RB Weekly AI Brief - Issue 30 - 30.10.2026

Covering the week of 30.10.2026 · Issue 30 of the RB Weekly AI Brief

Recurring themes: Regulatory & HTA Signals (3 of last 4 issues) · Regulation & Policy (3 of last 4 issues) · Healthcare & Life Sciences (3 of last 4 issues) · Models & Research (3 of last 4 issues)

AI News Roundup

Regulation & Policy

UK Budget Commits Over £700m to AI and Commissions NHS Productivity Review

The Budget delivered on 28 October states that the government has invested over £700 million in AI. Measures include an additional £120 million for the AI Security Institute, a £165 million National AI Biosecurity Initiative, an additional £140 million for the Sovereign AI Unit and £80 million for its challenge-led R&D procurement scheme, which names health and care among its priority areas. A new Chancellor's AI Challenge Prize, backed by £15 million, will hold its first competition on NHS challenges, with a top award of up to £2.5 million. The Budget also says Mark Bailie will lead an independent review of NHS productivity and AI, reporting in early 2027 to inform the next spending review.

So what? *The money for security and sovereign capability is announced, but the Bailie review is the item to watch: it feeds the next spending review and will examine where AI can lift NHS productivity. Pharma, market access and HEOR teams supplying AI-enabled tools or evidence to the NHS should track the review, as it may shape which AI applications the NHS prioritises and what productivity evidence it expects.*

HM Treasury

EU Council Backs Digital Omnibus Position on AI and Data Protection

On 21 October 2026 the Council agreed its negotiating mandate on the data-protection part of the Digital Omnibus, which amends the GDPR, updates cookie rules and sets up a framework for incident reporting. For AI, the mandate adds safeguards to the proposed derogation that allows sensitive data to be processed in the development and operation of AI systems, and clarifies the lawful ground for processing personal data. The Council also supports a principle-based definition of scientific research, with targeted exemptions for research organisations. Negotiations with the European Parliament will start once Parliament adopts its own position. The AI-focused part of the Omnibus, Regulation (EU) 2026/1744, was adopted on 29 June 2026 and is already in force; this data-protection file is the outstanding second track.

So what? *Pharma and HEOR teams that train or fine-tune AI models on patient-level data, including real-world evidence used in HTA submissions, now have a clearer sense of the direction on sensitive data and scientific research, but this is a negotiating position, not law, so wait for Parliament's position before revising data-processing agreements.*

Council of the European Union

MEPs Split on How to Govern AI as Parliament Debates Global Rules

Tech Policy Press reports that MEPs in Strasbourg debated global AI governance, with broad support for international cooperation, including with the US and China, and several lawmakers comparing AI risks to nuclear weapons. The Greens and Socialists rejected leaving AI rules to companies. Sergey Lagodinsky proposed a binding agreement built on the Hiroshima AI Process and Axel Voss an IAEA-style inspection regime, while the centre-right MEP Arba Kokalari argued that Europe must also lead in developing AI. On liability, the Commission has withdrawn the AI Liability Directive; Commissioner Jozef Síkela pointed to the AI Act and the revised Product Liability Directive as partial cover, while critics say victims still struggle to prove that an AI system caused harm.

So what? *With the AI Liability Directive withdrawn, liability for AI-related harm rests on the AI Act and the revised Product Liability Directive, which critics say leave proof of causation unresolved. Market access and HEOR teams should factor that uncertainty into risk frameworks for AI-generated evidence submitted to regulators and HTA bodies.*

Tech Policy Press

Regulatory & HTA Signals

NICE Draft Guidance Backs Five AI Prostate MRI Tools During Evidence Generation

On 23 October NICE published draft guidance recommending five AI technologies for use in the NHS, as options to help radiologists detect prostate cancer on MRI, during a two-year evidence-generation period: AI-Rad Companion Prostate MR, DeepHealth Prostate, ProstatID, Prostate Intelligence (Pi) and QP-Prostate. The independent committee found potential to improve the accuracy and consistency of interpretation, particularly in settings without specialist expertise, but said the evidence is limited and not always consistent across studies. Further NHS evidence is needed on diagnostic accuracy, patient outcomes, workflows, costs and performance across different groups. Two further technologies, hProstate and mdprostate, were not recommended for NHS funding because no clinical or cost-effectiveness evidence was available for them. The draft guidance and evidence generation plan are open for consultation.

So what? *NICE is pairing NHS access with a defined evidence-generation plan, and withholding it where no evidence exists. Market access teams for AI diagnostics should plan to collect real-world accuracy, outcome, workflow and cost data from the start, as that is what a final recommendation is likely to depend on.*

NICE

Healthcare & Life Sciences

AstraZeneca and Tremont AI Collaborate on Preclinical Safety AI

Boston-based Tremont AI, spun out of the Mahmood Lab at Mass General Brigham, has announced a research collaboration with AstraZeneca to develop and evaluate AI solutions for faster safety screening in preclinical toxicologic pathology studies. The work targets the review of tissue images, where pathologists must separate treatment-related changes from normal variation, and will test and improve models on rare findings that are hard to distinguish from background. The release describes research and evaluation only; it does not disclose financial terms, duration, results or a timeline.

So what? *With no results or timeline disclosed, this is an early signal rather than evidence. If AI-assisted toxicologic pathology is validated, it could change how safety signals are read before first-in-human studies, which would matter for the evidence that regulators and HTA bodies later see.*

Business Wire

Insilico Hits First Milestone in Servier AI Oncology Deal

On 22 October Insilico Medicine announced that it has reached the first early drug discovery milestone in its AI-driven oncology collaboration with Servier, triggering a milestone payment of undisclosed size. Insilico says it completed target selection for both programmes and advanced the first through predefined early discovery criteria within 10 months of launching the collaboration. Under the agreement signed in January 2026, worth up to US$888 million in total, Servier shares R&D costs and leads clinical validation and commercialisation. The milestone claims are the company's own.

So what? *This is a process signal, not clinical evidence: payments are being triggered by predefined early discovery criteria, with clinical validation still ahead. HEOR and HTA teams should treat the value narrative for AI-discovered oncology assets as unproven until clinical data exist.*

PR Newswire

Models & Research

OpenAI Releases GPT-6.1 Terra for Faster, Cheaper Agentic Work

On 26 October OpenAI released GPT-6.1 Terra, which it describes as a fast, efficient model combining detailed visual reasoning and precise computer use. It is available to Plus, Pro, Enterprise and Edu users in ChatGPT Work and Codex, and through the API at $1 per million input tokens and $5 per million output tokens, which OpenAI says is 50 per cent lower than GPT-5.6 Terra and GPT-6.1 Sol. OpenAI reports that Terra completes tasks more than three times faster than GPT-6 Astra and GPT-6.1 Sol, and says Sol and Astra remain the better choice for the most complex coding tasks. The speed, cost and benchmark comparisons are OpenAI's own.

So what? *Lower prices and faster completion reduce the cost of running AI agents on high-volume HEOR tasks such as literature screening, data extraction and dossier drafting, but the figures are vendor-reported, so test it on your own tasks and accuracy requirements before changing any workflow.*

OpenAI

Academic Paper Summaries

Selected from PubMed · Published within the last 12 months · New selections each week

Domain Paper — HEOR / Health Economics / Market Access

Knowledge-Practice Performance Gap in Clinical Large Language Models: Systematic Review of 39 Benchmarks.

Gong EJ, Bang CS, Lee JJ, et al. · Journal of medical Internet research · 2025

#ClinicalAI · #Regulation

This systematic review of 39 benchmarks used to test AI language models in medicine found a large gap between knowledge and practice. Leading models reach 84 to 90 per cent accuracy on exam-style questions, but only 45 to 69 per cent on benchmarks that test practical clinical tasks, and 40 to 50 per cent on safety assessment. About one in four benchmarks (10 of 39) did not report its methods well enough to be fully evaluated. The authors conclude that exam scores are insufficient proxies for clinical readiness and that deployment needs practice-based validation with human oversight. For healthcare and HEOR teams, it is a reminder to ask how a tool was tested, not just what it scored.

PMID: 41325597	PubMed →	DOI →

AI Research Paper 1

A clinical environment simulator for dynamic AI evaluation.

Luo L, Kim SE, Zhang X, et al. · Nature medicine · 2026

#ClinicalAI · #PatientOutcomes

The authors propose, but do not yet report building and testing, a Clinical Environment Simulator: a digital hospital in which an AI model's decisions change what happens next. A 'hospital engine' would track beds, staff workloads and equipment, while a 'patient engine' would simulate disease progression and treatment response. The aim is to evaluate AI on three things that static exam-style tests miss: reasoning as conditions change, decisions that use up capacity needed by other patients, and resilience when emergencies and system failures coincide. For healthcare executives, it is a proposal for how AI evaluation could better reflect real deployment conditions, not yet evidence about any tool.

PMID: 41820673	PubMed →	DOI →

AI Research Paper 2

Genomics of drug target prioritization for complex diseases.

Chen R, Duffy Á, Do R · Nature reviews. Genetics · 2026

#Genomics · #DrugDevelopment · #RealWorldEvidence

This review describes how large genetic databases, multi-omics data and machine learning are being combined to help choose which biological targets drugs should aim at, particularly for complex diseases. Bringing several lines of human genetic evidence together in a probabilistic framework can support more systematic target prioritisation, earlier prediction of adverse effects and the identification of existing drugs that might be repurposed. The authors position this as a way to address high late-stage attrition. For pharmaceutical teams, it is an overview of an approach that aims to lower late-stage risk, rather than a demonstration that it does.

PMID: 41198831	PubMed →	DOI →

This brief is produced using an AI-assisted workflow with expert human editorial review. Stories are curated and sources verified before publication. Readers are encouraged to consult source links directly.