

China's Dual-Track AI Strategy: Strengthening Regulation While Accelerating Life Sciences Innovation

September 1, 2026 | Opinion | By Anastacia Edwards-Kurianova, Life Sciences Executive and AI-native Transformation Strategist

China is taking a distinctive approach to AI governance, combining regulatory oversight with aggressive industrial adoption. For the life sciences sector, this dual-track strategy is particularly significant, as AI moves into drug discovery, healthcare, manufacturing and regulatory decision-making. The emerging framework offers lessons in scaling innovation while embedding accountability, traceability and risk management.



The global AI debate is often framed as a choice between acceleration and restraint: innovate first and regulate later, or regulate first and accept slower progress. China is pursuing a third path. It is regulating in order to scale.

The result is neither a laissez-faire market nor a single regulatory wall. It is a layered system in which industrial policy pushes AI into science, healthcare and manufacturing, while cyber, data and product rules define the corridor within which deployment can occur. For life sciences, where a model output may influence a clinical, quality or regulatory decision, this architecture matters more than any algorithmic leaderboard.

The duality became explicit in China's amended Cybersecurity Law, effective from January 2026. Its new Article 20 simultaneously supports AI research, algorithms, training data and computing infrastructure, while calling for stronger ethics, risk monitoring, assessment and security supervision. Development and control are not presented as opposing forces. They are components of the same national capability.

A regulatory stack, not a regulatory wall

China's framework is best understood by following the use case rather than searching for one all-purpose AI law.

The 2023 Interim Measures for the Management of Generative AI Services govern services offered to the public in China. They require lawful training data, protection of personal information and intellectual property, measures against discrimination, and improvements in accuracy and reliability. Crucially, research and internal organisational applications that are not offered to the domestic public fall outside these measures. That is not an exemption from governance. It is a form of regulatory calibration: public-facing generative services receive one layer of control, while internal life sciences applications remain subject to data, cybersecurity, intellectual-property, ethics and sector-specific obligations.

The 2025 Measures for Labelling AI-Generated and Synthetic Content add another layer. They require visible and embedded forms of identification across relevant generated content and place duties on both generation and distribution services. The immediate target is online information integrity, but the underlying principle is broader: synthetic outputs should carry provenance. In life sciences, provenance is not merely a disclosure issue. It is the beginning of auditability.

Data regulation supplies the foundation. The Cybersecurity Law, Data Security Law and Personal Information Protection Law govern how data are collected, processed and protected. China's 2024 cross-border data-flow provisions simplified some lower-risk transfers, but retained security assessment and other mechanisms for important data and threshold-based personal-information exports. For multinational life sciences companies, this means that global model development cannot be separated from data mapping, lawful access, transfer architecture and local accountability.

Acceleration is equally deliberate

The other track is not rhetorical encouragement; it is industrial execution with dates, infrastructure and measurable targets.

The State Council's 2025 "AI Plus" Opinion calls for AI to be deeply integrated across six priority domains. It targets adoption of next-generation intelligent terminals and agents above 70 per cent by 2027 and above 90 per cent by 2030. More important for life sciences, it promotes AI-driven scientific discovery, scientific foundation models, high-quality datasets, biotechnology integration, health assistants, auxiliary diagnosis and health management.

The pharmaceutical sector has its own implementation machinery. The Digital and Intelligent Transformation Implementation Plan for the Pharmaceutical Industry (2025-2030) links AI to the entire value chain and to full-lifecycle quality management. By 2027, it envisages more than 30 digital standards, 100 high-performance products, 100 representative application scenarios, 100 digital pharmaceutical and medical-device factories, and more than 10 pharmaceutical large-model innovation, validation and pilot platforms. The direction is unmistakable: AI is expected to move from isolated experiments into industrial infrastructure.

Healthcare policy follows the same logic. The 2025 implementation opinion on promoting and regulating "AI Plus Healthcare" applications identifies eight directions and 24 priority applications, spanning primary care, clinical diagnosis, patient services, public health, research, industry governance and the health economy. Alongside high-quality datasets, vertical models and application test bases, it calls for classified supervision, data security and privacy protection. Meanwhile, NMPA guidance for AI medical devices already requires evidence around dataset construction, annotation quality, diversity, algorithm performance, influencing factors, clinical evaluation and lifecycle control.

China is therefore not opening one undifferentiated fast lane. It is building multiple regulated lanes, each tied to the consequence of the use case.

From model performance to decision integrity

This distinction is critical because life sciences convert computational outputs into consequential decisions. A hallucination in a consumer chatbot may misinform. A hallucination embedded in target identification, protocol design, patient selection, batch disposition or clinical decision support may alter an evidence chain, a product-quality judgement or a patient's care.

Traditional AI governance often asks whether a model is accurate, secure and unbiased. Regulated life sciences must ask a harder question: can the organisation reconstruct and defend the decision made with it?

That requires decision integrity. The intended use must be bounded. Data lineage and dataset suitability must be demonstrable. Inputs, prompts, outputs and human interventions must be traceable where risk warrants it. Model changes must enter change control. Performance drift must be monitored in the population and environment of use. Most importantly, human oversight must identify who has the authority, competence and evidence to challenge the machine.

An AI system may recommend; accountability does not migrate to the algorithm.

This is where China's dual-track strategy offers a deeper lesson. Regulation can reduce innovation velocity when it is added after deployment as documentation and remediation. But when provenance, validation and oversight are designed into shared platforms, governance becomes reusable infrastructure. The first compliant deployment may require more discipline; the tenth can move faster because the evidence architecture already exists.

The executive challenge

For global life sciences leaders, three shifts follow.

First, map decisions, not merely models. The same foundation model can summarise literature, draft a protocol and support a patient-facing service; each use carries a different consequence and regulatory perimeter. Governance should follow the decision and its downstream impact, not the vendor label.

Second, combine technical assurance with GxP assurance. A model card is not a validation package, and a conventional validation protocol may not address bias, data drift or emergent behaviour. Organisations need one control framework that connects data science, quality, clinical, regulatory, privacy and cybersecurity evidence.

Third, localise execution without fragmenting accountability. Chinese data, hosting and product requirements may require local architecture, but the enterprise still needs globally coherent principles for risk classification, human authority, audit trails and post-deployment monitoring.

The strategic significance of China's approach is not that it has solved every tension between control and innovation. No jurisdiction has. It is that China is treating trust as productive capacity. Regulation makes deployment more legible; industrial policy makes deployment increasingly unavoidable.

The organisations that lead will not be those that run the most pilots, nor those that avoid risk by waiting. They will be those that turn governance into a repeatable operating system for innovation. In life sciences, speed is not the absence of control. It is the absence of preventable rework because control was built in from the beginning.

Anastacia Edwards-Kurianova, Life Sciences Executive and AI-native Transformation Strategist