BRIEFING · SEPTEMBER 2026

Artificial Intelligence in *UK BioTech*

What the evidence actually says, what the regulator expects, and what can safely be done now.

Neil Catton

Independent Technology Strategist · Technology, Ethics & Human Purpose

NEILCATTON.COM

Prepared as a market briefing. Every figure is traced to a regulator, a trade body, published legislation, named academic research, or an operator speaking on its own record.

ONE

A note on the numbers

The most repeated statistic in this sector was chased down and read in full before this briefing went to print. It is excluded anyway — not because it could not be found, but because of who produced it.

The claim is that AI-discovered drugs achieve a Phase I success rate of 80 to 90 per cent, against a historical industry norm of around 40 to 65 per cent. It is attributed — correctly — to a named 2024 analysis in *Drug Discovery Today* by Jayatunga, Ayers, Bruens, Jayanth and Meier. The paper was obtained and read in full, and the figure is real: 21 of 24 AI-discovered molecules that had completed Phase I trials by December 2023 succeeded.

It is excluded anyway. All five authors are Boston Consulting Group employees, and the paper states its research was funded by BCG's Health Care practice — the consultancy that advises pharma companies partnering with the biotechs it studied. Nor does the figure travel: the same analysis found a 40 per cent Phase II rate, "in line with historical industry averages," in the authors' own words. A number that beats history at Phase I and matches it at Phase II, produced by the firm paid to produce it, is not an independently verified finding.

What was excluded, and why

- → **The Phase I success rate above.** Read in full. Excluded because its five authors are employed by, and its research was funded by, the consultancy that produced it — and because the same analysis' own Phase II figure reverts to the historical average.
- → **"173 AI drugs in trials, zero approved."** Traced to a blog and a commercial content site. No methodology, no definition of "AI drug".
- → **Every AI discovery platform's own timeline and cost claim.** Excluded wholesale.
- → **All "AI will save $X billion in pharma R&D" projections.** Consultancy or vendor in every case located.
- → **Cell and gene therapy capacity and workforce figures from 2023.** Superseded; current surveys exist and were not obtained. A three-year-old figure presented as current is the error this series most often catches in others.
- → **Isomorphic Labs' scientific and platform claims.** Its £1.6bn Series B is a trade body financing fact and is used. Section two also sets out the company's own published claims for its Drug Design Engine — named throughout as the company's own account, not independently verified, rather than excluding the platform altogether.
- → **Quantified evasion rates from the biosecurity study in section six.** The paper is paywalled and the figures circulating in secondary coverage disagree with each other. The finding is used qualitatively with its full citation.
- → **A claim that the MHRA has backed a global framework for AI in pharmacovigilance rather than issuing UK guidance.** The only source is a paywalled trade headline. It would have been a significant finding. It is not verifiable and it is out.

Read in full, and excluded anyway

Jayatunga, Ayers, Bruens, Jayanth and Meier, "How successful are AI-discovered drugs in clinical trials? A first analysis and emerging lessons", Drug Discovery Today, 2024, article 104009, DOI 10.1016/j.drudis.2024.104009 — read in full before publication. Excluded on sponsorship grounds, not access grounds: see above. Current Cell and Gene Therapy Catapult survey data and a primary Genomics England position on AI were also sought and not obtained; neither carries a claim here.

One figure is attributed rather than primary

A life sciences AI adoption rate of 48.4% appears in a UK government publication, cited to ONS data from December 2025. It could not be located in the published ONS release, whose own headline figure is around 35% for businesses with ten or more employees and which publishes no life sciences breakdown. Where it is used here, it is attributed to the government document that carries it, with that discrepancy stated.

Sources for every section are given at the foot of that section and listed in full in section thirteen.

TWO

A record quarter, and its denominator

UK biotech venture investment reached a five-year high in the second quarter of 2026. One deal accounts for three quarters of it.

£2.05bn

UK biotech venture investment, Q2 2026 — the strongest quarter in five years

£498m

the same quarter excluding one AI drug discovery company's £1.6bn Series B

The BioIndustry Association reported the quarter as a genuine recovery, and on the underlying numbers it was one: £498m against £279m in the comparable quarter a year earlier is a real improvement, with eight seed deals averaging £6.4m, £190m in Series A and £225m in Series B and later. The UK took 61% of the £3.3bn European venture total, though that share is itself a function of the single deal.

Two things sit alongside it. There were no UK biotech initial public offerings in 2026, and no NASDAQ follow-on activity from UK biotech companies. And the largest single act of confidence in the sector was one company's bet on AI-driven drug discovery, not a broad reallocation across the industry.

WHY THE DENOMINATOR MATTERS TO A BOARD

A headline that reads "UK biotech investment hits five-year high, driven by AI" is accurate and misleading at the same time. A board reading it as evidence that the sector has funded a shift to AI is reading a single transaction. The question to ask of any sector adoption claim in this briefing, including the government's own, is what the denominator is.

The sector underneath it

UK biopharmaceuticals comprise 2,960 companies employing 163,600 people with £98.9bn turnover — two thirds of life sciences turnover from 40% of its companies. The Office for Life Sciences designates these "official statistics in development", says they are not comparable with 2021/2022 or earlier because of a methodology change, and states that "limited information is known about its coverage". They are the best figures available and this briefing does not treat them as precise.

What the money actually built

The £1.6bn figure above answers a financing question, not what an AI drug-discovery platform does day to day — and the sector's own regulator is little help here. The MHRA's principal statement on AI in medicines, already this briefing's source for the regulatory-strategy material in section seven, gives the discovery pipeline itself one sentence: AI may change "the pace at which new medicines can be developed... through the use of AI to reduce the costs and time invested in failed molecules." Everything else it publishes concerns AI as, or inside, a regulated product. Four concrete cases stand in for that silence — one live and adopted, one a claim about its own technology now backed by two paying pharmaceutical partners, one a different outcome from the same period, and one company whose own record holds the sector's clearest verified success alongside an inconclusive result.

THE PLATFORM, LAUNCHED THIS YEAR

Mystra AI: a chat interface onto 45,000 genetic studies

Genomics — a science-led, Oxford- and London-based TechBio with a US office in Durham, North Carolina — launched Mystra AI on 2 June 2026: a conversational platform letting researchers query its Foundational Data Collection, which it states spans over 45,000 genome-wide association studies and trillions of rows of data, to identify and prioritise drug targets in plain language. Its stated rationale: "targets with human genetic support are 2.6 times more likely to succeed in clinical trials," against an industry backdrop it puts at 95% of candidates failing. Named adopters at launch include BridgeBio Pharma and Relation Therapeutics, with an endorsement from AWS; a BridgeBio associate director is quoted saying the platform is "accelerating our day-to-day workflow by making it much easier and more efficient to mine genomics data." A live, adopted product — but the adoption and success-rate figures are the company's own account of itself.

THE CLAIM, AND WHO IS NOW PAYING FOR IT

Isomorphic Labs' Drug Design Engine

Isomorphic Labs — the company behind this briefing's £1.6bn figure — published its own account of the technology on 10 February 2026: a "unified computational drug-design system" it calls IsoDDE, which it states goes beyond AlphaFold 3 rather than simply building on it, claiming more than double AlphaFold 3's accuracy on challenging protein-ligand prediction, a 2.3x improvement on antibody-antigen structure prediction, and the ability to identify novel druggable binding pockets from a computer in seconds. This is Isomorphic's own claim about its own platform, not independently verified — treated here exactly as this briefing already treats Isomorphic's other scientific and platform claims: used, but named as such.

What is independently verifiable is who has since paid to use it. On 7 January 2024 Isomorphic signed two named multi-target research collaborations, confirmed directly from each company's own announcement: with Eli Lilly, for $45m upfront and up to $1.7bn in milestone payments, and with Novartis, for $37.5m upfront and up to $1.2bn. The platform's technical claims remain unverified. Two named pharmaceutical companies choosing to pay for access to it is a different, and independently checkable, kind of evidence.

THE OTHER BET, AND WHAT HAPPENED TO IT

Exscientia no longer exists as an independent company

In the same period that produced this quarter's record investment, the sector's other AI-native drug-discovery pioneer had a different outcome. Exscientia — Oxford-founded, and among the earliest companies to build the AI-drug-discovery category this briefing discusses — completed its merger into the US company Recursion on 20 November 2024. The combined, roughly 800-person organisation is headquartered in Salt Lake City; the Oxford office remains one of its stated sites, but Exscientia no longer operates as an independent UK company. One AI-native pioneer was scaled through a £1.6bn round in the reporting quarter; the other was absorbed into a US competitor in the two years before it.

The merger was not a sudden reversal of an otherwise clean run. A year earlier, in October 2023, Exscientia halted the Phase I/II trial of EXS-21546, an AI-discovered immuno-oncology candidate, after external data on the same drug target suggested the mechanism would struggle to reach a workable therapeutic index. The company's own stated reason was scientific rather than financial — but it means the AI-native pioneer had already wound down a clinical asset before the business itself was absorbed.

A FOURTH NAME, AND ITS MOST BALANCED CASE

BenevolentAI: one platform, two outcomes

London-headquartered BenevolentAI belongs in this briefing for the strongest independently verifiable AI-drug-discovery result available anywhere in this sector. Its knowledge-graph platform first flagged baricitinib, an existing rheumatoid-arthritis drug, as a candidate COVID-19 treatment, publishing the hypothesis in *The Lancet* on 4 February 2020. Eli Lilly ran the Phase 3 trial; the FDA granted Emergency Use Authorisation on 20 November 2020 — nine months from a computational hypothesis to authorised clinical use, later converted to full approval. This is not the company's account of itself. It is an FDA regulatory action and a trial run by one of the world's largest pharmaceutical companies.

The same platform's own internally discovered candidate tells the other half. BEN-2293, a topical treatment for atopic dermatitis, reported Phase IIa results on 5 April 2023 that met its safety endpoint — "safe and well-tolerated" — but were, in the company's own words, "not conclusive" on efficacy. Repurposing an existing, already-approved molecule succeeded in nine months. Discovering a new one from scratch has not, so far, worked. Both outcomes are from the same company and the same underlying technology, and both are named here rather than the more flattering half alone.

None of this is the full data-analytics picture. Our Future Health and UK Biobank — the population-scale genetic, NHS-record and lifestyle datasets covered in full in this briefing's HealthTech companion — are also the data layer that AI-driven biomarker and target-discovery research in this sector increasingly draws on: a biotech company mining Mystra AI or building its own models is, in practice, often working from the same underlying cohorts a hospital trust draws on for predictive care. That fuller treatment of the data itself is not repeated here; what belongs in this briefing is the observation that discovery-side AI and care-side AI are increasingly drawing on the same wells.

THE WHOLE CAST, IN ORDER

- → **February 2020** — BenevolentAI publishes the baricitinib hypothesis in *The Lancet*.
- → **November 2020** — the FDA grants Emergency Use Authorisation for baricitinib.
- → **April 2023** — BenevolentAI's own BEN-2293 reports safe but inconclusive Phase IIa results.
- → **October 2023** — Exscientia halts the EXS-21546 trial.
- → **January 2024** — Isomorphic Labs signs its Eli Lilly and Novartis collaborations.
- → **November 2024** — Exscientia's merger into Recursion completes.
- → **February 2026** — Isomorphic Labs publishes its Drug Design Engine claims.
- → **June 2026** — Genomics launches Mystra AI.

Four years, four companies, and a mixture of verified success, unverified claim, clinical setback and outright absorption — running well either side of the single quarter this section opened with.

Read together, the four cases above sit either side of the record quarter this section opened with. One platform is live, adopted, and grew out of a company that was not this quarter's headline deal. One platform is the headline deal's own account of itself — now with two named pharmaceutical partners paying to find out whether the claim holds. One pioneer of the category no longer exists as a UK company at all, and had already wound down a clinical asset before it was absorbed. And one company's own record holds both the sector's clearest verified success and an inconclusive result, from the same technology. A single quarter's investment figure captures none of that variation — which is the point the denominator warning above was making before any of these four names were on the page.

BioIndustry Association, *UK biotech financing April – June 2026*, published 20 July 2026. Office for Life Sciences, *Bioscience and health technology sector statistics 2023 to 2024*, published 2 October 2025. MHRA, *Impact of AI on the regulation of medical products*, 30 April 2024. Genomics, *Genomics Launches the Mystra AI Platform*, 2 June 2026. Isomorphic Labs, *The Isomorphic Labs Drug Design Engine unlocks a new frontier beyond AlphaFold*, 10 February 2026. Isomorphic Labs, strategic collaboration announcements with Eli Lilly and with Novartis, both 7 January 2024. Recursion Pharmaceuticals, merger completion announcement, 20 November 2024. Exscientia pipeline update reported by Fierce Biotech, 13 October 2023. BenevolentAI, press releases on the FDA Emergency Use Authorisation for baricitinib, 20 November 2020, and on Phase IIa results for BEN-2293, 5 April 2023. Eli Lilly and Company, investor news release on the baricitinib Emergency Use Authorization, 20 November 2020. Our Future Health, cohort profile, *International Journal of Epidemiology*, 15 October 2025.

THREE

The pressure point

The MHRA has started receiving responses to inspection findings that were written by a machine, and it has said so publicly.

On 29 June 2026 the MHRA Inspectorate published a post on its official blog titled "Use of AI for GxP inspection responses: setting standards without stifling innovation". It is not a consultation, a guideline or a warning letter. It is the regulator describing what it has been sent.

> THE MHRA INSPECTORATE, 29 JUNE 2026
>
> *"AI can support better regulatory outcomes and improve patient safety. However, we have also encountered responses containing references to MHRA guidance that doesn't exist, citations of inappropriate regulatory frameworks, and responses to serious deficiencies that appear designed to mislead rather than address underlying problems."*

The post gives a worked example. An AI-generated response to a serious patient safety deficiency contained material inaccuracies that delayed the resolution of serious compliance failures, and expanded the regulatory review required from four hours to more than twenty.

Why this is the finding, and not a curiosity

An inspection response is the mechanism by which a compliance failure gets corrected. A finding is raised, the organisation responds, the regulator accepts or challenges, the fix happens. If the response itself is unreliable, the correction loop breaks at the point where it matters most — and in the documented case, it broke on a serious patient safety deficiency.

This is also not an attack from outside. It is a governance failure inside a regulated organisation, occurring in a document nobody had classified as a technology risk. There is no AI system on any risk register described as "the tool that drafts our regulatory correspondence".

> WHAT HAS NOT CHANGED
>
> "Organisations have always been responsible for the accuracy of their responses to the MHRA and persons responsible for them have always been expected to verify factual claims." The Inspectorate's stated requirements are that a submission is factually accurate and verifiable, technically reviewed by experienced personnel, signed off by someone with authority and accountability, supported by evidence, and appropriate to the regulatory context — "regardless of how you draft your submissions".

That is a defensible regulatory position and this briefing does not criticise it. Accountability has not moved, and the MHRA is right that it has not. But it is enforced through inspection rather than guidance, which means the burden of demonstrating that a review actually happened sits entirely with the organisation, with no published standard to point at.

MHRA Inspectorate, "Use of AI for GXP inspection responses: setting standards without stifling innovation", 29 June 2026. GxP covers good manufacturing practice, good clinical practice and good pharmacovigilance practice.

FOUR

What the sector is facing

The UK rewrote its clinical trials law for the first time in two decades, and the instrument does not address the technology.

The Medicines for Human Use (Clinical Trials) (Amendment) Regulations 2024 were signed into law on 10 April 2025 and applied from 28 April 2026, revised from an original date of 10 April 2026 because of a technical processing issue. They are good law. They embed Combined Review — a single application for ethics and regulatory approval — introduce a risk-proportionate approach, and create for the first time a statutory duty to register a trial on a WHO-recognised public register and publish its results.

Janet Messer, Director of the Approvals Service at the HRA, described the changes as reflecting "our commitment to making it easier to do high-quality research that people can trust". That is a fair description.

Fourteen parts, and no mention

The instrument runs to fourteen parts, amending everything from ethics committees and trial authorisation to labelling, enforcement and the schedules. No Part, regulation or heading in it addresses artificial intelligence, automated processing or algorithmic decision-making.

> THE DETAIL THAT SHARPENS IT
>
> Part 6 amends the pharmacovigilance provisions directly. Pharmacovigilance is the one place in the medicines lifecycle where AI is already doing production work — case intake, duplicate detection, coding, signal triage — inside a statutory safety obligation. The legislator had that section of the regulations open on the table and did not address the technology already operating inside it.

This briefing states that absence at the level it was checked: the instrument's published contents and headings. It is not a claim about every line of the text, and it is not a criticism of the drafters, who were completing a reform begun years earlier.

Six bodies touch this, and only one regulates a company's use of AI

The MHRA is not the only agency in this picture. Six bodies have a hand in AI and UK bioscience, and the gap between them is not that nobody is watching — it is that each one watches a different thing, and only one of them regulates a company's use of AI at all.

MHRA	Regulates the medicines lifecycle and inspects GxP. Its published position on AI in submissions is an inspectorate blog post. Its AI strategy states roughly three full-time equivalents rising to about seven and a half for AI product regulation.
HRA	Oversees clinical trials of AI health technologies in the NHS, through Research Ethics Committees and the Confidentiality Advisory Group.
HSE	Statutory regulator for contained use of genetically modified organisms under the 2014 Regulations. Premises must be notified before first use; classes 2 to 4 are notifiable or require consent. Silent on AI. DEFRA and the devolved governments handle deliberate release.
DSIT	Owns the UK Screening Guidance on Synthetic Nucleic Acids — the UK document standing closest to AI-designed biology — and, through the AI Security Institute, publishes measurements of what frontier models can do in biology.
UKHSA	Owns biosurveillance, microbial forensics and diagnostic scale-up — the Biothreats Radar used across fourteen government departments, metagenomic surveillance through mSCAPE, and wastewater surveillance. Not a regulator of biotech companies.

ICO	The one body whose remit already covers developing and using AI, because every AI system processes personal data. Required by statute since 12 May 2026 to produce a code of practice on artificial intelligence and automated decision-making. Section seven sets out what that means.

Two distinctions matter here. UKHSA is the agency that would *detect* a consequence — a national biosurveillance and microbial forensics capability, funded and expanding — which is a different job from preventing one. And the body writing the UK's actual AI rulebook is the data regulator, not the medicines regulator. Nothing in the first five rows places a company's use of AI inside a regulatory perimeter. The sixth does.

What the state is actually doing about AI in biology

The absence of a regulatory perimeter is not the same as inattention, and it would be unfair to the agencies above to imply that it is. The government is measuring the problem, and publishing what it measures.

The AI Security Institute, part of DSIT, published its first Frontier AI Trends Report on 18 December 2025. It records that models now exceed PhD-level baselines on chemistry and biology knowledge, "can now consistently produce detailed and accurate protocols for a range of complex scientific tasks", and are "up to 90% better than human experts at providing troubleshooting support for wet lab experiments". Non-experts using frontier models to write protocols for viral recovery had 4.7 times higher odds of producing a feasible one. AISI is careful about what that means: its evaluations are "controlled, task-based settings" that "may not reflect or generalise to real world effectiveness", and are "a snapshot in time".

The Cabinet Office's July 2026 implementation report commits £115 million to UK defences against AI-related threats including biological security, commits to a second AISI trends report, and commits to "scope options to legislate... including to proportionately mitigate dual-use risks from AI-enabled biological capabilities".

> A RISK THIS BRIEFING HAD NOT OTHERWISE REACHED
>
> The same report notes that biological datasets "may be vulnerable to 'poisoning attacks' that corrupt training data, while laboratory automation systems and connected research infrastructure may be exposed to hardware manipulation." For an organisation running automated laboratory workflows, that is an information security question sitting inside a science budget — and it is unlikely to be owned by either function.

The Medicines for Human Use (Clinical Trials) (Amendment) Regulations 2024, contents as published on legislation.gov.uk. Health Research Authority, 10 April 2025 and 26 September 2025. MHRA, *Impact of AI on the regulation of medical products*, 30 April 2024, last updated 31 July 2026. HSE, *GMOs and the law*, page updated 19 May 2026. AI Security Institute, *Frontier AI Trends Report*, 18 December 2025. Cabinet Office, *UK Biological Security Strategy Implementation Report*, 14 July 2026.

FIVE

The challenges of using AI here

Five risks, each drawn from a regulator's own account, published legislation, or the sector's own data.

ONE — THE REGULATORY RECORD

The correction loop is where it breaks

An inspection response is the mechanism by which a compliance failure gets fixed. If that response is unreliable, the loop is broken at exactly the point where it matters. The MHRA has documented a case where it cost time on a serious patient safety deficiency, turning four hours of review into more than twenty. This is the risk a board can act on this quarter, and the control that answers it costs almost nothing.

TWO — PHARMACOVIGILANCE

Production work under a statutory duty

Case intake, duplicate detection, coding and signal triage are where AI already does real work in this sector, inside a legal safety obligation with an inspection regime attached. The clinical trials rewrite amended the pharmacovigilance provisions and said nothing about it, and no UK guidance sets out what a compliant deployment looks like. An organisation running AI here is operating against an inspection expectation, not a standard.

THREE — NO STANDARD TO POINT AT

The burden sits entirely on the sponsor

With no UK framework, an organisation cannot demonstrate compliance by reference to a published standard. It has to construct a defensible position itself, hold it, and prove on request that the human review was real rather than nominal. Section 80 of the Data (Use and Access) Act 2025, in force since 5 February 2026, makes the quality of that review the thing the law turns on.

FOUR — THE CONCENTRATION RISK

One deal is not a sector position

Three quarters of a record quarter went to one company. Around 70% of AI-related investment over five years concentrated in oncology, neurology and COVID-19. There were no UK biotech IPOs in 2026. A board reading the headline as evidence of a broad, funded shift across the industry is reading a single bet in three therapeutic areas.

FIVE — THE CONSTRAINT NOBODY LEGISLATES FOR

Your own data stops you before the law does

UKHSA told its own Advisory Board in May 2024 that "labelling and classification of data as well as permissions are not currently sufficiently well controlled to grant access to 3rd party LLMs", and put six to eight months of data labelling work between itself and a commercial assistant. A year later it still named "fragmented data structures, inconsistent labelling and classification of data, and lack of consistent permissions" as what obstructs enterprise AI. Its auditors rated the controls moderate, noting that "limited guidance from central government made the frameworks immaturity understandable."

That is a public body being unusually straight about a problem the private sector shares and rarely publishes, and the government's own life sciences AI plan names non-AI-ready data as its first barrier. An organisation that cannot say what its data is classified as, or who may see it, cannot lawfully put a model near it — whatever the ICO's code eventually says.

MHRA Inspectorate, 29 June 2026. Clinical Trials Amendment Regulations 2024, Part 6. Data (Use and Access) Act 2025, section 80. BioIndustry Association, Q2 2026. *AI Adoption Plan: Life Sciences*, 8 June 2026. UKHSA Advisory Board papers, May 2024 and March 2025.

SIX

The threat side

Two threats. Neither is an attack, and both are internal to the organisation.

The first: the record itself

Section three sets it out. In every previous briefing in this series the threat was something done to the organisation — fraud, impersonation, better-equipped counterparties. Here it is the organisation's own drafting process, in a document class nobody had treated as a technology risk. The fix is procedural and cheap: name who verifies each regulated submission, check the citations, and record that both happened. That control would have caught the documented failure.

The second: a voluntary standard, and a published way past it

In October 2025 *Science* published work by Wittmann, Alexanian, Bartling, Beal, Clore, Diggans, Flyangolts, Gemler, Mitchell, Murphy, Wheeler and Horvitz on nucleic acid biosecurity screening. The researchers asked whether AI protein design tools could produce redesigned variants of proteins of concern that preserved function while evading the screening systems used by DNA synthesis suppliers. Eric Horvitz, Chief Scientific Officer at Microsoft Research, summarised the result: "The answer to that question was yes, they could." Patches were subsequently developed and distributed, greatly improving detection; the authors describe that as an initial step rather than a fix.

The UK has a standard those suppliers are expected to meet. DSIT publishes the *UK Screening Guidance on Synthetic Nucleic Acids for Users and Providers*, first issued on 8 October 2024 and last updated on 2 September 2026 — six days before the date of this briefing. It asks providers to screen customer identity and legitimacy, and to screen sequences: "at a minimum, DNA or RNA molecules (single or double-stranded) 50 nucleotides or longer should be screened."

Voluntary	The guidance sets out "expected standards of responsible practice" and is framed throughout as recommendations, not legal requirements. The government has committed to "determine if guidance for nucleic acid synthesis screening should be moved to a statutory footing". No date is attached to that decision.
One mention	The guidance refers to AI exactly once, noting that increased access to synthetic nucleic acids "coupled with developments in converging tools and technologies – such as AI – and an increase in knowledge and expertise, may exacerbate the risks." It does not address AI-designed sequences.

WHY THIS IS A BOARD QUESTION AND NOT A POLICY QUESTION

Every organisation that orders synthetic DNA relies on its supplier's screening, and there is now a published UK standard to hold that supplier to. Three questions cover it: does the provider screen customers and sequences to the UK guidance at 50 nucleotides and above, are the published biosecurity patches applied, and when was that last confirmed. The guidance does not tell a buyer how to verify any of it, which is precisely why the buyer has to ask.

This briefing does not reproduce the study's quantified evasion rates. The paper is paywalled and the figures in secondary coverage disagree with one another; the numbers are named in section one as an exclusion.

Wittmann, B. J., et al., including Horvitz, E., "Strengthening nucleic acid biosecurity screening against generative protein design tools", *Science* 390, 82–87 (2025). DSIT, *UK Screening Guidance on Synthetic Nucleic Acids*, 8 October 2024, updated 2 September 2026. AI Security Institute, *Frontier AI Trends Report*, 18 December 2025. Cabinet Office, *UK Biological Security Strategy Implementation Report*, 14 July 2026.

SEVEN

The UK legal position

There is no UK rulebook for AI in the medicines lifecycle. There is a rulebook for AI and personal data, it is being written now, and every organisation in this briefing is inside it.

In force	**Clinical Trials Amendment Regulations 2024.** Applies from 28 April 2026. Combined Review in law, a risk-proportionate approach, and a statutory duty to register trials publicly and publish results. No AI provision in any Part or heading.
In force	**UK GDPR Articles 22A–22D**, inserted by section 80 of the Data (Use and Access) Act 2025 from 5 February 2026. For most significant automated decisions the full range of lawful bases now opens up, including legitimate interests, subject to four safeguards — the person must be informed, be able to make representations, obtain human intervention, and contest the decision.
Unchanged	**The exception that matters in this sector.** Where a significant automated decision is based entirely or partly on special category data — which includes health data — the prohibition still applies, and the decision still requires a specific Article 9 basis such as explicit consent.
Loosened	**Scientific research provisions.** Research is now defined as any that "can reasonably be described as scientific, whether publicly or privately funded". Broad consent to an area of scientific research is permitted, and personal data may be re-used for research without a privacy notice where that would involve disproportionate effort, subject to safeguards.
In force	**SI 2026/425.** The Data Protection Act 2018 (Code of Practice on Artificial Intelligence and Automated Decision-Making) Regulations 2026. Made 16 April 2026, in force 12 May 2026. Requires the Information Commissioner to prepare a code of practice on developing and using artificial intelligence, and on automated decision-making.
Standing	**GxP inspection expectations.** Submissions must be factually accurate and verifiable, technically reviewed by experienced personnel, and signed off by someone with authority and accountability.

Two things the sector has read the wrong way round

The first is section 80. It is widely described as having made automated decision-making permissible. That is true in general and misleading here. The prohibition was narrowed, not removed, and what it now catches is precisely the decision based on special category data. In a sector whose data is overwhelmingly health data, the liberalisation largely does not reach the decisions that matter. The research provisions moved the other way, and moved genuinely: broad consent and a wide statutory definition of scientific research are real enablers for commercial biotech. The regime tightened and loosened at once, in different places.

The second is where the framework is coming from. The sector is watching the MHRA and the National Commission; neither is producing a code with statutory backing. The Information Commissioner is, under a duty in force since May 2026, and is separately engaging eleven foundation model developers on the lawfulness of their training data. Whatever that code says will apply to every AI system in this briefing, because every one processes personal data.

> WHAT IS STILL REQUIRED, AND WHAT IS NOT
>
> **Required:** verification. Every factual claim and citation confirmed by a named person with authority, who signs — an expectation that predates AI and applies however a submission was drafted.
>
> **Not required:** disclosure. Nothing obliges a sponsor to declare a model was involved, so an inspector cannot ask about a system they do not know exists.

legislation.gov.uk. Data (Use and Access) Act 2025, s.80, in force 5 February 2026. SI 2026/425, in force 12 May 2026. ICO, *The DUAA 2025 — what does it mean for organisations?*, updated 19 June 2026, and *AI and biometrics strategy update*, March 2026. MHRA Inspectorate, 29 June 2026.

EIGHT

Beyond the UK

Every other major medicines regulator has written down how AI may be used across the lifecycle. A UK sponsor will be read against those documents whether or not it built to them.

The European Medicines Agency, September 2024

The EMA adopted a reflection paper on the use of artificial intelligence in the lifecycle of medicines (EMA/CHMP/CVMP/83833/2023) on 30 September 2024, following consultation from July to December 2023. It addresses AI and machine learning "at any step of a medicines' lifecycle, from drug discovery to the post-authorisation setting", for human and veterinary medicines. It is explicitly non-binding — a reflection paper sets out current thinking rather than mandatory requirements.

The EMA and the FDA, January 2026

In January 2026 the two agencies jointly published *Guiding principles of good AI practice in drug development*. Ten principles: human-centric by design; a risk-based approach; adherence to standards; clear context of use; multidisciplinary expertise; data governance and documentation; model design and development practices; risk-based performance assessment; life cycle management; and clear, essential information. On the risk-based approach, the document states that AI technologies should "follow a risk-based approach with proportionate validation, risk mitigation, and oversight based on the context of use".

They are advisory, framed to "inform, enhance, and promote" good practice and to underpin future guidance across jurisdictions. Olivér Várhelyi, European Commissioner for Health and Animal Welfare, described them as "a first step of a renewed EU-US cooperation".

> WHAT THIS BRIEFING DOES NOT CLAIM
>
> The MHRA is not named as a party to the EMA and FDA principles. This briefing states what the document is and who published it. It does not assert that the UK declined to participate, because that is a different claim and it has not been established.

The practical response

A UK organisation running trials or filing dossiers in the European Union or the United States will be assessed against those principles whether or not it built to them. The response is not to wait for a UK equivalent, whose timing nobody can predict. It is to adopt the ten principles as an internal standard now. They are published, free, written by the regulators a UK sponsor most often files with, and adopting them costs nothing but the decision to do so — which also gives an organisation something to point at when the UK framework eventually arrives.

EMA, *Reflection paper on the use of artificial intelligence in the lifecycle of medicines*, EMA/CHMP/CVMP/83833/2023, published 30 September 2024. EMA and FDA, *Guiding principles of good AI practice in drug development*, January 2026.

NINE

Today

Seven things that can be done now. None requires a UK framework to exist first, and all of them are things an inspector could ask about tomorrow.

01	**Find out where AI is already drafting**	Regulatory correspondence, deviation reports, CAPA responses, safety narratives, protocol text. Ask the question in the departments that write to the MHRA, not the ones that buy software. The documented failure happened in a document class nobody had on a technology risk register.
02	**Put a named verifier on every regulated submission**	Not a sign-off box. A named person who confirms every factual claim and citation, with the confirmation recorded. This is the single cheapest control in this briefing and it answers the MHRA's stated expectation directly.
03	**Check the citations, specifically**	The documented failure was fabricated references to guidance that does not exist and citation of inappropriate frameworks. A citation check on regulated submissions is a short, mechanical control that catches precisely the thing the regulator reported receiving.
04	**Adopt the EMA and FDA ten principles as an internal standard**	Published, free, and written by the regulators you file with. Adopting them costs nothing, gives an internal reference where none exists, and positions the organisation well whenever a UK framework arrives.
05	**Map AI in pharmacovigilance against the statutory duty**	Case intake, duplicate detection, coding, signal triage. Establish what is automated, what a human actually reviews, and what evidence exists that the review happened. This is where a legal safety obligation and an unregulated technology already overlap.
06	**Hold your DNA supplier to the DSIT guidance**	Ask whether they screen customers and sequences to the UK guidance at 50 nucleotides and above, whether the published biosecurity patches are applied, and when that was last confirmed. The guidance is voluntary and does not tell a buyer how to verify it, so the buyer has to ask. If it moves to a statutory footing, you will already have the answer.
07	**Get the data permissions in order first**	UKHSA found its labelling and permissions too weak to let a third-party model near its data, and was still naming the same problem a year later. Before any enterprise AI decision, establish what your data is classified as and who is permitted to see it. It is the constraint that binds first, and the ICO code will assume you have done it.

None of the seven requires new procurement, a vendor conversation, or the outcome of any pending regulatory process.

Derived from the evidence in sections three to eight.

TEN

Tomorrow

Three levels, defined by where accountability sits rather than by how capable the system is.

The government's own AI Champion for life sciences states the same principle plainly: "the final accountability, ethical judgment, and execution will always rest with a human professional." The three levels below are a way of holding a board to that sentence.

LEVEL ONE

Assistive — the person decides

Prepares, summarises, surfaces. Accountability unchanged. Literature review, document drafting, data wrangling. Almost all defensible value sits here today — and this is exactly where the inspection-response failure happened, because nobody treated drafting as a regulated activity. The first level is not the safe level. It is the level where the control is assumed rather than designed.

LEVEL TWO

Augmentive — the person confirms or overrides

Recommends within bounds, with reasoning and confidence, and the override is recorded. Signal triage, case prioritisation, eligibility screening. The review must be demonstrably meaningful, which since 5 February 2026 is a legal requirement rather than a design preference.

LEVEL THREE

Adaptive — the envelope is what is overseen

Decides within a defined envelope. Oversight moves to the envelope; accountability moves to whoever sets and monitors it. This is a reasoned projection, not an observation. No current UK framework would support it inside a GxP-regulated process.

Three horizons, and an honest caveat

Nought to twelve months. The MHRA has stated an expectation publicly. The next step in that pattern is inspectors asking the question routinely, and findings written where the answer is weak. An organisation with a named verifier and a citation check answers it in a sentence.

One to three years. The ICO code of practice lands, and with it the first UK document with statutory backing that tells an organisation how to develop and use AI on personal data. Alongside it, either the UK writes a medicines-specific framework or the EMA and FDA principles become the de facto standard because that is what sponsors already file against. Both routes favour organisations that adopted the ten principles early, at no cost.

Three years and beyond. The likeliest development is a requirement to declare material AI involvement in regulated submissions and evidence generation. It is the smallest change that would let a regulator supervise what it currently cannot see. A second decision is already on the table and undated: whether nucleic acid synthesis screening moves from voluntary guidance to a statutory footing.

The first horizon is a reasonable reading of a published regulatory expectation. The second is a projection with two named routes. The third is speculation, offered as such and argued rather than asserted.

AI Adoption Plan: Life Sciences, 8 June 2026. MHRA Inspectorate, 29 June 2026. EMA and FDA, January 2026. Data (Use and Access) Act 2025, section 80.

ELEVEN

Benefit realisation

The business case for AI in this sector is a modelled one. The evidence that it is not being paid for out of the compliance budget does not yet exist anywhere.

A worked example

The government's *AI Adoption Plan: Life Sciences*, published on 8 June 2026 by Dave Hallett, the sector's AI Champion, reports early modelling suggesting that AI-driven research from discovery through to preclinical "could deliver time and cost savings of at least 25–50%", and sets a target of trial-ready drugs within 100 days by 2030. It cites an OECD estimate of total factor productivity gains exceeding 6% in pharmaceuticals over the next decade under high adoption.

Those are modelled figures and this briefing labels them as such throughout. Against them sits a 2026 review characterising AI in drug discovery as "in a transitional phase characterised by high investment but limited validated clinical impact", and noting that approximately 90% of candidates entering clinical development still fail — "a rate broadly comparable to traditional pharmaceutical attrition despite AI-assisted optimisation strategies". Acceleration before the clinic has not yet changed what happens inside it.

The plan is also candid about the barriers, and its candour is worth more than its targets: fragmented, "non-AI-ready" data requiring expensive wrangling; compute demand outstripping supply; "a missing middle of machine learning engineers and data stewards"; and public trust, which the plan addresses directly by naming the UK Biobank data leak and the controversy over Palantir's NHS contract rather than talking around them.

What to measure instead

Verification burden	Hours spent checking AI-assisted regulatory content, against hours saved drafting it. If the first exceeds the second, the saving is a transfer.
Citation integrity	The share of AI-assisted submissions in which a factual claim or citation failed review. This is the direct analogue of the failure the MHRA documented.
Cycle time to close	Time from inspection finding to accepted response, tracked before and after AI assistance was introduced. Four hours becoming twenty is this measure, from the regulator's side.
Declared scope	The share of regulated processes where AI involvement is documented rather than assumed absent. This is the measure that makes the other three possible.

WHY THE SECOND SET MATTERS MORE

The first set is a business case. The second is the evidence that the business case is not being paid for out of the compliance budget. The MHRA's documented example — four hours of regulatory review becoming more than twenty — is a benefit realisation failure quite as much as it is a governance one, and no organisation currently measures it.

AI Adoption Plan: Life Sciences, Dave Hallett, 8 June 2026. Khairil et al., *Pharmaceuticals* 19(6):916, 10 June 2026 — peer-reviewed, characterised here as a review rather than a primary dataset. MHRA Inspectorate, 29 June 2026.

TWELVE

A way in

Eight workstreams. The first three form a natural first step.

01	**Readiness assessment**	Where AI already sits across GMP, GCP and GPvP processes, sanctioned and unsanctioned, and what evidence exists for each	4–6 weeks
02	**Use case identification and triage**	Which uses are assistive, which are augmentive, and which sit inside a statutory safety obligation without anyone having said so	3–4 weeks
03	**Regulatory and legal position paper**	A defensible written position on verification, disclosure, Articles 22A–22D and data permissions, mapped to the EMA and FDA ten principles and ready for the ICO code	3–4 weeks
04	**Board and executive governance**	Who decides, who verifies, who signs, and what the board sees before an inspection rather than after one	4 weeks
05	**Threat and integrity response**	Citation controls on regulated submissions, and supplier assurance on biosecurity screening	4 weeks
06	**Benefit realisation framework**	The four measures in section eleven, instrumented from the start	3 weeks, parallel
07	**Build, buy or partner evaluation**	Supplier assurance that survives an inspection question about who verified the output	4–6 weeks
08	**Twelve-month mobilisation roadmap**	Sequenced, costed, and written to be handed to whoever runs it	2 weeks

THE FIRST STEP

The first three workstreams form a natural first step of six to eight weeks.

There is no commencement date for a medicines-specific AI framework, and this briefing will not invent one. But the ICO's code of practice is being written under a statutory duty already in force, with draft automated decision-making guidance ahead of it, and two further decisions are live and undated: whether nucleic acid screening becomes statutory, and whether the UK writes a medicines framework at all. The deadline that exists today is the next inspection — the one date nobody controls, and the reason six to eight weeks starting now is the right answer rather than a convenient one.

Contact: nc@neilcatton.com · neilcatton.com · writing.neilcatton.com

THIRTEEN

Sources

Every source read directly. Where a document could not be read in full, that is stated.

Regulators and government

- MHRA Inspectorate, "Use of AI for GXP inspection responses", 29 June 2026.
- MHRA, *Impact of AI on the regulation of medical products*, 30 April 2024, last updated 31 July 2026.
- Health Research Authority, "New UK clinical trials regulations signed into law", 10 April 2025.
- Health Research Authority, on the National Commission on the Regulation of AI in Healthcare, 26 September 2025.
- Dave Hallett, AI Champion for the Life Sciences sector, *AI Adoption Plan: Life Sciences*, gov.uk, 8 June 2026.
- Office for Life Sciences, *Bioscience and health technology sector statistics 2023 to 2024*, 2 October 2025.
- Office for National Statistics, *Artificial intelligence in UK businesses: 2023 to 2026*, 20 July 2026 (BICS Wave 159).
- National Commission into the Regulation of AI in Healthcare, gov.uk collection, last updated 19 June 2026.
- European Medicines Agency, *Reflection paper on the use of artificial intelligence in the lifecycle of medicines*, EMA/CHMP/CVMP/83833/2023, 30 September 2024.
- EMA and FDA, *Guiding principles of good AI practice in drug development*, January 2026.
- Department for Science, Innovation and Technology, *UK Screening Guidance on Synthetic Nucleic Acids for Users and Providers*, published 8 October 2024, last updated 2 September 2026.
- AI Security Institute (DSIT), *Frontier AI Trends Report*, 18 December 2025.
- Cabinet Office, *UK Biological Security Strategy Implementation Report, July 2025 – July 2026*, published 14 July 2026.
- Health and Safety Executive, *GMOs and the law*, and the GMO (Contained Use) Regulations 2014. Page updated 19 May 2026.
- UK Health Security Agency, Advisory Board paper on artificial intelligence, May 2024 (Steven Riley), and *Update on artificial intelligence in UKHSA*, March 2025 (Nick Watkins, Deputy Director, Data Science and Geospatial). *Pathogen Genomics Strategy*, 24 January 2024, last updated 7 August 2025. Its biological security outcomes are set out in the Cabinet Office report above.
- Information Commissioner's Office, *The Data Use and Access Act 2025 — what does it mean for organisations?*, page updated 19 June 2026, and *AI and biometrics strategy update*, March 2026.

Legislation

- The Medicines for Human Use (Clinical Trials) (Amendment) Regulations 2024, contents as published on legislation.gov.uk. Applied from 28 April 2026.
- Data (Use and Access) Act 2025, section 80, inserting UK GDPR Articles 22A–22D, in force 5 February 2026.
- The Data Protection Act 2018 (Code of Practice on Artificial Intelligence and Automated Decision-Making) Regulations 2026, SI 2026/425. Made 16 April 2026, in force 12 May 2026.

Industry, trade bodies and research

- BioIndustry Association, *UK biotech financing April – June 2026*, 20 July 2026.
- Genomics, *Genomics Launches the Mystra AI Platform, Adopted by Industry Giants and Up-and-Coming Biotechs*, 2 June 2026.
- Isomorphic Labs, *The Isomorphic Labs Drug Design Engine unlocks a new frontier beyond AlphaFold*, 10 February 2026.
- Isomorphic Labs, strategic multi-target research collaboration announcement with Eli Lilly, 7 January 2024.
- Isomorphic Labs, strategic multi-target research collaboration announcement with Novartis, 7 January 2024.
- Recursion Pharmaceuticals, merger completion announcement, 20 November 2024.

- Fierce Biotech, "AI drug hunter Exscientia chops down 'rapidly emerging pipeline' to focus on 2 main oncology programs", 13 October 2023 — reporting Exscientia's own stated reasoning for discontinuing EXS-21546; the company's own release was not read directly.
- BenevolentAI, press release on the FDA Emergency Use Authorisation for baricitinib, 20 November 2020, stating the hypothesis was first published 4 February 2020.
- Eli Lilly and Company, investor news release on the baricitinib Emergency Use Authorization, 20 November 2020.
- BenevolentAI, top-line Phase IIa results for BEN-2293 in mild-to-moderate atopic dermatitis, 5 April 2023.
- Our Future Health, cohort profile, *International Journal of Epidemiology*, 15 October 2025.
- Wittmann, Alexanian, Bartling, Beal, Clore, Diggans, Flyangolts, Gemler, Mitchell, Murphy, Wheeler and Horvitz, "Strengthening nucleic acid biosecurity screening against generative protein design tools", *Science* 390, 82–87 (2025). Cited qualitatively; paywalled, and its evasion rates are not reproduced here.
- Khairil, Koay, Jerry, Khatib, Che Ramli and Kumar, "AI in Drug Discovery: Clinical Failures, Regulatory Reality, and the Validation Crisis Behind the Hype", *Pharmaceuticals* 19(6):916, 10 June 2026. Characterised throughout as a review, not a primary dataset.
- Jayatunga, Ayers, Bruens, Jayanth and Meier, "How successful are AI-discovered drugs in clinical trials? A first analysis and emerging lessons", *Drug Discovery Today*, 2024, article 104009. Read in full. Excluded: all five authors are Boston Consulting Group employees and the paper states BCG's Health Care practice funded the research; see section one.

Two named pharmaceutical collaborations sit behind the Isomorphic Labs figure in section two, alongside a dated and company-stated reason for Exscientia's EXS-21546 discontinuation and a fourth company profile for BenevolentAI's baricitinib and BEN-2293 results. The Jayatunga et al. paper above was read in full and is excluded on sponsorship grounds, not access grounds.

FOURTEEN

About the author

Neil Catton is an independent technology strategist with more than thirty-nine years of technology leadership across twenty-plus sectors. He has held CTO and Director-level roles at BT, Wipro, Fujitsu, Sopra Steria and Capita, and works across a fractional portfolio spanning public safety, insurance, health technology and early-stage companies.

His career began in a safety-critical regulated environment — British Nuclear Fuels — and has returned to regulated delivery repeatedly since, across defence manufacturing, national health platforms and public safety. He is not a life scientist, and this briefing does not pretend otherwise. What it addresses is the question he has spent a career on: how an organisation proves, to a regulator, that a decision it now makes with a machine was actually reviewed by a person.

He is the author of four books, including the trilogy *The Next Evolution*, *The Cognitive Crucible* and *The Shadow System*, writing at the intersection of technology, ethics and human purpose. He writes at writing.neilcatton.com.

Contact: nc@neilcatton.com · neilcatton.com · writing.neilcatton.com

How this briefing was produced

Every figure is traced to a regulator, a trade body, published legislation, named academic research, or an operator speaking on its own record. Figures that could not be traced were excluded and named in section one rather than quietly dropped — including the sector's most-quoted statistic, whose source paper could not be obtained. Where a figure is attributed to a document that cites it rather than to the original source, that is stated. Where a claim rests on an absence, the absence is described at the level it was actually checked. Two corrections were made before publication. An earlier draft argued that no UK measurement of AI failure inside this sector existed; that was wrong, the MHRA's own inspectorate had published one, and the argument was rebuilt around it. A later draft described the biosecurity threat as a supply-chain question with no UK institutional anchor; that was also wrong, and sections four and six now set out the six bodies that hold this and the voluntary UK standard closest to it. A third correction followed: earlier drafts described section 80 of the Data (Use and Access) Act 2025 as having made automated decision-making permissible, without recording that the prohibition still applies where the decision rests on special category data, and omitted the statutory code of practice on AI that the Information Commissioner has been required to prepare since May 2026. Section seven now carries both.