

The AI Arms Race and what it means for Competition Law: A new era or new focus

14 April 202616 min read

We are not in the habit of writing breathless technology briefings. That is not our role.

But the industrial reorganisation now underway around artificial intelligence is arguably the most consequential structural shift for competition policy since the rise of the digital platform economy in the 2010s. It is also submitted that in some respects it is more challenging, because it is happening faster, across more sectors simultaneously, and with a degree of vertical integration that makes traditional market definition genuinely difficult.

After studying the PitchBook's December 2025 *AI Outlook: The Great Competition Wars Have Begun*^[1], a research report analysing various sectors, presenting unusually granular evidence of where capital, market share, and M&A activity are concentrating, we deemed it of sufficient insight to pen this article.

Our aim with this thought piece is to equip boards, general counsel, and deal teams with a clear assessment of where the competition law risk actually lies, not where the headlines suggest it does.

The Shape of the Problem: Who Controls What, and Why It Matters

The AI economy is not one market. It is a stack, and competition concerns differ sharply depending on which layer you occupy.

At the base sits the compute and datacentre layer, where the global build-out for AI infrastructure is approaching one trillion dollars annually. That figure alone tells us something important about barriers to entry: this is not a market that a well-funded startup can meaningfully contest from scratch. The capital requirements are enormous, the lead times for power and cooling infrastructure are measured in years, and the physical constraints, grid capacity, permitting, energy supply, are binding in ways that no amount of venture capital can overcome quickly. Clean energy integration, grid management, and distributed energy resource management are becoming necessary complements to datacentre scale, and the players best positioned to secure those inputs are, unsurprisingly, the ones who already have them.

Above compute sit the foundation models. This is the layer that, we submit, presents the most acute competition concerns. PitchBook's data confirms what practitioners have long suspected: foundation model providers have captured the bulk of AI deal value, capital formation remains "heavily concentrated" around a limited number of large-scale model builders, and these providers are *"becoming the default infrastructure for a growing share of enterprise AI workloads"*. Their scale, model performance, and integration depth give them what the report describes as *"clearer business durability than most application-layer startups"*.

The economic logic is straightforward and familiar from prior platform cycles: as inference becomes a recurring utility service, as workflows increase usage density, and as enterprises lock into multiyear computing platform commitments, switching costs compound and the installed base becomes progressively harder to displace.

For competition lawyers, the language of "default infrastructure" and "multiyear platform commitments" should trigger immediate recognition. These are the structural preconditions for dominance, not necessarily dominance today, but the kind of durable market power that, once established, is exceptionally difficult to unwind through ex post enforcement alone.

Above the model layer are the application, and sector-specific layers, where the dynamics differ but are no less important. Here, the PitchBook analysis reveals a pattern of rapid adoption by sector incumbents who are acquiring AI capabilities to defend entrenched positions, in healthcare, agriculture, cybersecurity, enterprise software, and e-commerce, among others. The competitive risk in these sectors is that AI will entrench the advantages of those who already control distribution, data, and customer relationships.

Barriers to Entry: The New Moats

We have spent years advising clients on what constitutes a durable barrier to entry. The AI cycle introduces some familiar moats in new guises, and at least one that is genuinely novel.

Data. The PitchBook report is emphatic that durable advantage in AI rests on *"unique data moats"* and deep integration into enterprise workflows. In enterprise SaaS, the marketing and analytics niches are *"saturated with undifferentiated tools,"* and the startups that lack proprietary data and workflow ownership face rapid commoditisation. In fintech, the CFO stack exhibits the same pattern: without proprietary data, domain depth, or workflow ownership, products converge on similar automation use cases and pricing pressure erodes margins. In agtech, the major crop science companies, Corteva, Bayer, Syngenta, have built precision platforms (Granular, Climate FieldView, Cropwise) that integrate machine and imagery data at scale, creating what the report calls *"the default enterprise solution"* while marginalising independents that lack integration into these dominant ecosystems. These are barriers: the more data a platform accumulates over time, the more formidable they become.

Compute and energy. We have already noted the trillion-dollar infrastructure build-out. What deserves emphasis is that access to compute is not simply a function of money. GPU supply is constrained. Energy capacity is constrained. Cooling technology is still maturing. The firms best positioned to secure these inputs are those with existing data centre footprints, long-term energy contracts, and the balance sheets to make speculative capacity commitments years in advance. Building a next-generation database or vector store is described by PitchBook as *"a capital-intensive R&D endeavour requiring patient, long-term capital"*. Warehouse robotics deployment remains *"expensive, limiting the domain to leaders such as Amazon and Walmart"*. We submit that in each instance, the barrier is not merely financial; it is also structural.

Distribution. In code generation, the dominant incumbents, Microsoft's GitHub Copilot and Cursor, enjoy *"unparalleled data and distribution advantages,"* a combination that the report expects will lead to *"widespread commoditization and value destruction for undifferentiated startups"*. In e-commerce, answer engines and platform partnerships are re-routing product discovery in ways that advantage the owners of those channels. Across the stack, AI models are consumed via APIs, making the infrastructure to secure, monitor, and scale those interfaces *"a non-negotiable, high-growth layer"*. Whoever controls the API layer controls the terms on which downstream innovation occurs, a distribution chokepoint that competition authorities will need to understand in technical detail.

Vertical Integration and the Risk of Foreclosure

The AI supply chain runs, in simplified terms, from chip manufacturers through cloud and compute providers, through model developers, down to application-layer companies and end users. At each handover point, there is a potential for vertical leverage, and the PitchBook evidence suggests that consolidation incentives are strong at every junctio

Consider the pattern across sectors:

In **cybersecurity**, AI protection capabilities are being rapidly absorbed by incumbent security platforms through M&A. The PitchBook report notes that the outlook for 2026 and 2027 *"remains strong in aggregate, though dominated by incumbent-led acquisitions rather than by standalone startups,"* and that the opportunity for independent startups *"is narrowing as these features become standard across application and cloud security suites"*. Once model-defence features are bundled into broader platform offerings, the incentive to disfavour third-party alternatives through licensing terms, API design, or marketplace placement is well understood in competition law.

In **healthcare**, second-tier AI scribe companies are expected to be acquired *"at highly discounted valuations"* by electronic health records incumbents or larger healthcare IT platforms. The EHR market is oligopolistic, Epic and Oracle Health dominate distribution, and the embedding of native AI scribing functions directly into these platforms creates a classic vertical foreclosure risk for independents that depend on EHR integration to reach clinicians.

In **agriculture**, Corteva, Bayer, Syngenta, and John Deere have vertically integrated precision agriculture platforms that connect directly to farmers, bypassing traditional retailer agronomists. Seven major retailers control approximately 60 to 90 per cent of crop input sales, and digital advisory is now embedded in the platforms of the crop majors themselves. Independent drone-based monitoring startups face what amounts to a distribution foreclosure problem: without integration into these dominant ecosystems, they cannot reach the customer base at scale.

In **e-commerce**, the emergence of LLM-native ad networks, the opening of Amazon's DSP and SSP infrastructure as a service, and the shift to answer engines as the primary discovery layer together represent a rewiring of the commerce stack that incumbents are best placed to shape. Consumer-facing, domain-specific search platforms face *"structural headwinds in attribution and monetization as horizontal platforms such as ChatGPT and Perplexity integrate commerce functionalities"*.

At the **model layer** itself, the prospect of enterprises locking into multiyear platform commitments creates the most systemic lock-in risk. Where those commitments bundle compute, safety tooling, agent management, and application accelerators into a single relationship, the switching costs can become prohibitive in practice even if they are not insurmountable in theory. Volume-based discount ladders and committed-spend rebates, familiar from cloud markets, are the mechanism through which this lock-in is likely to deepen.

The self-preferencing risk is real. Where a platform owner operates at both the model layer and the application layer, or controls both the inference utility and the safety/governance stack that wraps around it, the incentive to advantage affiliated products through scheduling priority, API design, or pricing structure is obvious. These are not speculative harms; they are the same theories that have sustained enforcement in digital platform markets, transposed to a new industrial context.

Merger Control: Speed, Serial Acquisition, and the Nascent Competition Problem

We turn now to what is, in our view, the most pressing institutional challenge: whether existing merger control frameworks can keep pace with the tempo of AI-driven consolidation.

The PitchBook evidence paints a vivid picture. The report describes a *"platform war"* in which incumbents are *"driven to acquire"* point solutions to *"plug immediate GenAI and security gaps,"* with acquisition characterised as *"the fastest route to market"*. This language matters because it describes a strategic logic, defensive ecosystem consolidation, that is precisely the kind of conduct that merger control exists to scrutinise.

Serial acquisitions are the dominant pattern in several sectors as detailed in the PitchBook report. In agtech, 2025 exits occurred *"entirely through M&A"* (38 transactions) and buyouts (five transactions), with zero IPOs. The acquirers are Corteva, Bayer, Syngenta, John Deere, consolidating digital and autonomous capabilities. In cybersecurity, AI protection is undergoing rapid consolidation through acquisition by major security vendors. In healthcare, PE-owned healthcare IT companies are driving AI-capability acquisitions, with R1 RCM's acquisition of Phare Health cited as a recent example. In enterprise software, incumbents such as GitHub and Atlassian *"must acquire these point solutions to plug immediate GenAI and security gaps"*.

Each of these transactions, taken individually, may fall below jurisdictional thresholds or appear competitively benign. Taken cumulatively, they can neutralise an entire stratum of nascent competition. This is the serial acquisition problem that competition authorities in the UK, EU, and US have been discussing for years but have yet to address with fit-for-purpose tools. The AI cycle may force the issue.

Kill acquisitions present a related but distinct concern. Where an incumbent acquires a startup not to integrate its technology but to prevent it from becoming a competitive threat, the competitive harm is the loss of future rivalry. The difficulty, as always, is evidentiary: proving that the target would have become a meaningful competitor absent the acquisition requires counterfactual analysis that courts and regulators find inherently speculative. The PitchBook data provides at least a circumstantial basis for concern: the report repeatedly identifies sectors where startups are commercially maturing, gaining traction with enterprise customers, and then being absorbed by the very incumbents whose market positions they threaten.

We do not suggest that every AI acquisition is anticompetitive. Many will be genuinely pro-competitive, accelerating diffusion and enabling product improvement. But the industrial dynamics documented in the PitchBook report, defensive acquisition strategies, ecosystem lock-in, and serial consolidation by a small number of strategic buyers, provide a strong case for authorities to invest in improved monitoring, refined nascent competition tools, and the willingness to deploy interim measures where integration risks creating irreversible structural harm.

Sector-Specific Flashpoints

The competition risks are not uniform across the economy. Some sectors warrant particular attention, either because AI adoption is especially rapid or because the pre-existing market structure amplifies concentration dynamics.

Healthcare. The combination of oligopolistic EHR incumbents, regulatory barriers to entry, high switching costs, and the foundational nature of ambientscribe technology creates a high-risk environment for vertical foreclosure. The anticipated absorption of independent companies into EHR platforms will concentrate the clinical workflow layer around a very small number of providers. In medtech, AI-powered imaging and smart implants are gaining traction, and incumbents are already acquiring remote monitoring platforms, Philips acquiring BioTelemetry, UnitedHealth acquiring Vivify Health, suggesting a pattern of vertical integration that could foreclose independent device and diagnostics innovators.

Enterprise software. The displacement of legacy analytics and BI platforms by agentic AI and natural language querying is a genuine paradigm shift, moving from dashboard navigation to conversational, in-workflow insights delivery. The incumbents that own the system of record, the CRMs, ERPs, and data warehouses that enterprises cannot easily replace, are strongly incentivised to acquire emergent AI search, compliance, and automation tools to keep customers captive. Whether that integration is conducted on open, non-discriminatory terms or through proprietary interfaces designed to lock out rivals will determine the competitive outcome.

E-commerce and advertising. The shift to answer engines as the primary product discovery mechanism is already measurable: Adobe Analytics reports a tenfold increase in referral traffic from answer engines, with Google conceding low-value queries. As OpenAI, Anthropic, and others seek scaled monetisation, LLM-native ad networks will emerge as a logical extension of the platform, with Amazon's and Walmart's retail media businesses providing the template. The risk is that a small number of answer engine operators will control the default pathways for commercial traffic in the same way that search engines did in the prior era, with the same attendant risks of self-preferencing, exclusionary conduct, and opaque ranking decisions.

Agriculture. This is a sector where vertical integration is already far advanced. Corteva, Bayer, Syngenta, and John Deere between them control precision agriculture platforms, digital advisory, crop input distribution, and equipment data pipelines. The PitchBook report states that independent drone-based crop monitoring faces value destruction risk as more than 50 competitors offer commoditised solutions while the major platforms *"maintain powerful industry moats by integrating machine and imagery data at scale"*.

Cybersecurity. AI protection for models and applications is both a growth market and a consolidation market. The 2025 acquisition wave, by SentinelOne, Palo Alto Networks, and other major vendors—signals sustained appetite for model-defence capabilities, with M&A identified as the most likely exit path for emerging vendors. The bundling of AI protection into broader platform suites is efficient, but it also risks foreclosing innovative niche solutions if platform APIs or marketplace placement disfavour third-party integrations post-acquisition.

Regulatory Adequacy and Emerging Theories of Harm

We should be candid about the limits of the evidence base. The PitchBook report is an investment and industry analysis, not a regulatory assessment. It does not catalogue specific enforcement actions by the CMA, the European Commission, or the DOJ and FTC, and we have not supplemented it with external enforcement data for the purposes of discussion. Any definitive assessment of agency responsiveness would require a broader evidentiary foundation.

That said, the market dynamics documented in the report have clear implications for the adequacy of existing frameworks.

Abuse of dominance and self-preferencing. The structural features of the foundation model layer, default infrastructure status, API-mediated access, inference-as-utility pricing, and multiyear platform commitments create fertile ground for discriminatory conduct. Where a platform owner controls both the inference utility and adjacent layers (safety tooling, agent orchestration, application accelerators), selective degradation of access for rivals can effectuate a margin squeeze that is invisible in headline pricing but decisive in competitive outcome. The evolution from single-shot inference to autonomous, multistep agent workflows will create new chokepoints around scheduling, memory allocation, and tool access that do not map neatly onto existing precedent. These are novel theories of harm in form, though cognate to established non-discrimination and refusal-to-deal doctrines in substance.

Essential facilities. We raise this with appropriate caution, because the legal threshold for essential facilities claims is high in all jurisdictions. But the PitchBook evidence on compute and energy constraints, trillion-dollar infrastructure, constrained GPU supply, grid capacity limits, suggests that access terms for capacity-constrained inputs will be competitively determinative in some markets. Where a small number of vertically integrated providers control GPUs, data centre space, and on-site generation, the analytical conditions for an essential facilities enquiry may be met.

Practical Implications

For companies at bottleneck layers. If you operate at a chokepoint in the AI stack, foundation models, compute, developer platforms, EHR systems, and precision agriculture platforms, you should expect regulatory scrutiny of your API terms, bundling practices, default settings, and data-access policies. Sensible risk mitigation includes transparent, stable, and non-discriminatory API terms; documented interoperability commitments; and governance structures that separate upstream access decisions from downstream competitive incentives.

For acquirers. Serial acquisition strategies in sectors such as cybersecurity, healthcare IT, and agtech are accumulating competitive significance even where individual transactions fall below enforcement thresholds. Internal documents, integration planning materials, and board presentations that describe acquisitions in terms of ecosystem defense, competitive neutralisation, or market foreclosure will be discoverable and will inform the analysis. Transaction teams should be briefed accordingly.

For startups and investors. The PitchBook evidence suggests that the exit environment in many AI subsectors is overwhelmingly M&A-driven, with strategic acquirers dominating. Where the acquirer is an incumbent with an existing dominant position, the transaction will attract greater scrutiny. Founders and their advisers should factor merger control risk into exit planning earlier than is conventionally the case, particularly in sectors where the buyer universe is narrow.


Concluding Observations

We do not think the sky is falling. The existing toolkit of competition law, abuse of dominance, vertical agreements, and merger control is conceptually adequate to address the risks we have described. The theories of harm are, in most cases, familiar ones applied to new industrial facts.

What concerns us is timing. The AI cycle is moving faster than prior platform cycles, the capital commitments are larger, and the vertical interdependencies are deeper. The window between the emergence of a competitive threat and its absorption by an incumbent is narrowing, and once multiyear platform commitments, default infrastructure status, and ecosystem lock-in have taken hold, the practical scope for ex post remediation shrinks considerably.

The law is not unsettled in its principles. It is unsettled in its application to the specific industrial facts of the AI stack, inference pricing, API-level discrimination, agent orchestration, and the competitive significance of sovereign capital. These are the areas where novel theories of harm are most likely to emerge, and where early engagement with regulators, careful compliance design, and thoughtful transaction structuring will deliver the greatest returns.

Food for thought.



Ahmore Burger-Smidt

Head of Regulatory

[View Profile](#)