.32	--	\$3.60

Figure 4: illustrative advertising value chain economics.

In a world where each layer of the stack is independent, the best businesses to own are DSPs and adtech infra providers.⁴ However, reality is more complex as the biggest businesses vertically integrate other layers of the stack to expand their margins. Today the industry is segmented into “walled gardens” like Meta, Google, and Amazon that control the full advertising stack and earn monopoly-like profits within their platforms, and open ecosystems like email, SMS, or OOH (out-of-home) with fragmented providers. Publishers are increasingly becoming walled gardens, evidenced by recent ad exchange launches at [Netflix](#), [Paypal](#), [Spotify](#), [Uber](#), [Disney](#), [Fox](#), [NBC](#), [Warner Brothers](#) and countless other F500 companies.

Roku is an interesting case study. Roku launched its advertising business in 2012, ostensibly as an SSP helping smart TV manufacturers connect their ad inventory to various exchanges through Roku's TV OS. In 2017, the company launched its free streaming service, [The Roku Channel](#), becoming a publisher. In 2020, they launched [OneView](#) and became a DSP. In 2024, they expanded further up the value chain with [Roku Exchange](#). Roku's 90m active accounts streamed 34 billion hours of video in 2024, generating \$2B ARR for the company at a 14% EBITDA margin and supporting a market capitalization of \$10B.⁵




Figure 5: Roku's vertical integration over time: started in black (SSP) and expanded to gray (publisher, DSP, exchange) over time.

⁴ And publishers, of course, but we consider publishers to be outside the core advertising value chain.

⁵ Roku reports gross revenues (including publisher share) as its top-line. These figures strip out ~zero margin hardware revenues and publisher revenue-shares from the top-line in order to be comparable with Applovin and TheTradeDesk's reporting methodology.

Applovin is another interesting case study. Founded in 2012, the company started [ostensibly as a DSP](#) helping brands and agencies optimize ad spend across new mobile-first channels. Starting in 2014, the company doubled down on its highest-LTV customers - mobile game developers - and eventually became the dominant channel for mobile games to acquire new users. In 2018, after helping many leading titles grow to tens of millions of active users, Applovin launched an SSP to help those same developers monetize their games, and in 2022 it launched an exchange to internalize even more economics.⁶ Vertical integration has allowed Applovin to increase gross margins from 63% at IPO to 82% today, and they are now focused on expanding horizontally into ecommerce. Applovin's ad network reaches 1.6B users, generating \$6B ARR and \$2B FCF for the company and supporting a market capitalization of \$120B.




Figure 6: Applovin's vertical integration over time: started in black (DSP) and expanded to gray (SSP, infra) over time.

TheTradeDesk started as a pure-play DSP focused on streaming and connected TV in 2009, and operates with substantially the same strategy today. They cite their independence as a key differentiator for brands and agencies, as opposed to captive DSPs who are incentivized to prioritize their first-party channels even when they drive subpar results for brands. The company has grown revenues +12x since its IPO in 2016 and the stock is up +25x over the same period, valuing the company at over \$35B. Their platform currently integrates across 220 exchanges, publishers and SSPs, generating \$2.5B ARR and \$1B EBITDA. Unlike Applovin and Roku who benefited from vertical integration, TheTradeDesk's gross margins are relatively flat at ~80% since the IPO—though there are clear signs of [margin compression](#).⁷




Figure 7: Unlike Roku and Applovin, TheTradeDesk has not vertically integrated over time and operates as a pure-play DSP.

To summarize, the **best businesses in advertising are Google and Meta**: they source content from users for free, and within their walled gardens they control every layer of the stack from SSP through to DSP. Mark Zuckerberg even wants to [replace agencies](#) with AI-powered self-serve ad creation and campaign management tools. The next best business (Applovin) created its own walled gardens with a fully-integrated SSP/DSP/exchange purpose-built for the growing mobile gaming ecosystem. The next best business (TheTradeDesk) is an omnichannel DSP that facilitates massive ad budgets at the top of the funnel. The next best business (Roku) is an SSP that vertically-integrated as quickly as possible.

⁶ Applovin also operated as a publisher between 2018, when they began acquiring mobile game studios, and May 2025, when they [divested](#) their mobile games portfolio for \$800m.

⁷ TheTradeDesk announced [Ventura](#), a smart TV operating system (SSP) competing directly with Roku - one of its biggest SSP partners - in November 2024. Roku stock [fell 8%](#) on the news but later recovered. As of May 2025, Ventura has not yet launched and [looks to be shutting down](#) based on LinkedIn headcount and lack of OEM partnership disclosures in recent earnings calls.

Where are the Opportunities for startups?

The biggest companies will enable and aggregate **new forms of advertising inventory** that deliver superior results for brands and agencies. The playbook here is to identify **pockets of under-monetized user attention** and enable publishers to monetize their audience with ads without disrupting existing workflows (ideally, improving them) or cannibalizing existing engagement (ideally, growing it).

[Daisy](#) is enabling this future for influencers, unlocking a new form of advertising supply - **likes, reposts, and comments on social media** - and turning it into a full-scale performance channel. Brands and agencies come to Daisy to have their content “boosted” by a network of influencers, [driving CPMs lower than working with a single influencer](#). Influencers come to Daisy to monetize their following by engaging with sponsored posts from other influencers in the network, earning a steady source of income rather than lumpy brand deals—as one influencer put it, “*with Daisy I make \$50 every time I brush my teeth*”. Daisy reduces the months-long process of sourcing, negotiating, and executing an influencer marketing campaign down to a few clicks and enables influencers to get paid in minutes rather than months.

Companies with this strategy need to build critical tooling to turn informal impressions into a true performance channel. At a minimum, you need **routing** to decide which ads to surface on specific inventory, **serving** to physically display or execute the ad, **attribution** to track results and demonstrate value-add to brands and agencies, and **payments** to distribute proceeds to publishers. If you provide these tools, brands and agencies are more than willing to pay 20% fees to access high-quality ad inventory. Crypto rails - stablecoins and zkTLS - enable payments and attribution at dramatically lower costs, while generative AI enables the efficient routing and serving of ads across millions of users.

Where are the other pockets of under-monetized attention?

Rembrand recently [raised a \\$25m Series A](#) to build the ads network for **in-scene media**. The company’s platform helps creators embed **AI-generated “virtual product placements”** in videos in post-production, monetizing passively without disrupting the audience’s attention with traditional read-out or insert ads. Anzu [raised a \\$48m Series B](#) to build the ads network for **in-game media**, allowing game developers to monetize passively without disrupting user gameplay, e.g. by running ads on billboards in the background of a first-person shooter map. Nexad raised a [\\$6m Seed](#) to build an ads network for **in-conversation media**, allowing AI agents and chatbots to insert context-aware, sponsored ads seamlessly into their responses to user queries. Whether it’s social media influencers, game developers, or AI chatbots, the playbook is the same: make it easy for them to monetize their in-platform ad inventory, in a way that doesn’t disrupt the core user experience, and you have a winning formula to capture the category.

Crypto wallets and dApps are another source of rapidly-growing, under-monetized source of consumer attention with over 50m self-custody wallets active globally. These products have the benefit of high engagement (users come back multiple times per day or week) and high proximity (users go there to transact and spend money), making them an attractive marketing channel from a CPM perspective. Companies like [Bluwhale](#), [HypeLab](#), [Addressable](#) and [Spindl](#) collectively raised \$25m to build an ads network for the crypto industry—while the latter was [acquired by Coinbase](#) earlier this year, the race to build an independent web3-native ads network remains open with no clear winner. Such a network could help leading wallets and dApps add a new, stickier source of revenues, while also helping them grow their audience by opening up a new performance marketing channel.

While it's easy to look at the opportunity through rose-colored glasses - and indeed, we believe “dApplovin” will be a \$100B+ business eventually - entrepreneurs today face a real risk of being a couple years too early. There's two major challenges: 1) the leading wallets, which own distribution, make more money on swap fees than they could via ads, and 2) wallet-level onchain data is not very useful for ad targeting—most active users have dozens of wallets across dapps and chains, making it difficult to stitch together a single user profile, so crypto ad networks end up targeting a small subset of “whale wallets” which more often than not are institutional in nature. Two potential go-to-market strategies we see are: 1) build a new, ad-supported wallet that uses ad revenues to sponsor users' transaction fees on approved dapps, or 2) building an ads network that connects crypto news and analytics sites (which are inherently difficult to monetize) with a set of whitelisted dapps that advertisers are comfortable showing on their site.

Another interesting strategy is directly owning part of the **open source mobile stack** and integrating ads natively into the user experience. The distribution strategy here is to partner with OEMs and pay them up-front fees and/or revenue-share to **pre-install your app** on every phone they sell. For example, [Glance](#) (owned by InMobi) partners with Android smartphone manufacturers to place lock-screen ads on the phones of 230 million users in India and Indonesia, [Phoenix](#) reaches 400 million users in Africa with a similar distribution model for its mobile-browser-turned-superapp, and [ShareIt](#) reaches 2 billion users in Asia with its bluetooth-powered filesharing app that came pre-installed on Lenovo phones. You could build a critical utility app like a browser (a la [Wootzapp](#)) or file sharing app (a la [ShareIt](#)), integrate directly into the operating system (a la [Glance](#)), or potentially even build your own operating system (a la [KaiOS](#)).

Another source of underutilized attention on the internet is **e-commerce checkout pages**. Over 100 million people make an e-commerce purchase every day and - like crypto wallets - checkout screens have the benefit of high proximity because users are there specifically to transact. **Lifecycle marketing** startups like [Flywheel](#) help brands incentivize past customers to post brand-affiliated content on their social media at various points during the purchasing funnel. This effectively turns word-of-mouth into a performance marketing channel, where brands can measure the return on incentivizing customers to accelerate viral, word-of-mouth growth. While Flywheel focuses on first-party (lifecycle) marketing at checkout, another approach is to show relevant third-party ads at checkout: Tiger-backed [Rokt](#) was recently [valued at \\$3.5B](#) with this model, and Felicis-backed [Disco](#) [raised \\$20m](#) for the same.

While ecommerce is a large and growing segment, the **majority of commerce still happens in-person**: the average American spends ~4 hours per week shopping at retail outlets (e.g. grocery, convenience stores, malls), during which their attention can be directed towards driving incremental sales. One interesting approach, pioneered by [Qsic](#), is to partner with retailers to power in-store audio ads played over retail store speakers. The company then analyzes granular POS data to measure the uplift in sales for specific brands vs a control group, with the average retailer seeing a +14% uplift in sales attributable to audio ads, and recently raised a [\\$25m Series B](#) with plans to reach 12K stores by year-end 2025. Other adtech startups building retail media networks or [RMNs](#) include Firework, which [raised \\$150m](#) to power ads in [retailers' livestreams](#), Swiftly [raised \\$100m](#) to power ads on retailers' [apps and websites](#), and Kevel [raised \\$23m](#) for modular [omnichannel RMN infrastructure](#) across digital and physical channels.

Another massive and largely-untapped source of consumer attention is **live sports**. Of 260 million American adults, 160 million (~60%) regularly watch live sports and 37 million (~15%) bet on them. US sports TV and streaming rights hit \$30B for the first time in 2024. Because sports leagues lack capabilities to develop advertising products internally, they primarily rely on long-term licensing deals with the large walled-garden publishers, such as the NFL's deals with [Amazon](#) for Thursday Night Football (\$1B/yr), [Youtube](#) for Sunday Ticket (\$2B/yr), and [Netflix](#) for Christmas Day games (\$150m/yr). Publishers are willing to pay up for these deals knowing that live sports attract a massive, highly-engaged audience, with

a high-propensity to spend—the most valuable type of inventory. Sports leagues, teams, and stadiums are hungry to monetize novel ad formats (e.g., sponsored replays, virtual LEDs) rather than leave money on the table with licensing deals. Startups in this space like [WSC](#), [Supponor](#), [Kiswe](#), [FanAI](#), [LiveLike](#), [TAPPP](#), [Playback](#), [Streamlayer](#), and [Feature.io](#) have raised \$400m from venture capitalists.

Another strategy is aggregating **new types of screens**. Consumers are increasingly spending time staring at digital screens, e.g. while sitting in a self-driving car, wearing their AR goggles, working out on their treadmill, or sitting on a flight. Today, these channels are either very early in [experimenting](#) with ads or the platforms are owned by the publishers themselves (e.g., United Airlines' [Kinective Media](#)). There's obvious benefits to aggregating supply across multiple publishers and formats and building independent ads networks for the entire channel. This thinking has led to large VC investments in OOH (out-of-home) and CTV (connected TV) over the past few years. MNTN, a DSP for connected TV that's affiliated with Ryan Reynolds, [went public](#) at a \$1B valuation last month: the company generates >\$250m ARR, [>90% of which comes from small- and medium-sized businesses](#), effectively democratizing access to TV advertising to the long tail of businesses via its software tools.

The most interesting trend here is **AR/XR smart glasses**, which saw [≥1.5 million new device sales](#) in 2024 with some analysts projecting [30-50% CAGRs](#) through 2030 as incumbents like [Meta](#), [Apple](#), [Samsung](#), [Snapchat](#), [Sony](#), and [Huawei](#), as well as startups like [Xreal](#) and [Mentra](#), look to cement their spot in the next great computing paradigm: spatial computing (yes, I'm writing this from my Vision Pros). Smart glasses are a killer advertising platform for three reasons. First, they have an extreme amount of context-awareness because they “see” everything that you see; we can imagine an AR glasses app that tracks the stores you visit in real life and sells that data to brands paired with your hashed user credentials in order to re-target you on traditional social, mobile or CTV channels. Second, they have an extremely fine-tuned ability to measure users' attention (and therefore ad effectiveness) by tracking users' eye movements with infrared cameras; we can imagine an AR glasses app that measures attribution metrics for out-of-home advertising and uses its revenues to fund rewards for users. Lastly, they can serve extremely immersive ads by virtue of having a screen positioned mere inches away from your face; we can imagine an AR glasses app that injects ephemeral virtual ad overlays over billboards and empty wall spaces as users walk by them. Like the early days of mobile, the practical realities of advertising in AR will ultimately be determined by the limitations set by the leading operating systems.

While we're still in the early days of AR, our portfolio company [Over The Reality](#) (OVR) has built a full-stack AR/XR marketing platform that's already being used by brands like [Gucci](#) and [Balmain](#) today. Users can use OVR's [mobile app](#) to create 3D maps by scanning real-world locations with their smartphone camera. Brands and agencies can use Ovr's [web app](#) or [Unity SDK](#) to create AR ads, or buy 3D assets on the [marketplace](#). Since launching in 2021, Ovr's community has created [135K 3D maps](#) and [22K 3D assets](#), amassing the largest global 3D map database outside of [Niantic](#), the developer of PokemonGo. As the install base and developer tooling for AR glasses matures, the next great ads opportunity unlocks.

The biggest tectonic shift in advertising today is Google's [\\$250B+/yr search monopoly](#) being [unraveled](#) by the rise of [generative AI](#). In a world where **LLM-based chat interfaces** replace Google Search, a quarter of the global advertising market will be up for grabs. The winning LLM providers will need to buy or build the same pieces of infrastructure that Google runs internally: indexing, targeting, serving, and attribution. Zach Cohen and Seema Amble from a16z recently coined the term [Generative Engine Optimization](#) to describe the 25+ startups helping brands improve their visibility in LLM responses at the *indexing* layer.⁸ It's worth noting that the three pure-play SEO companies - Yext, Semrush and SimilarWeb - generate a

⁸ Today this usually means presenting content in structured schemas, adding citations, endnotes, and summaries, etc.

combined <\$1B/yr of revenues (vs Google's \$250B) and have all [destroyed value in the public markets](#). LLMs are even more powerful with deeper semantic understanding than Google's keyword search, and so it's likely that GEO - while [effective today](#) - will play an even smaller role than SEO does in search.

[Nexad](#) is positioning themselves as the “Applovin for LLMs” at the *targeting* and *serving* layers. LLM developers share user and query context with Nexad, and call Nexad's API which serves up sponsored recommendations or ads on-demand. The company claims to already reach [30m+ end-users](#) through its integration with six LLM apps.⁹ [Henry.Social](#) is attacking the *attribution* layer and looks more like an [affiliate network](#) rather than a traditional advertising network: brands specify the commission they're willing to pay for referred sales, and then LLMs can then index the catalogue of affiliate rates and use it to inform which products they recommend to users. Over time, LLM providers will adjust models for some chosen set of tradeoffs between optimizing for relevance to users (maximizing subscription revenues) vs optimizing for affiliate fees (maximizing advertising revenues). Per Ben Thompson's recent interview with Sam Altman, this seems to be [the direction OpenAI is leaning](#):

*“The kind of thing I'd be much more excited to try than traditional ads is people use Deep Research for e-commerce, for example, and... we could come up with **some sort of new model, which is we're never going to take money to change placement or whatever, but if you buy something through Deep Research that you found, we're going to charge like a 2% affiliate fee...** That would be cool, I'd have no problem with that. And maybe there's a tasteful way we can do ads, but I don't know. I kind of just don't like ads that much... I'm not going to say what we will and will never do because I don't know, but I think there's a lot of interesting ways that are higher on our list of monetization strategies than ads right now.”*

To summarize, the first type of adtech business we're excited about **unlocks new, valuable ad inventory** in high-growth channels. From social media influencers and LLM chatbots, to self-driving car dashboards and AR smart glasses, wherever consumer attention flows there is often an opportunity to serve ads and measure their impact in a high-fidelity way, creating new ways for brands to reach engaged audiences.

Applovin and, in our view, Daisy, are **special cases** of businesses in this category with the potential to grow beyond \$1-10B to generate truly outsized \$100B+ outcomes. The two companies have one very important trait in common: they create a **closed-loop ecosystem** where every new dollar introduced into the system recirculates around the network multiple times, paying a platform fee each step of the way. This is ultimately driven by user and developer behavior: Applovin started out as an DSP, helping game developers acquire users for new games and then expanded to become an SSP to help monetize games; as game developers made money through Applovin on their mature games, they reinvested the proceeds back into Applovin to acquire more users for their newer titles by advertising on other successful mobile games. Daisy has the inverse approach, whereby influencers come to Daisy to monetize their following, and subsequently reinvest their earnings back into Daisy to grow their audience even further. The other use cases don't share this quality – you probably wouldn't promote a FPS within another FPS, or an LLM within another LLM – and therefore have weaker network effects and a lower ceiling for growth.

The second category of adtech business we're excited about uses **proprietary data** to power **10x better attribution** for an **existing channel**. From billboards, to television, to mobile phones, even mature channels can see significant growth when better and more-precise attribution technologies arise. With advancements in AI enabling **machines to understand what users are doing** better than ever before, alongside advancements in cryptography enabling **users to prove what they're doing** easier than ever before, the way we measure and attribute ad effectiveness is set to fundamentally change.

But first, we need to understand how attribution works today.

⁹ One of the six is [Dippy.AI](#), owners of Bittensor subnets SN11 (Dippy) and SN58 (Dippy Speech).

How does digital Attribution work today?

[Cookies](#) are the primary method for ad attribution on the open web. Mechanically, they are small <4KB strings that are generated by a server, temporarily cached in users' browser memory, and included in any future HTTPS requests from the browser to that server.¹⁰ Typically, they contain a unique session identifier, a specified duration, a session counter, and an affiliate or campaign identifier in the case of marketing cookies. For example, a cookie representing the second visit (*sc=2*) by a specific browser (*mkt_uid=3f75e...*) who was referred by an affiliate (*cid=aff42*) within the past 30 days (*max-age=2592000 seconds*) would inject this text into HTTPS requests to the brand's website:

Set-Cookie: mkt_uid=3f75e3d2-1bd4-4c9a-b98c|sc=2|cid=aff42; Domain=.tracker.net; Max-Age=2592000; Path=/; SameSite=None; Secure

Globally, [more than 40%](#) of websites use cookies, with an average of [55-78 cookies](#) per webpage. The typical user caches hundreds or thousands of active cookies in their browser at any given time. **First-party cookies** live on a single website and often store information like a shopping cart snapshot (for ecommerce sites) or a JWT authorization token (to enable auto login for repeat visitors). Most cookies are **third-party cookies** or "pixels" that track users across many websites and often store information like an audience segment (e.g., "NY sports fan") or a unique third-party identifier (e.g., a hashed [RampID](#)). Google is by far the largest issuer of third-party cookies with a presence on >50% of all websites with cookies; a handful of other players have a presence on only 3-10% of websites with cookies.¹¹

The landscape faced by advertisers on the open web today is this: **two-thirds** of traffic runs on Chrome, where permissive rules let brands retarget users with active cookies for up to 400 days. The remaining **one-third** runs on Safari, Firefox, and Brave, which block all 3P cookies but allow for [LiveRamp-style alternative IDs](#) to track users who log across multiple sites with a common email address for 7-180 days.¹² These browsers have [native ads networks](#), but they remain relatively small. When neither cookies nor alternative IDs are available, advertisers must fall back on contextual or probabilistic signals like page keywords, device/IP fingerprints, or real-time intent with weaker targeting and lower conversion.¹³

Google has a complicated history with third-party cookies. It owns Chrome, the dominant browser which is [by far the least-restrictive browser](#) for cookie developers, and generates ~\$10B in annual revenues from third-party cookie-enabled ads.¹⁴ In 2020, amidst [pressure](#) from [regulators](#), Google announced an initiative to [make third-party cookies obsolete](#) "within two years". In Q1'24 the company ran an A/B test on 30 million users (1% of Chrome's user base) and found that [disabling third-party cookies reduced programmatic advertising revenue by 21-34%](#), which could be improved to an "only" 18-20% reduction by implementing Google's new [Privacy Sandbox APIs](#). As a back-of-the-envelope estimate, a 20% drop in third-party ads revenues would mean Google's shareholders losing \$2B/yr in revenues and ~\$11B in market cap¹⁵—roughly 0.5% of Google's current value. Instead, in July the WSJ [broke the news](#) that Google had internally [sunset its plans to phase out 3P cookies](#) and would continue with the status quo. It remains unclear how regulators, particularly in the UK and EU, will react to the move.

¹⁰ Technically cookies are not stored files; they are defined in the Set-Cookie HTTP header line and stored in a SQLite database.

¹¹ The Trade Desk, Amazon, Meta, Cloudflare, Quantcast (InMobi), Liveramp, Criteo, Oracle, Comscore. Note that a large presence does not imply a large business: Comscore trades at a \$22m market cap (<2x EV/EBITDA) and Quantcast was acquired for ~\$30m.

¹² Alternative identifier networks like LiveRamp are still valuable for cookie-enabled Chrome users because 1) unlike cookies, identifiers persist across devices and browser sessions, 2) deterministic matching enables finer-tuned targeting than, for example, targeting users based on audience segment, and 3) they enable retargeting of users on non-browser channels like connected TV.

¹³ For example, Safari/Firefox users who are not logged in, or the [~30% of EU users decline to be tracked](#) under [GDPR compliance](#).

¹⁴ Google reported \$31B in FY2024 revenues for Google Network. We estimate 1/3rd comes from web display ads that rely on cookies for targeting and attribution, with the remainder coming from AdSense for Search (search embedded on third-party websites) and AdMob (mobile apps) which relies on device-level identifiers rather than browser-based identifiers like cookies.

¹⁵ Assuming Google's current EV to trailing revenue multiple of ~5.5x.

With the uncertainty surrounding Google's approach to cookies over the past five years, many publishers have diversified by integrating with [third-party identifier networks](#).¹⁶ Websites can embed several (free) JS modules that ingest users' PII, such as an email or phone number, SHA-256 hashes the data locally in the browser, and sends it back to LiveRamp's servers, which stores hundreds of millions of user profiles in a pseudonymous graph powered by the hashed PII.¹⁷ If a user visits two websites with RampID model and uses the same login, LiveRamp matches the identical hashes and associates the user's activity across both sites. Because the match is deterministic and based on (pseudo-)unique PII, ads powered by solutions like RampID command a [significant CPM premium](#) over more lossy matching mechanisms like contextual or geographically targeted ads. For example, brands can take a list of user logins from their website and use a provider like LiveRamp to re-target those users across different websites and devices. This explains how, for example, you might subscribe to a fitness blog and in the subsequent week see an outsized number of gym ads on news sites (via cookies) or streaming on TV (via alternative identifiers).

However, you've also likely experienced ads unrelated to anything you've interacted with in the past, especially if you use a privacy-focused browser. There's the concept of **match rate** that measures how often publisher/SSPs are able to successfully match cookies with advertiser/DSPs to serve a relevant ad. Due to consumer privacy protections, every player in the ad supply chain - publishers, SSPs, exchanges, and DSPs - has to issue their own individual cookie to the browser, and sync their session identifiers in a pairwise process called a [cookie sync](#). Cookie syncs fail a quarter of the time, and since the typical ad auction requires multiple of them, the end-to-end cookie match rate is much lower at [31-45%](#) per Google. Alternative identifier networks can deterministically match user profiles on websites where users share the same login credentials or personal information, however in practice most page visits don't fit this criteria so the overall match rate seen by advertisers is [~28%](#), similar to cookies (though the company believes they can deliver 75%+ match rates through [direct integrations](#) with Google's and other DSPs). In the words of LiveRamp's Chief Connectivity & Ecosystem Officer Travis Clinger:

"The cookie may end up being like 3G. I think it existed a couple years longer than anyone used it as we all upgraded to 5G, and it was just an entirely better network. And that's kind of what I see [LiveRamp's] ATS as versus the cookie."

Mechanically, re-targeting via cookies in a programmatic ad auction works like this:

1. Your browser loads a webpage and identifies available ad slots. A small Javascript file called the header bidding wrapper (HBW) scans the page for ad slots. In some jurisdictions, users must consent to cookies before the HBW runs. 2. HBW requests bids from multiple SSPs for every available ad slot. The XHB payload often includes slot metadata (e.g., banner size, media types), privacy consents provided by the user, auction rules (e.g., bidding period, bid floors), and other publisher data (e.g., site keywords, contextual data) sent directly from the in-browser HBW to the SSP. If alternative identifiers like LiveRamp's RampID are available, they are included too. The browser auto-attaches each SSP's cookie if it already exists; if not, each SSP's response creates a new cookie (cookie_A) with its internal identifier for the session. 3. Cookie sync: each SSP returns a pixel pointing the browser to each DSP. Each SSP's response to the browser includes a pixel URL for each DSP. The browser loads these URLs, enabling DSPs to set their own cookies (cookie_B) and make a pairwise association with the SSP's identifier (cookie_A). If the sync is successful, both the SSP/exchange and the DSP agree that cookie_A and cookie_B represent the same browser session. 4. SSPs build an OpenRTB bid request file for each DSP and send it via JSON. The bid request file includes cookie_B, each DSP's internal identifier for the browser session that was synced in the previous step, as well as all of the fields in the initial HBW message to the SSP (step 2), plus additional fields added by the SSP like detailed IP address and network metadata or alternative identifiers like RampID.	5. The DSP scores the impression and submits a bid to the SSP/exchange. The DSP pulls data about the user and session from their internal database, scores the impression with their own internal methodology, and replies to the exchange with their bid in CPM terms and the content to be displayed if they win. 6. Three distinct decision-layers resolve the open ad slot. First, each SSP runs a first-price auction among its DSPs and returns the highest CPM to the HBW. Then, the HBW selects the highest CPM from participating SSPs and forwards it to the ad server. Finally, the ad server compares winning-bid CPM with the publisher's direct or guaranteed deals. 7. The ad server returns the URL for the winning ad's content to the browser. The browser fetches the winning content from the ad server or a CDN and renders it for the user. Sometimes the ad server sends in-line HTML or JS content rather than a link to an external URL. 8. The winning DSP returns a pixel inside the ad reflecting attribution. The browser's request to the pixel URL includes ad and auction IDs plus cookie_B, allowing the DSP to log the visit for attribution, billing, and frequency- capping. If a user clicks on the ad, the DSP redirects the URL and logs the visit for attribution. 9. The SSP/exchange bills the DSP for the impression and/or conversion. Both sides exchange impression log files and resolve significant discrepancies (typically hourly or daily) and the SSP bills the DSP for the final amount, typically paid on net 30 or net 60 terms.
--	---

¹⁶ Alternative identifier networks include RampID by LiveRamp (\$2B market cap), UID2 by TheTradeDesk (\$35B market cap), CORE ID by Epsilon (acquired for \$4.4B), ID+ by Zeotap (\$130m raised), ID5 ID by ID5 (\$27m raised), and ConnectID by Yahoo.

¹⁷ Technically there is another layer of [salting and encryption](#) so that LiveRamp never stores the raw hashes produced by browsers.

Adspace auctions look like Blockspace auctions.

If you're in crypto, the programmatic advertising value chain probably looks familiar—it mirrors the block production process—though the reverse is probably more accurate. First, users take an action, i.e. sign a transaction or load a webpage. The **gossip layer** distributes information to a set of **bidders** (searchers or DSPs) that use proprietary data and methodologies to value the impression/transaction. The **auctioneers** (SSPs or block builders) select the winning bid and, after taking a fee, forwards it to the demand relay. The **demand relay** (HBWs or relays) forward the highest bid to the **delivery layer** (ad server or block proposer) who serves the ad or transaction to the user. At 250-500ms from initial page load to an ad being rendered in the browser, the process is as fast as the [fastest block production machines](#) today.¹⁸

	User Request	Gossip	Supply Relay	Bidder	Auctioneer	Demand Relay	Delivery	Settlement	Finality
Advertising	Browser	HBW	SSP	DSP	SSP	HBW	Ad Server	SSP/DSP	250-500ms*
Solana Gulf Stream	Wallet	RPC	Leaders*	Searcher	Leader	--	Leader	--	400ms
Ethereum PBS	Wallet	RPC	--	Searcher	Builder	Relay	Proposer	--	12,000ms

Figure 8: adspace vs blockspace auctions.

The key difference between the two stacks is, of course, the permissionless nature of public blockchains. In advertising, there is no equivalent to a public mempool where DSPs (searchers) can parse through the set of available ad inventory (transactions). Instead, DSPs' access to inventory is intermediated by SSPs with which they must integrate tech stacks with and form commercial agreements in order to participate.

High-throughput blockchains like Solana were designed [without a public mempool](#), and therefore mirror the programmatic advertising much closer than Ethereum's [PBS](#). In both programmatic ads and Solana, ad or transaction inventory in the gossip layer must be **relayed through the auctioneer** (i.e., SSPs or block leaders), rather than being accessed directly by searchers or DSPs from a public mempool. This dynamic gives the auctioneer undue power to manipulate the auction in a way that's un-auditable by auction participants, and therefore vulnerable to abuse. The best example is Google being [sued by the DOJ](#) for abusing their power as an integrated SSP + DSP + exchange under [Project Bermanke](#): Google charged advertisers the second-highest bid, paid publishers the third-highest bid, and pocketed the difference for an internal "slush fund" used to artificially inflate its bids from Google's own first-party DSP.

Unlike in programmatic ads where DSPs only see the inventory that SSPs allow them to see, every validator on Solana sees the full set of gossiped transactions *eventually*—but as the saying goes, time is money. Solana's analogy to Project Bermanke came to light in Q1'24, when Jito saw [community pushback](#) for forwarding transactions to Solana's public validator port on a 200ms delay relative to their own first-party block engine. In both cases, Google and Jito leveraged privileged positions in the stack to extract value. Ethereum, with proposer-builder separation, reduces this risk at the cost of much slower auctions.

As auctions in both adspace and blockspace run into physical latency limits, i.e. the speed of light through fiber, there's diminishing returns to decreasing auction times beyond the current 400-500ms. Rather than increasing speed, the next-generation of auctions focus on increasing **depth of inventory**. In ads, **omni-channel auctions** allow DSPs to bid on CPMs across many channels in parallel. In blockchains, [cross-chain MEV](#) is arising as searchers price bundles of transactions across multiple chains in parallel. These auctions are emerging and more complex than their monochannel predecessors, but will inevitably take market share as they enable more efficient near-real-time auctions of both adspace and blockspace.

¹⁸ Solana's recently-announced [Alpenglow upgrade](#) will push auction times down to 150ms, faster than real-time ads.

What happens to Attribution in the age of AI?

What happens to the advertising industry in a world where users no longer surf the web and instead most web traffic is driven by AI agents running in [headless browsers](#)?

The first-order implication is that **match rates plummet**. Agents spin up fresh browser instances for every request, so there is no persistent browser session for cookies to track. For alternative identifier networks, unless an agent has 1) access to users' logins or PII and 2) permission to share it with third-party sites, these networks will not be able to make a deterministic match to an existing user profile. Companies like [1Password](#) and [Dashlane](#) are trying to solve this problem with agent-friendly credential security platforms, and could become the next-generation instantiations of businesses like LiveRamp and Epsilon CoreID. If an agent is indeed able to securely access and share user credentials, then such an agent will be able to spin up arbitrarily many headless browser instances and inflate traditional attribution metrics like CPMs, CPAs and CPCs—effectively making them as useless a metric as 'active wallets' for blockchains.

The second-order implication is the prophesied [Death of the Open Web](#). As the percentage of web traffic attributable to bots and agents continues growing, the value of ad inventory on webpages that do not require a login trends towards zero; meanwhile, the value of ad inventory on walled garden platforms like Google, Facebook and Netflix becomes correspondingly more valuable. Eventually, every webpage will either be login-gated or overrun by bots. And guess who owns logins? That's right, Google and Facebook. While [75% of logins](#) still use traditional email and password, over [55% of 18-25 year olds](#) prefer to login with email or social accounts. One way or another, Google and Facebook are poised to keep winning.

The third-order implication is an increased focus on **verifiable human engagement**. Ad fraud is rampant because everyone in the value chain except for advertisers gets paid based on either a % of ad dollars spent or a fixed \$ amount per impression or click with no regard to efficiency. This drives a structural incentive towards reporting inflated metrics by including bot-driven activity. Companies like [CHEQ](#) and [Human Security](#) raised a combined \$200m to fight ad fraud by detecting and blocking bots, but this is a relatively small opportunity. The bigger idea for users to anonymously attest to their humanity using biometrics-enabled secure computing elements. [Persona](#) and [Worldcoin](#), each worth >\$2B, are pursuing this strategy using smartphones and [commercial orbs](#), respectively, to detect human vs bot users and enable developers (including but not limited to adtech) to build on top of a shared identity graph.

The final implication is a shift towards **lower-funnel attribution**, i.e. rewarding publishers for actual sales rather than top-of-funnel impressions or middle-of-funnel clicks. In the past this was done through affiliate marketing—a [relatively small](#) (\$10B/yr) and medium-growth (10% YoY) channel.¹⁹ But in the age of AI, affiliate links suffer the same fate as cookies and alternative identifiers: agents can trivially spin up new browser sessions and strip out affiliate links and/or insert their own affiliate link to farm commissions.²⁰ Honey quietly [did this for years](#) with a simple browser extension until the practice was recently [banned by Google](#). Agents running locally on users' smartphone and laptop operating systems are 100x+ more powerful than browser extensions and cannot be banned en masse by Google or other similar platforms. Agents will eventually decimate trust in the affiliate marketing channel on the open web, requiring a fundamentally new way to attribute digital impressions and measure advertising effectiveness. For over a year now, we've been [writing about](#) and [investing in](#) companies building and leveraging a new technology called **zkTLS** that we believe will sit at the core of digital attribution for the next decade.

¹⁹ In affiliate marketing, a brand assigns a unique URL for each publisher (influencer), and whenever a user clicks on the URL anywhere on the web the brand tracks the user through checkout and attributes any sales to the publisher.

²⁰ While Type I error is described here, the Type II error (a legitimate agent being flagged as fraudulent) is equally problematic.

zkTLS powers digital Attribution in the age of AI.

We discussed above how cookies are cached in browser memory and sent to a specific third-party server via HTTPS headers. HTTPS is secured by ephemeral symmetric-key encryption [via TLS](#) that effectively creates a “[private phone call](#)” between exactly two counterparties: both the browser and server can 1) authenticate who they’re communicating with on the other end of the line, and 2) ensure that no third parties can decrypt or “listen in” to the conversation. Today, adtech uses a process called [cookie syncing](#) to share data cached in a single browser session with several third-party vendors for identification. The process is unreliable, prone to privacy abuses, and reduces both the effectiveness and speed of auctions.

zkTLS, pioneered by our portfolio company [Opacity Network](#), allows adtech businesses to achieve the same result as cookie syncing in a way that’s more secure and privacy-preserving than cookie syncing. Instead of caching data in browser memory and relaying it in HTTPS request headers to multiple servers, zkTLS inserts a third-party notary node that can verify data between client and server and attest to it arbitrarily many times at ~zero marginal cost. The simplest analogy for zkTLS is “Plaid for Everything”: it enables users to share their account data with third-party developers in a safe and verifiable manner, not just for financial account data but for any login-gated information or action on the web or in mobile apps.

Within adtech, zkTLS enables better targeting (i.e., users can reveal data about themselves) and also better attribution (i.e., users can verify that they took certain lower-funnel actions on third-party platforms). Most of the companies mentioned above - including Daisy, T-Rex, Flywheel, Henry.Social - are using zkTLS for lower-funnel attribution within their domains (social, web, ecommerce, and LLMs, respectively). The best example of a project using zkTLS for both targeting and attribution is [EarnOS](#): users earn rewards for interacting with brands across all digital channels, e.g., social media, ecommerce, and mobile apps. Brands and agencies use EarnOS to allocate marketing spend in a way that targets the high-quality users and drives verifiable outcomes according to the campaign parameters they set. Since launching in Q1’2025, EarnOS’ [beta](#) has enabled **26 million verified outcomes across 2 million verified users** for [brands including](#) Uber, Adidas, Minecraft, XBOX, Marvel, North Face, BMW, L’Acoste and Apple Music. Every time a user engages with a new brand or campaign on the platform, connects a new account, or refers a new user, EarnOS deepens its user-data graph and is able to target users with more relevant campaigns offering even higher rewards in the future. The result is a flywheel that makes EarnOS’ value proposition stronger as the network gets bigger, eventually creating the [Internet’s Reward Program](#).




Figure 9: EarnOS [beta platform](#)

Brands and performance marketing agencies that aren’t leveraging zkTLS-powered platforms to drive verifiable engagement risk falling behind as AI agents render traditional attribution signals meaningless.

Attribution beyond the Browser: the Nielsen playbook.

Another venture-scale opportunity we see is leveraging novel sources of data to provide **10x better targeting and/or attribution** for existing channels. We call this the *Nielsen* playbook, after the company that got its start providing targeting and attribution data for the radio and TV industries nearly a century ago. Arthur Nielsen created a device called the *Audimeter* that recorded which radio stations listeners were tuned to and for how long they listened to each station. Nielsen collected listening data, combined it with audience-level demographic data from surveys, and became the leading data provider to both publishers like radio stations measuring their reach, and to brands measuring their return on ad spend. With granular, reliable data on the types of users and level of engagement brands were reaching with specific radio (and later TV) ads, brands were able to massively increase marketing budgets for these channels to a peak of >\$20B per year each in the US alone. Nielsen grew alongside the industry, IPO'ing at a \$9B market cap in 2011 (\$26/share) and was [taken private in 2022](#) at roughly the same price (\$28).

From the perspective of the ad industry - that is, ignoring user privacy constraints - better targeting and attribution is strictly positive-sum: it's better for brands to have more certainty about their growth drivers, better for DSPs who see a greater amount of volume, and better for SSPs and publishers to have more granular data with which to measure and optimize inventory. Because they add value to everyone in the ecosystem, Nielsen-like companies can grow quickly to **ingrain themselves as critical data providers** to both buyers and sellers of ads. The key question is 1) how do these companies source proprietary data, and 2) can they scale quickly enough to become the de facto data provider for a given channel?

Today, Nielsen's data is sourced from a network of [~50K households](#) who earn a few hundred dollars annually for contributing their data through physical in-home devices (descendents of the [Audimeter](#)) and completing demographic surveys. This suggests Nielsen pays ~\$10-20m per year to data contributors, and we assume the total data costs (including processing and storage) amount to ~\$50m annually. Nielsen sells this data at a **70x+ markup**, generating more than \$3.5B in annual revenues from its efforts.

Many startups have tried to disrupt Nielsen's dominance in radio and TV by sourcing data from cheaper sources and passing a portion of the cost savings back to customers. One path is to source data by striking enterprise deals with cable companies and television OEMs in order to support firmware-level integrations to pull data directly from the existing hardware in customers' homes, e.g. cable set top boxes and smart TVs. The two leading companies with this approach, [VideoAmp](#) and [iSpot.tv](#), pull data from a combined footprint of 100m+ devices and are operating with a war chest of \$1B of venture capital raised. Only recently have these companies started to become viable alternatives to Nielsen: in 2019, CBS tried playing hardball by letting their [nine-figure contract with Nielsen lapse](#), but relented only nine days later. Five years later, CBS's parent company Paramount Global [churned from Nielsen to VideoAmp](#), citing a quintupling in Nielsen's prices in certain parts of their business. Nielsen's monopoly is finally under threat.

Our portfolio company [Rumi Labs](#) is pioneering the next-gen approach to sourcing TV viewership data. Rather than sourcing data via deploying dedicated hardware (Nielsen) or via OEM firmware integrations into existing hardware (Videoamp et al), Rumi leverages **edge AI models running on smartphones** to detect what users are watching on TV in real-time. Rumi achieves this by incentivizing a network of contributors to run software that fingerprint audio/video content with a unique code, to which audio samples from users' smartphones can be quickly matched to—like **Shazam, for videos**. Because Rumi can identify what users are watching in real-time at the edge (vs legacy providers that send data back to their cloud and report aggregated results later), it enables **both post-ad attribution & pre-ad targeting**: mobile apps can surface ads that are contextualized for what users are watching in TV in real-time.

Why is this valuable? The average American watches TV and streaming for 3-4 hours per day, with roughly half that time spent simultaneously scrolling on a smartphone. Despite the overlap, the content on users' smartphones is **entirely disconnected from the content on their TVs**. Rumi bridges this gap, enabling mobile app developers to build ads and other experiences with real-time context-awareness of what users are watching on TV—for example, users scrolling a social feed can be shown ads to vacation at the place where the scene they're currently watching on TV was filmed, or ads showing AI-generated avatars of users wearing the same outfit as the main character in the show. Other startups like [Ryff](#) (for video) and [Sounder](#) (for audio) are also focused on real-time targeting and attribution technologies.

Where are other pockets of under-utilized, high-value data?

- [375ai](#) uses a network of **physical AI edge nodes** to capture granular data for the out-of-home (OOH) advertising sector. Through its partnership with [Outfront](#), one of the largest owners of OOH inventory in the US (e.g., billboards, banners, transit ads), 375ai's network of edge nodes already covers >70% of Americans, surfacing insights like which specific [vehicle types](#) a certain billboard reaches in a privacy-preserving manner. In 2024, the company [raised \\$5m](#) and launched its [incentivized testnet](#) which currently has 180k active nodes across 100 countries.
- [XPLN](#) recruits users to participate in trials where they surf the web or smartphone while wearing **AR glasses with eye-tracking technology**. They use this data to score ads "in the wild" based on its page placement, surface area, semantic analysis of adjacent content, etc but without tracking user profiles via cookies or similar mechanisms. The company [raised \\$8m](#) in 2024 and announced it had signed on 300 brand customers. As the growth of the install base of AR glasses inflects upwards with the launch of new headsets by [Google](#), [Apple](#), and [Meta](#) in 2026-2027, we expect to see DePINs where users contribute eye-tracking data for token rewards.
- [TREX](#) users run a browser extension that rewards users for minting a **verifiable proof-of-engagement** whenever they complete an attributable action online (e.g., watch a video). As their install base grows, the sample size of TREX's dataset increases and it becomes representative of the broader population, and therefore valuable to both the supply-side (publishers) and demand-side (brands) of advertising. The challenge on today's web, however, is not so much about verifying *engagement* as it is proving **human engagement**. The largest incentivized browser extensions, Grass and Dawn, estimate that only [2-3% of their install base](#) are humans—the vast majority of nodes are run by bot farms. TREX can mitigate this risk by partnering with zkTLS-provider Opacity to require users to verify their identity on third-party platforms or with worldID-provider World for users to verify their unique identity by scanning their irises. The company is developed by HK-based venture studio [EVG](#) and raised \$17m in pre-seed financing.
- [Placer.AI](#), [Safegraph](#) and [CubelQ](#) source **pseudo-anonymized location data** by paying mobile developers to embed an SDK within their mobile apps. Because people-to-smartphones is a roughly 1:1 mapping, smartphone geolocation data can be used to infer human traffic at a specific location and therefore attribute a certain number of real-world impressions for out-of-home ads. These companies pay developers [1-3 cents per user per month](#) to contribute their users' location data, and claim to monetize the aggregated dataset at a rate of up to 8-33 cents per user-month. Mobile-app based DePINs like [WiFi Map](#) and [Silencio](#) are using these SDKs to fund token burns.

Attribution fuels the entire advertising engine. Startups that source proprietary, real-time data on consumer attention - and do so at sufficient scale to be representative to large brands and publishers - grow the pie for the entire industry and can become massive businesses in the process.

Should we own the Factories, or the Warehouses?

We believe it's much **more valuable to own the means of distribution** (i.e. the content “warehouse”) than the means of production (i.e. the content “factory”) in advertising. Factories, with increasing returns to scale on a fixed cost base, seem like good investments; but as new technologies arise, they run the risk of becoming uncompetitive and obsolete without significant new capital expenditures. Warren Buffett famously [learned this lesson](#) 50 years ago with his investment in Berkshire Hathaway's eponymous textile business and investors have been re-learning it ever since. In marketing, the old means of production - hiring traditional media agencies, contracting with social media influencers, and incentivizing users to create UGC - are all at risk of being [displaced by generative AI](#) that renders all prior investments in the factory near worthless. Warehouses are a different type of business than factories: they don't offer the same returns to scale, but they can 1) manage their gross margins in a downturn, 2) carry a far lower risk of technological disruption, and, most importantly, 3) can become a distribution **network** with density.

We believe content factories - startups that **enable brands to create new, more-engaging forms of content** - are not a venture-backable business. Content factories can be powered by AI - for example [Flam](#) for mixed reality, [Creatopy](#) for banners, [Creatify](#) for videos, and [Sploots](#) for clips - or they can be powered by using crowdsourced participation to power a “mechanical turk” content production machine. Either way, the risk is that incumbent platforms like Facebook and Netflix succeed on their [public plans](#) to internalize ad creation, or that LLM providers like OpenAI or Grok provide a substitutable alternative.

As we said before, we believe the value is in the distribution. A class of startups we call *Guerilla Marketing Platforms* **coordinate large groups of individuals** to create **novel types of ad inventory**. We've already discussed [Daisy](#), which coordinates influencers “boosting” of each other's sponsored content, effectively turning likes, comments and reposts into a new performance marketing channel. Whop's [clipping marketplace](#), where thousands of contributors watch streams, edit viral clips, and post them on social media, sees [>100 million views per day](#) suggesting they pay out >\$30m/yr to clippers. These companies scale extremely quickly by starting out with experimental budgets from web2 brands and performance marketing agencies, and rapidly taking wallet share as they see superior results.

The biggest risk to guerilla marketing platforms is **Goodhart's Law**: when a measure becomes a target, it ceases to be a good measure. Crypto readers will undoubtedly be familiar with platforms like [Kaito](#), whose [mindshare](#) leaderboard [devolved](#) from dominating CT and driving \$100M+ market cap moves in Q1'25 to the so-called “mindshare leader” trading at a [<\\$20M market cap](#) in Q2'25. In 2024, quest platforms like [Layer3](#) and [Galxe](#) that reward users for engaging with projects' social accounts followed a similar trajectory. While these platforms initially drove organic user growth for projects, as they scaled their user base they quickly devolved into effectively bot farm aggregators. Most projects will still use platforms like Layer3 or Galxe for application testing, i.e. to test their product with “live” traffic and iron out bugs, but the amount they're willing to pay for this is a small fraction of their marketing budget. The only way to get around this is to continuously invest in **lower-funnel attribution capabilities** that tie marketing spend to actual sales or gross profits, rather than top-of-funnel impressions or clicks. For example, if Kaito could track onchain purchases of a certain token and attribute it to the specific influencer's tweets who mentioned the ticker in the buyer's X feed (i.e., without the clunkiness of a ref link). It's not clear how this would be possible, but if it is, Kaito could reward only the yappers that drive token buy pressure.

The second biggest challenge for guerilla marketing platforms is **speed of coordination**. On the demand side, advertisers are accustomed to <0.5s auctions for digital ad space. On the supply side, coordinating hundreds or thousands of people in seconds or minutes, even with smartphones, is extremely difficult.

Opportunities in Adtech x Fintech.

With >\$1T flowing through the advertising industry each year, there must be opportunities to build companies that enable better, faster, smarter flows of money throughout the industry.

One interesting approach is [Attain](#) (fka [Klover](#)), which offers a <\$200 interest-free cash advances to users who agree to share their purchase data through Plaid. The company aggregates and anonymizes the data to power a SaaS platform for brands and retailers to analyze consumer purchasing behavior. Over 10 million users opt-in to provide merchant-level data to Klover, with 50k of those users providing more granular receipt-level data. A portion of the revenues from its SaaS business is used to fund a points-style [rewards program](#). Attain raised a [\\$25m Series B](#) in 2022 and [grew revenues 150% in 2024](#).

Attain is an interesting experiment in unit economics. Payroll loans are typically outstanding for only ~2 weeks at a time, which means some users may take out only one advance per year while others borrow money every month. Attain monetizes users' data in perpetuity after the initial advance, so the former - one-time borrowers - actually have a much higher LTV than repeat borrowers.²¹ The net result is that Attain is incentivized to acquire users once but not lend to them repeatedly (some might argue this drives customer-alignment, since payroll loans should be a temporary measure for borrowers to improve their cash flow and credit profile and graduate to cheaper, revolving loans over time). Publicly-traded payroll lenders like [Dave](#) and [MoneyLion](#) experience annualized credit losses in the 1-3% range. Assuming the average user borrows \$4k per year (20 advances), Attain would need to monetize its data at a rate of ~\$80 per user per year to breakeven on a gross margin basis—a high hurdle for a low-income user base.

To mitigate this, Attain offers “fast withdrawals” where users pay a fee between [\\$1.5 and \\$20](#), i.e. up to 10% of the advance, to receive funds immediately vs wait 2-3 days for ACH settlement. The closest public comp, Dave, reported that [99% of borrowers](#) used fast withdrawals at least once, but with no indication of frequency. Assuming two-thirds of loans are fast withdrawals with an 8% fee, the average user pays \$160 in early withdrawal fees (vs \$80 of credit losses) and Attain's fintech business runs at 50% gross margins. By monetizing users' data, Attain hopes to push its gross margins to 70-80% like a software business.

But is this reasonable? Public comp [Cardlytics](#) is much larger than Attain with 165+ million MAUs and monetizes purchase data at a measly rate of ~\$1.7 per user-year. Cardlytics sources data from wholesale partnerships with large banks: it pays [~45% of its net revenues](#) for banks to provide merchant-level data, on a 24-hour lag, using hashed card numbers as identifiers. Attain, via Plaid, can provide data in near real-time (<1hr) and uses stronger identifiers, i.e. hashed email addresses (1:1 or few) vs cards (1:many), but it runs into the same fundamental limitation around sourcing merchant-level - which is primarily useful for retailers - rather than SKU-level receipt data, which is the holy grail for brands and advertisers.²² Cardlytics paid [half a billion dollars](#) in 2021 (5x its current market cap...) to acquire [Bridg](#), a SaaS platform that sources SKU-level data from merchants at point-of-sale. For its part, Attain launched the [Merryfield](#) rewards app through which [50K users](#) - less than 1% of the user base - contributes receipt-level data.

Even if these companies manage to acquire SKU-level data across their entire use base, the hurdle for making a dent in their unit economics is high. To reach a 75% gross margin using our assumptions above, Attain would need to monetize users' purchasing data at a rate 20-25x higher than Cardlytics does today. The only viable path is powering real-time targeting and attribution for the adtech industry.

²¹ Technically Plaid requires a re-authorization every 12 months, but this does not require another loan issuance.

²² Merchant-level data only shows users spent \$[x amount at [y] merchant on [z] date. SKU-level data itemizes individual purchases.

[Revving](#) is an invoice factoring platform that recently raised a [\\$7m seed round](#) alongside a \$100m debt facility to help adtech platforms shorten cash conversion cycles by selling their receivables at a discount. The company integrates directly with ad exchanges and affiliate marketplaces, ingests real-time sales-level data, uses its risk engine to price the outstanding receivables, purchases them at a discount, and collects future payments directly from upstream partners (e.g., DSPs). This allows adtech platforms to get cash up-front rather than the standard 30-120 day payment terms. Other adtech-focused lenders like [FastPay](#), [OAREX](#), and [Harper](#) advance funds only after an invoice has been created (up to 30 days post-ad), whereas Revving's uniquely advances funds in near real-time after an ad has been served.

Embedded lenders like Revving, [Parafin](#), and [Wisetack](#) benefit from a similar flywheel to attribution: they benefit everyone in the ecosystem, and can therefore scale extremely quickly. For the software platform that owns the customer relationship - i.e., DSPs and ad exchanges for Revving, gig economy platforms and ecommerce marketplaces for Parafin, and SaaS platforms for home services business for Wisetack - embedding a third-party financing option helps you provide a better experience to your suppliers, who are now better-capitalized and in a position to grow their business with you. For borrowers, having a financing option embedded into the existing software and payments stack means easier, faster, and cheaper loans. Parafin launched in 2021 and is now [originating \\$1B per year](#) as a 4-year old company. Revving launched at a similar time and is aiming to [loan £1.8B](#) to UK-based adtech businesses over the next three years. The largest first-party embedded lenders live inside payments processing businesses and show the growth potential of the business model: [Square Capital](#) and [Stripe Capital](#) originate billions every quarter.

That said, these embedded lending businesses are structured - and should be valued - differently:

		Revving	Parafin	Wisetack	Stripe & Square
Platform	Platform Type	Adtech DSPs & Exchanges	Gig Economy & Ecommerce	Vertical SaaS for SMBs	Payments for SMBs
	Platforms				
Lending	Lending Product	Invoice Factoring	Revenue-Based Advance	BNPL Installment Loan	Merchant Cash Advance
	Collections	At invoice settlement: 30-120 days	Starts at platform payout (15-30 days); full repayment in 6-9 months	At installment expiry: 6, 12 or 24 months	Starts at payment settlement: (day 1); full repayment in 12-18 months
	Loss scenario	Advertiser or DSP refuses to pay or shuts down	SMB churns from platform or shuts down	SMB churns from SaaS platform or shuts down	SMB switches payment providers or shuts down
Unit Economics	Fees	1-5%	8-12%	5-10%	9-15%
	Gross Yield	12-18%	15-25%	10-20%	8-12%
	Credit Losses	<1%	2-3%	3-5%	2-3%

Figure 10: B2B embedded lending platform economics.

Invoice [factoring](#) businesses like Revving have the lowest credit losses because they **do not take credit risk against the borrower**. Unlike the other businesses, Revving gets paid whether or not the business it buys the receivable from continues as a going concern—so long as its upstream partner settles its [AP](#). That said, factoring businesses earn the lowest fees - typically buying invoices at 1-5% discounts - which means they operate at similar or lower gross margins than traditional (unsecured) lending products.

Lenders that take credit risk against borrowers rely on the **strength of the connection between borrowers and the embedded software platform**. There are two ways these lenders lose money: borrowers either [go bankrupt](#) (involuntary churn), or they churn from the platform (voluntary churn). There's a saying in boxing: it's not the hardest punch that knocks you out, but the one you don't see coming. The same is true for lenders: involuntary churn can be underwritten by analyzing operating history, but voluntary churn is difficult to predict. For example, it's harder for a food truck to churn from Doordash and find alternative order management and logistics solutions, than it is for a home services business to churn from ServiceTitan and find alternative CRM and invoicing solutions; as a result, Parafin experiences lower credit losses and is overall a better lending business than Wisetack.

What Adtech ideas are not likely to become Big Businesses?

After 20 pages, here's a few adtech ideas that we **don't** think can become VC-scale businesses:

- **Content factories**, as opposed to distribution networks. Discussed at length above.
- **SSPs** optimize yield for publishers, but have no say as to how advertising dollars are spent by the demand-side, and therefore little market power. Per our analysis above, SSPs take home only 3 cents in profits for every dollar of ad spend, albeit with low capex requirements. From 2004-2014, roughly a dozen venture-backed SSPs were founded and raised a collective \$1B in venture capital.²³ Only three of them - [Magnite](#), [PubMatic](#) and [Criteo](#) - are publicly traded today, and their combined \$4B market cap implies investments in SSPs as a whole have under-performed the S&P 500 for the past two decades.²⁴ The SSPs that manage to create venture-scale businesses will be those that follow the *Roku playbook*, i.e. have a first mover advantage in an emerging channel like connected TV, and expand as quickly as possible to become a full-stack publisher, DSP, and exchange. **We're not convinced there are "more Rokus" out there to build.**
- **Ad servers** serve an important but not super valuable role, and typically earn less than 1 cent in revenues for every dollar of ad spend. The most interesting company in the space is [Kevel](#), which raised a [\\$23m Series C](#) to enable retailers to build ad products that leverage their first-party data. While the ad server product might support \$100m ARR in the best case, growing beyond that will require expanding up the value chain to directly managing inventory (SSP) or spend (DSP).
- **Customer Data Platforms (CDPs)** are the infrastructure that enables effective and compliant targeting of ads. A CDP is effectively a data lake that ingests data from various sources, creates a unified list of customer profiles and attributes, encrypts any PII data, and converts the list to a format that can be easily ingested by a DSP where a brand or agency manages marketing spend. Hightouch [raised an \\$80m Series C](#) to build a "composable" CDP—in the words of their lead investor, *"they realized that cloud data warehouses are the new customer data platforms."* While we see an opportunity for next-gen CDPs to be built on [composable warehouses like Chakra](#), the "app layer" of CDPs is inherently less valuable in a world where user data is natively composable.
- **QR-code based attribution** requires users to scan a QR code with their smartphone camera. This is popular in markets like China, where users use an app-specific camera to interact with a closed commerce ecosystem like WeChat. In the US, users with native camera apps scan QR codes that connect to the open web by injecting a URL into users' native smartphone browser. Studies have found that [~2% of QR code scans are to malicious URLs](#) and over a quarter of all phishing attempts leverage QR codes, limiting their potential for mainstream consumer adoption. QR codes are simply not safe unless embedded in a closed ecosystem with existing distribution.
- **Lifecycle marketing** companies like [Braze](#) and [Flywheel](#) help brands engage and re-target existing customers. While these channels often deliver superior CACs and attribution capabilities, they're fundamentally limited in scope to existing users vs a much bigger pool of potential users. It's not clear to us how a platform can someday generate \$1B+ ARR re-targeting existing users, though we could be wrong especially for businesses with strong referral flywheels.

²³ Includes [GumGum](#), [Seedtag](#), [Sovrn](#), [Sharethrough](#), [Yieldmo](#), [OpenX](#), [Criteo](#), [PubMatic](#) and [Magnite](#)'s predecessors.

²⁴ All three of the SSP stocks trade below their IPO price.

About EV3

Escape Velocity (EV3) invests \$500k-\$2m in founders exploring novel crypto use cases, no matter how early. We're contrarians at heart and work hard to understand the companies and industries we invest in from first principles.

In the past twelve months we've seeded crypto-enabled adtech businesses including:

- [Daisy](#), powering influencer marketing and social media ads
- [EarnOS](#), powering verifiable campaigns on the open web
- [Over The Reality](#), powering ads in augmented reality
- [Rumi Labs](#), powering real-time attribution for TV and streaming
- [375ai](#), powering real-time attribution for out-of-home
- [Opacity](#), a general-purpose zkTLS infrastructure provider

If you're a founder, investor or researcher interested in the future of adtech, especially in relation to crypto, we'd love to chat: reach out to us at founders at ev3 dot xyz.

Thank you to Ian Ettinger (Daisy Pay), Tom Riordan (EarnOS), Dawufi Finkelstein (Opacity Labs), Anirudh Pai, Philip Qiu (Henry Social), and Harry Dewhirst (375ai) for their feedback on this essay.