The Arbitrage Window

AI, sport, and the market for auditable judgment

Joel A. Adejola · April 2026

joel@statswing.com

CAPITAL SPORT AI AUDITABLE JUDGMENT

For decades, serious analytical talent and institutional capital treated sport as adjacent to the real work: culturally enormous, commercially valuable, but not where the deepest decision institutions were built. That settlement is breaking.

Capital allocated to real estate, defense, finance, energy. Talent went into physics, medicine, law, public service. Sport was culture, entertainment, identity — important to billions, but not where the institutions that decided things were built.

Sport has become a convergence arena. Capital is allocating to it with the structural conviction it brought to real estate after the REIT cycle of the 1990s and to defense after 2001. Performance science is increasingly continuous with biotech, robotics, and human optimization: Honor-developed Lightning won the April 2026 Beijing robot half-marathon, with reporting crediting the smartphone maker's hardware and liquid-cooling work.[1] Media and public measurement have made sport the most legible decision system on earth: every choice a club makes is rated in real time by tens of millions of strangers.

The institutional question is no longer whether sport matters as a serious arena. It is which institutions will form at the intersections sport now creates, and how soon.

Performance science, biotech, robotics, media, and the public-measurement layer are all moving toward sport at once. The intersection that has moved fastest is capital. Capital is in. The decision infrastructure that capital relies on in other domains — the auditable reasoning chain between a number on a spreadsheet and a committee memo — has not arrived at the same speed.

That gap is the window.

01 / CAPITAL HAS OUTRUN ITS SOFTWARE

Sport is now an institutional asset class — not in the magazine-cover sense, but in the sense that the firms managing pension obligations and sovereign wealth allocate to it. Apollo Sports Capital completed its investment to become majority shareholder of Atletico Madrid in March 2026, with Apollo describing Atletico as the platform's flagship majority equity investment and committing up to EUR 100 million of additional strategic capital to support the club's long-term plans.[2]

This is not one firm's bet. Ares Management raised $3.7 billion of dedicated sports, media, and entertainment capital.[3] CVC Capital Partners established Global Sport Group, bringing together investments across properties including La Liga, Ligue 1, the WTA, Volleyball World, Six Nations, Premiership Rugby, and United Rugby Championship.[4] Arctos Partners has raised dedicated institutional capital for sports franchise stakes, and Sixth Street has become an active sports investor across franchises, clubs, and structured partnerships.[5] The pattern is systemic.

The numbers around football make the shift harder to dismiss. PitchBook reported in late 2025 that European sports private-equity activity had already exceeded EUR 10.6 billion for the year, roughly triple 2024's full-year level, and separate PitchBook analysis put private-equity, venture-capital, or private-debt participation at roughly 36.5% of clubs in Europe's Big Five leagues.[6] UEFA has tracked hundreds of clubs worldwide in multi-club investment structures.[7] UEFA's 70% squad-cost rule is in full effect, and the Premier League's squad-cost rule comes into effect from 2026/27.[8] Every transfer is becoming a constrained optimization problem under explicit financial rules.

The women's-sports signal confirms that institutionalization is not confined to men's football: Morgan Lewis has described new investor opportunities in the sector, and McKinsey projects the US women's sports market could reach at least $2.5 billion by 2030.[9]

The shape should be recognizable to anyone who has watched capital enter a domain before. When real estate institutionalized through REITs in the 1990s, the demand surface produced Yardi and CoStar. When defense procurement institutionalized after 2001, it produced Palantir and later Anduril. When legal procurement met current-generation AI, it produced Harvey, which announced an $11 billion

valuation in March 2026 after press reports put ARR at $190 million by the end of 2025 and later annualized revenue above $200 million.[10] The pattern holds: institutional capital arrives in a domain, demands the operational software it cannot operate without, and the companies that build that software become very valuable.

The pattern is not universal, and it is worth being precise about why. Maritime shipping, mining, and private credit each received institutional capital without producing a single consolidated operational-software layer in the same way. The pattern requires three conditions: genuinely absent operational software, regulatory complexity that rewards data layering, and durable contractual cash flows. Sport now has all three. The cash flows do not come from transfer fees, which are speculative and one-off; they come from the long-duration structure of media rights, league distributions, sponsorship, and matchday revenue — the revenue base that gives a club its institutional weight in the first place. The regulatory complexity is tightening, not loosening. And the operational software is, for the most consequential decisions, missing.

Apollo's own white paper, published in December 2025, calls sports a $2.5 trillion global ecosystem and emphasizes media-rights cash flows, under-levered franchises, early institutional development, and opportunity across the capital stack.[11] Apollo is not the protagonist. It is the vocabulary lesson. The relevant buyer does not talk primarily about dashboards, scouting vibes, or startup TAM. It talks about underwriting, origination, cost of capital, capital-stack flexibility, cash-flow durability, and platform capability.

The capital is arriving. The decision layer it requires — the layer that turns a number on a screen into a defensible recommendation — has not arrived with it.

---

02 / DATA ACCESS IS SOLVED; DECISION INTEGRITY IS NOT

Football's data layer is more sophisticated than it has ever been. StatsBomb covers clubs and federations with granular event data. SkillCorner builds physical tracking models from broadcast footage and has raised significant capital. Hudl/Wyscout sells scouting video and data at global scale across professional football.[12] Front offices have access to more structured football information than at any point in the sport's history.

The decision layer is emerging. Well-capitalized, credentialed teams have entered with serious acquisition strategies and significant analytical depth. The gap between what those entrants currently provide and what institutional capital actually requires remains large. Stating the gap precisely takes a distinction that is easy to write down and harder to internalize.

Current solutions provide decision *support* — better inputs for human judgment. A dashboard, data feed, scouting video layer, model score, and search tool can all improve the information available to a decision-maker. That is meaningful. It is not the same as decision infrastructure.

Decision support improves the inputs to judgment. Decision infrastructure makes the reasoning process itself repeatable, auditable, and defensible.

Consider the buyer. When a private equity firm acquires a football club for hundreds of millions of pounds and needs sporting due diligence, the analyst preparing the investment-committee memo does not need a better dashboard. They need a sporting-diligence memo that can be attached to an investment memo: visible methodology, comparable cases, stated uncertainty, and a recommendation the institution can defend. Decision support gives that analyst better inputs. Decision infrastructure gives the institution a defensible process. The first reduces individual error; the second creates institutional accountability. The price point, the buyer profile, the deliverable format, and the contractual relationship around each are different.

This is not an anti-consulting argument. Consulting can be the beachhead because trust often begins with bespoke work. The difference is what compounds afterward. A consulting engagement compounds into relationship, reputation, and case experience. Decision infrastructure compounds into method, audit trail, product surface, and institutional memory.

The reason this distinction matters in practice — not just in definition — is that the buyer's failure mode is different in each case. A scout using decision support and getting it wrong has a bad day. An investment committee approving a transfer based on decision infrastructure and getting it wrong has a paper trail. The paper trail is what allows the institution to learn, apportion responsibility, and defend its decisions to its own capital partners. Buyers who can absorb individual error differently from institutional error will pay differently for the two products.

Sport has a strange skin-in-the-game problem. The asset is a human performer inside a role, a tactical context, a development curve, and a market. The allocator cannot directly inspect the asset by becoming it. Even ex-player expertise does not automatically explain whether another player will translate across role, league, team model, coach, injury history, and price. That makes the observer's method unusually important. The buyer is not only buying a view on whether a player is good. The buyer is buying a reason to trust the chain that produced the view.

A football-native example anchors the structural point. Set-pieces are football's closest thing to a playbook: bounded, rehearsed, role-assigned moments inside a fluid game. Corners, throw-ins, long balls, second balls, height, and duel competence are not decorative details; they reveal a squad's hidden beliefs about who is trusted to head, mark, recover, strike, protect the counter, and execute under rehearsed constraint. As set-pieces and margins became more important to how teams create advantages, the measurement primitive underneath aerial evaluation became consequential.

The standard aerial-duel definition records a contest only when both players leave the ground. In STATSWING's single-match study, that definition excluded approximately 80% of contested aerial situations.[13] The metric does not measure what many practitioners think it measures. It misses exactly the wrestling, pinning, first-contact, and second-ball situations that set-piece and long-ball football depend on. Decision support built on the flawed definition inherits the flaw without surfacing it. Decision infrastructure that examines the definition — that asks whether the metric measures the football question — is a different product entirely.

The pattern is not limited to aerials. Aerials audit the event definition. Possession adjustment and its reanalysis audit the normalization. Mechanics audits the inference from action output to transferable skill. The work repeatedly asks what is collected, what is normalized, what is inferred, and whether the inference carries the warrant practitioners assign to it.

Notice what this implies for intelligence platforms assembled through acquisition. A platform that buys multiple analytics companies inherits the metric definitions of each acquisition. It assembles intelligence from parts that may not share foundational assumptions. Platforms assembled through acquisition can become powerful. They also inherit the assumptions of the products they acquire. If the foundational definitions are wrong, integration can scale the error faster than it corrects it.

Dr. Ian Graham, former Director of Research at Liverpool FC, observed at a StatsBomb conference that more than 50% of "big transfers" — those at GBP 10 million or above — fail by a minutes-played criterion.[14]

Football has solved data access. It has not solved decision integrity.

DECISION-INFRASTRUCTURE STACK *Where the buyer's burden changes from better inputs to accountable process.*

01 **Data access**
Feeds, video, event data, tracking, search

02 **Decision support**
Dashboards and model scores improve the inputs to judgment

03 **Decision infrastructure**
Method, uncertainty, comparable cases, and recommendation

04 **Institutional accountability**
Audit trail, decision memory, and defensible process

03 / WHY THE WINDOW IS OPEN

The interesting question is not whether AI changes things; everyone already believes that. The interesting question is what specifically changed, what gap opened, and why sport is the arena where the gap has become most visible.

Begin with the operator. AI agents, given a clear specification, are competent. They retrieve, structure, compare, analyze, and write. The remaining bottleneck — which used to be execution time and is now nothing of the kind — has migrated up the stack. It sits in the operator's *abstraction*: the willingness and capacity to hold complexity over multi-hour stretches, the focus required when the agent is offering tempting work that is not load-bearing, the patience for the dialectic that refines the spec. None of these are technical capacities. All of them are human.

Ethan Mollick, the Wharton professor whose research has become one of the more widely read sources on AI's workplace effects, quantified one version of this with BCG in 2023: GPT-4 users completed 12.2% more tasks, 25.1% faster, at 40% higher quality than colleagues without access — but the capability was uneven across tasks.[15] Mollick called it the *jagged frontier*. The operator's judgment about *where* to deploy the tool became the remaining scarce resource.

There is a diffusion lag underneath this. Anthropic's Economic Index, published in March 2026, shows that even the most exposed occupational category — Computer & Math — exhibits only about a third of its theoretical ceiling in observed usage on Anthropic's platform.[16] The gap between what AI can do and what professionals are doing with it is large. The March 2026 learning-curves update adds a complication: high-tenure users show persistent usage differences, and geographic convergence appears slower than earlier diffusion curves might have suggested.[17] The right-tail compounding interpretation is plausible, but it is an inference from one provider's platform data, not independent labor-market fact.

The deeper source of the arbitrage, however, is not technology. It is human inertia. The gap exists because most organizations — and most individuals inside them — resist the possibility that the entire frame within which they have been working may need to be rebuilt from the definitions up. Not tweaked. Not enhanced with a copilot. Rebuilt. AI is the mechanism that makes the right-tail operator's advantage possible. The source is the difficulty of accepting that one's prior model of how work should be done may no longer hold.

Jack Dorsey and Roelof Botha articulated a version of this at Block in March 2026. Their thesis: organizational hierarchy is an information-routing protocol built around human span-of-control limitations, and AI can replace the routing function itself.[18] In sport, the version is precise. Most analytics groups in football are using AI to enhance existing metric definitions and existing recruitment workflows. The arbitrage belongs to whoever is willing to rebuild the definitions.

The throughput of an institution is a function of how epistemically sound its primitives are. Institutions built on AI propagate epistemic unsoundness at speed. Rebuilding from sound foundations is not optional; it is the precondition for the intelligence layer to be trustworthy. The aerial example matters for this reason. Once the market starts caring about set-pieces, height, duels, and second balls, an inherited aerial-duel primitive stops being a harmless metric quirk. It becomes a capital-allocation risk.

Ajibola Lawal called one configuration of this earlier in 2026 the *Young King*: the operator who enters a vertical they were not formally credentialed for and competes because they have judgment, an LLM, and the willingness to be told they are wrong.[19] The window stays open as long as the median operator inside the vertical has not adopted the same posture. Once a local *Wise Old King* appears — the credentialed expert who has rebuilt as a centaur — the Young King's edge closes against that one.

Sport is where this gap is most visible because the asset transfer in sport happens at the decision point. Buying a player for GBP 50 million is making a capital allocation decision about the asset itself, in a way that PE portfolio acquisitions, where assets are abstracted behind financial instruments, are not. The proximity gives the decision layer direct, measurable impact on outcomes, and direct visibility into whether the methodology is sound. A natively built system with deep domain specificity can compete with a platform assembled through acquisition because the buyer can evaluate the methodology directly. There is no fund-of-funds opacity to hide in.

The arbitrage is not "we have AI and they don't." It is not "we work faster." It is the fit between domain judgment, production speed, primitive quality, and an institutionalizing market that has not rebuilt its decision layer. The window opens because the specificity of the problem — auditable decision infrastructure for institutional buyers of sporting assets — sits adjacent to the integration overhead of acquisition-assembled intelligence stacks, and the operating model fits the specificity unusually well. The resistance to rebuilding from foundations is the window. AI is the mechanism.

---

04 / STATSWING AS FIRST FIELD SITE

STATSWING is where this thesis has first been tested. It is a sports intelligence institution — research-first, product-led — built to serve the decision-infrastructure gap.

In early 2026, the first public version went live. The launch-state surface covered roughly 14,800 players across 21 leagues by our internal count, and the public product has since crossed 100,000 organic visits by STATSWING's public/operator-reported figures. The research published around the first surface covered contested aerial measurement, epistemic certainty in recruitment, and the assumptions buried inside possession-adjusted statistics. Underneath both, an internal research-orchestration substrate connects proprietary research, agent-assisted analysis, and the public surface.

I built the first version, and the institution remains accountable through me. What I brought to the agents was not technical sophistication. It was three years of football scouting work, three semesters of philosophy on epistemic justification, a year of investigative reporting at the International Centre for Investigative Reporting, time in student government, and roughly three hundred Writing Center sessions watching how people actually think through problems.[20] None of that material was invented by the agents. What the agents did — and what most organizations have not yet absorbed in their working practice — is collapse the time between having a clear sense of what something should be and the artifact existing in the world.

That collapse is the arbitrage. AI did not supply the judgment; it compressed the distance between judgment and an institution-shaped artifact. An operator with domain knowledge, training in evaluating the trustworthiness of claims, and meaningful proficiency with current-generation AI agents can now produce public, testable, institutional-grade work at a speed and scope the market has not yet seen from a single team, let alone a single operator.

The background matters because it produces a specific kind of judgment. Epistemic justification trains the instinct to ask when a metric is measuring something other than what practitioners think it measures, when a data definition obscures the mechanism it claims to capture, when a scout's report is reliable and when it is performing reliability. That instinct is what makes the decision-infrastructure framing possible from inside the work.

That is as far as the evidence currently goes. The stronger claim has to be earned through buyer-grade artifacts, inspected methodology, decision memory, and recurring institutional work. If those do not materialize, my claim should shrink: STATSWING may be an unusually serious football research project, not evidence of a transferable institutional pattern.

---

05 / WHAT WOULD PROVE THIS WRONG

The thesis can fail in four ways.

First, the sample may stay at one. STATSWING may show a public instance of primitive audit, AI-enabled production, and decision-infrastructure framing in football without proving that the model generalizes across buyers, sports, or verticals.

Second, the best-capitalized buyers may build the layer in-house. Apollo-like firms may treat specialized sports underwriting as a core platform capability, not a function to buy from outside specialists.[11] The external-vendor case then depends on speed, independence, domain specificity, published methodology, and auditability.

Third, the diffusion lag may close faster than expected. The March 2026 Anthropic update supports caution more than certainty: it shows tenure-linked differences and slower geographic convergence, but it does not prove that early adopters will keep pulling away across all knowledge work.[17]

Fourth, a well-capitalized rollup may already have assembled the missing layer. The relevant question is whether breadth has produced coherent method. If a rollup publishes compatible definitions, traceable recommendations, and uncertainty practices at scale, the native-coherence advantage narrows quickly.

None of those objections kills the thesis today. Each can kill it later. They define what to watch, what to build, and what would count as being wrong.

---

06 / WHAT TO WATCH NEXT

The first market test is purchasing behavior. By Q4 2027, at least one club, multi-club ownership group, PE-backed sports asset, or sports investor should hold a recurring external contract for structured, auditable sporting-decision diligence, distinct from data feeds, scouting video, dashboards, or ordinary advisory work. The thesis strengthens if that purchasing category becomes visible. It weakens or time-shifts if diligence remains bespoke, hidden, or bundled inside familiar consulting and data relationships.

The second test is whether sporting judgment attaches to capital process. By the end of 2028, at least one sports acquisition, governance process, or major capital allocation should publicly reference external sporting-decision diligence, methodology, or an independent intelligence provider. The constraint is visibility: this work may already happen privately around multi-club ownership and sports-capital

underwriting without becoming part of the public institutional apparatus.

The third test is provenance demand. By the end of 2028, at least one major sports analytics or decision-intelligence provider should publish a methodology audit trail: versioned definitions, traceable recommendation chains, uncertainty disclosure, or comparable governance artifacts. If audit trails become a buying criterion, the category is forming. If buyers reward outcomes without asking how the reasoning chain was produced, provenance may be philosophically right and commercially premature.

STATSWING has its own burden. It should produce a transfer or sporting-diligence artifact that an institutional reader would plausibly attach to an investment memo. Insight is not enough. The artifact has to carry methodology, comparable cases, uncertainty, and recommendation in a form legible to a sporting director, investment team, owner, or board.

It also has to become visibly inspectable. An institution arguing for auditability has to publish a methodology trail for its own research primitives before the market demands one by default: definitions, revisions, uncertainty, retrodiction tests, and known limitations. It should build decision logs so past judgments can be revisited and scored against the assumptions that produced them. Without memory, decision infrastructure becomes one-off persuasion.

The commercial test is whether the method survives revenue. Consulting, research, diligence, or bespoke work may be the beachhead, but the durable asset is method plus audit trail. The thesis strengthens if STATSWING becomes one of the first external decision-infrastructure vendors with a recurring institutional contract in sport. It weakens if commercial work requires abandoning the decision-infrastructure thesis and dissolving into generic advisory labor.

The broader test is whether the operating model travels. By April 2029, at least three operator-led institutions should exist in newly institutionalizing verticals — legal diligence, regulated procurement, public-sector audit, healthcare decision support, or similar categories — with paying institutional clients, published methodology, and recurring revenue. If this remains idiosyncratic to one operator and one football market, the claim narrows.

The watchlist is equally important. Apollo-like investors may build sports decision infrastructure internally, acquire it, partner with specialists, or outsource narrow diligence. Well-capitalized rollups may close the primitive-audit gap if they publish coherent definitions, traceable recommendations, and uncertainty practices at scale. AI diffusion may also close the window if usage rises symmetrically and median operators catch up quickly. Regulation may strengthen the thesis if squad-cost rules, multi-club governance, LP scrutiny, and acquisition diligence make sporting method explicitly inspectable; it weakens the thesis if buyers satisfy oversight through lawyers, bankers, and existing advisors without changing the decision layer.

---

07 / PROVENANCE AND JUDGMENT

Provenance is not only a question of where data came from. It is the chain between observation and judgment.

The aerial-duel metric that misses most contested aerial situations in a single-match study breaks at the event-definition layer. Possession adjustment can break at the normalization layer. Mechanics can break at the inference layer, when action counts are treated as evidence of transferable skill without enough attention to execution quality, role, and context. A transfer that fails because the recruitment team missed a known weakness visible in the data is not only a measurement failure. It is a failure in the chain between measurement and judgment.

In January 2026, a forensic audit of AI-era survey papers examined 5,514 citations across fifty papers published between September 2024 and January 2026. For nearly one in five citations, the digital chain of custody was severed; the referenced paper could not be verified through the citation chain. Ilter called this "The 17% Gap." It is a preprint, not settled literature, but the evidence for chain-of-report decay in the AI era is no longer merely theoretical.[21]

This matters beyond academia because it is the visible edge of a structural shift. Most of the institutions we built around the previous information economy — the journals, the diligence reports, the investigative units, the underwriters, the analytics teams inside clubs — were structures for managing the scarcity of information and the cost of producing it. The institutions of the next economy will manage something different: the trustworthiness of testimony. Producing words is now trivial. The question of who said what, on what evidence, in whose voice, has become the load-bearing problem.

C.A.J. Coady's *Testimony* helped establish testimony as a serious basic epistemic source: not reducible in any simple way to perception, memory, or inference.[22] The AI-era question is what happens when chains of mediated testimony lengthen, pass through systems whose reliability we cannot directly assess, and propagate at machine speed. Elizabeth Fricker has argued more radically that LLM outputs should not be classified as testimony at all but as instrument readings, because current AIs lack the mental states and communicative intentions that testimony requires.[23] If she is right, the provenance problem is more severe than the main argument here suggests.

Other domains are already responding with standards. The C2PA standard — the Coalition for Content Provenance and Authenticity — is building literal provenance infrastructure for media. Supported media and document files can carry verifiable signatures of origin and edit history.[24] Institutions are not only arguing about trust. They are building machinery for it.

A company that produces structured, auditable reasoning for a major player acquisition decision — one that shows its evidence, assumptions, metric definitions, comparable cases, uncertainty, and revisions — is a provenance institution. It does not sell data alone. It sells an audit trail for testimony about a player's value. The buyer is not paying for the recommendation alone. The buyer is paying for the ability to evaluate why the recommendation was made, by what method, on what evidence, with what acknowledged limitations.

The larger claim is that structures for managing trustworthiness in a world of trivial production will become one of the central institution-building problems of the next decade. Sport is one arena. Law is another. Finance, research, governance, and public administration will each need their own.

When production is commodified, what remains is judgment. The institutions that certify judgment — that show their work, version their definitions, and stand behind their recommendations — are the ones that will matter.

For correspondence: joel@statswing.com.

— Joel A. Adejola, Lawrence, Kansas, April 2026.

---

08 / NOTES

1 "Winner of Beijing robot half-marathon smashes human world record by 6 minutes," *South China Morning Post*, April 2026, https://www.scmp.com/tech/article/3350610/winner-beijing-robot-half-marathon-smashes-human-world-record-6-minutes; Victoria Song, "A humanoid robot finished a half marathon in 2 hours and 40 minutes," *The Verge*, April 2026, https://www.theverge.com/tech/914869/humanoid-robot-half-marathon-record. ↩

2 "Apollo Sports Capital Completes Transaction to Become Majority Shareholder of Atletico de Madrid," Apollo, March 12, 2026, https://www.apollo.com/institutional/insights-news/pressreleases/2026/03/apollo-sports-capital-completes-transaction-to-become-majority-shareholder-of-atl-tico-de-madrid-3254644. ↩

3 "Ares Management Raises $3.7 Billion of Sports, Media and Entertainment Capital," Ares Management, September 14, 2022, https://www.aresmgmt.com/news-views/ares-management-raises-37-billion-sports-media-and-entertainment-capital. ↩

4 "CVC Announces Global Sport Group's First New League Investment with the Acquisition of Equine Network," CVC Capital Partners, 2026, https://www.cvc.com/media/news/2026/cvc-announces-global-sport-group-s-first-new-league-investment-with-the-acquisition-of-equine-network/. ↩

5 "Arctos Sports Partners Announces Fund I Close With More Than $3.0 Billion," Business Wire, October 14, 2021, https://www.businesswire.com/news/home/20211014005255/en/Arctos-Sports-Partners-Announces-Fund-I-Close-With-More-Than-%243.0-Billion; "Sixth Street Closes Minority Investment in the New England Patriots," Sixth Street, https://sixthstreet.com/investment_announce/sixth-street-closes-minority-investment-in-the-new-england-patriots/. ↩

6 "How PE firms are hoovering up sports media rights in Europe," PitchBook, https://pitchbook.com/news/articles/how-pe-firms-are-hoovering-up-sports-media-rights-in-europe; "Private equity shakes up football tactics after recent troubles," *Financial News*, citing PitchBook analysis, https://www.fnlondon.com/articles/private-equity-shakes-up-football-tactics-after-recent-troubles-fcf9123d. ↩

7 UEFA, *The European Club Finance and Investment Landscape* microsite, 2026, https://ecfil.uefa.com/; UEFA, "New report highlights record revenues and increasing investment into European football," February 26, 2026, https://www.uefa.com/news-media/news/02a2-200452a66064-0cfd3f86b94f-1000--new-report-highlights-record-revenues-and-increasing-investment-into-european-football/; UEFA, *The European Club Finance and Investment Landscape*, 2024/25 edition, multi-club investment section, https://editorial.uefa.com/resources/028a-1a2f899177e2-b3619612eaa4-1000/uefaeuropeanclubfinanceinvestmentlandscape_150224.pdf?pubDate=20250322. ↩

8 "Article 94: Squad cost rule," UEFA Club Licensing and Financial Sustainability Regulations 2025, https://documents.uefa.com/r/UEFA-Club-Licensing-and-Financial-Sustainability-Regulations-2025/Article-94-Squad-cost-rule-Online; "Premier League clubs agree Squad Cost Ratio Rules," Premier League, https://www.premierleague.com/en/news/4468470. ↩

9 "New Opportunities for Investors in the Growing Women's Sports Sector," Morgan Lewis, July 2024, https://www.morganlewis.com/pubs/2024/07/new-opportunities-for-investors-in-the-growing-womens-sports-sector; "Closing the monetization gap in women's sports: A $2.5 billion opportunity," McKinsey, https://www.mckinsey.com/industries/technology-media-and-telecommunications/our-insights/closing-the-monetization-gap-in-womens-sports-a-2-point-5-billion-dollar-opportunity. ↩

10 Company references: Yardi, https://www.yardi.com/; CoStar Group, https://www.costargroup.com/; Palantir, https://www.palantir.com/; Anduril, https://www.anduril.com/; Harvey's valuation announcement, https://www.harvey.ai/blog/harvey-raises-at-dollar11-billion-valuation-to-scale-agents-across-law-firms-and-enterprises; Forbes reporting on Harvey reaching $190 million ARR by the end of 2025, https://www.forbes.com/sites/iainmartin/2026/02/09/legal-ai-startup-harvey-in-talks-to-raise-200-million-at-11-billion-valuation/; Business Insider reporting annualized revenue above $200 million, https://www.businessinsider.com/harvey-ceo-winston-weinberg-200-funding-round-2026-3. ↩

11 "The Financing Gap in Sports: Unlocking a $2.5 Trillion Opportunity," Apollo, December 2025, https://www.apollo.com/institutional/insights-news/insights/2025/12/the-financing-gap-in-sports-unlocking-a-dollar-2-5-trillion-opportunity. ↩ ↩2

12 Provider references: Hudl StatsBomb, https://www.hudl.com/products/statsbomb; StatsBomb teams page, https://statsbomb.com/who-we-help/sports-teams/; SkillCorner $60M release, https://skillcorner.com/articles/skillcorner-x-silversmith; Hudl Wyscout, https://www.hudl.com/products/wyscout. ↩

13 "Aerials," STATSWING Research, https://www.statswing.com/research/aerials/. ↩

14 Dr. Ian Graham quotation and transfer-failure framing, Analytics FC blog, October 18, 2021, https://analyticsfc.co.uk/blog/2021/10/18/measuring-transfer-success-through-minutes-played/. ↩

15 Ethan Mollick, "Centaurs and Cyborgs on the Jagged Frontier," *One Useful Thing*, 2023, https://www.oneusefulthing.org/p/centaurs-and-cyborgs-on-the-jagged; Dell'Acqua et al., "Navigating the Jagged Technological Frontier," SSRN, 2023, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4573321. ↩

16 "Anthropic Economic Index: AI and the Labor Market," Anthropic, March 2026, https://www.anthropic.com/research/labor-market-impacts. ↩

17 "Anthropic Economic Index: March 2026 Report / Learning Curves Update," Anthropic, March 2026, https://www.anthropic.com/research/economic-index-march-2026-report. ↩ ↩2

18 Jack Dorsey and Roelof Botha, "From hierarchy to intelligence," Block, March 2026, https://block.xyz/inside/from-hierarchy-to-intelligence. ↩

19 Ajibola Lawal, "Cut Your AI Coat Oversize This Year," Medium, 2026, https://medium.com/@SHIAMANI/cut-your-ai-coat-oversize-this-year-ae55dcab5c85. ↩

20 International Centre for Investigative Reporting author archive for Joel Adejola, https://www.icirnigeria.org/author/jadejola/. Other biographical claims are author-held and should be checked against the publication bio if needed. ↩

21 Ilter, "The 17% Gap," arXiv preprint, January 2026, https://arxiv.org/abs/2601.17431. ↩

22 C.A.J. Coady, *Testimony: A Philosophical Study*, Oxford University Press, https://academic.oup.com/book/25355. ↩

23 Elizabeth Fricker, article on LLM outputs, testimony, and instrument readings, *Inquiry*, 2025, https://www.tandfonline.com/doi/abs/10.1080/0020174X.2025.2553297. ↩

24 Coalition for Content Provenance and Authenticity, C2PA standard, https://c2pa.org/. ↩