The Artificial Intelligence Bill 2026: A Legal Analysis

Lyla Latif
latif@uonbi.ac.ke

Table of Contents

Cite as: Latif, L., 'The Artificial Intelligence Bill 2027: A Legal Analysis' *Briefing Paper No. 10, 2026*. Committee on Fiscal Studies. University of Nairobi.

BACKGROUND

Would I be right to say that Senator Nyamu has introduced an AI bill that knows what AI can do to a person but not what it does to a country? One important thing about the Artificial Intelligence Bill 2026 that was gazetted on 19 February 2026 is that it arrives in the context of a growing public and legal conversation about AI governance in Kenya, which has been building without legislative response for several years. When we look around us AI systems are already making consequential decisions about us. What I can say with confidence is that Kenya's mobile lending sector and a range of other digital credit platforms, operates on algorithmic credit scoring systems that determine in real time whether a Kenyan citizen accesses credit, at what cost, and under what terms. These systems process behavioural data, transaction history, social network data, and device metadata to generate creditworthiness scores. Accordingly, they are, by any technical definition, AI systems performing high-stakes decisions. The Central Bank of Kenya's own Digital Credit Providers Regulations 2022 acknowledge the existence of algorithmic lending but do not regulate the decision-making logic that is embedded in those algorithms. AI, therefore, is making consequential decisions affecting Kenyan citizens. The bill's arrival is therefore necessary.

1. Substantive Features of the Bill and Why They Matter

- Part II of the bill establishes the Office of the Artificial Intelligence Commissioner as a State Office under Article 260(q) of the Constitution, with corporate personality and independence in the performance of its functions under Section 4(4). This is the bill's most important structural contribution. An independent constitutional office, as distinct from a ministerial directorate, has genuine insulation from political interference.

- The appointment process under Section 5, which runs through the Public Service Commission and requires parliamentary approval of the President's nominee, is more transparent than many comparable regulatory appointments in Kenya's recent legislative history.

- The qualification criteria under Section 6 are appropriately broad, requiring a master's degree in AI, computer science, information technology, engineering, data science, law, ethics, or a related field, combined with ten years of relevant experience including in AI governance, data protection, technology policy, human rights, and risk management. The inclusion of law and ethics alongside technical disciplines signals that the office is conceived as a governance institution rather than a technical inspectorate.

- Part V contains the substantive governance provisions. Section 25 establishes a four-tier risk classification system distinguishing unacceptable risk, high risk, limited risk, and minimal risk AI systems. The classification of high-risk systems under Section 25(2)(b) encompasses healthcare, education, agriculture, finance, security, employment, and public administration. These are the right picks for Kenya as each of them involves AI deployment that has direct consequences for whether a Kenyan citizen accesses credit, receives a job offer, obtains healthcare, or is flagged by a law enforcement algorithm.

- The obligations attached to high-risk systems under Section 26 are substantive. These obligations include pre-deployment risk assessments, human oversight, human rights impact assessments under the Data Protection Act, five-year retention of training data records and performance metrics, transparency and explainability requirements, and annual compliance reporting to the Commissioner. Section 33 requires providers and deployers of AI systems likely to affect employment to conduct workforce impact assessments and implement mitigation measures including reskilling programmes. This provision reflects a political economy awareness appropriate to a labour market where AI-driven displacement of routine work is an immediate threat rather than a speculative one.

2. The Problem of Overlapping Mandates and Existing Law

Before assessing what the bill misses, it is necessary to work through what Kenyan law already covers in this space, because the bill creates a new regulatory layer on top of existing instruments without adequately addressing how that layer relates to what is already there.

- First, is the Data Protection Act 2019. It governs the collection, processing, storage, and transfer of personal data, establishes the Office of the Data Protection Commissioner, and creates rights including the right to access, rectification, erasure, and objection. The AI Bill at Section 26(1)(e) simply requires compliance with the Data Protection Act in relation to personal data processing. This is legally insufficient. The Data Protection Act was designed for conventional data processing relationships. Its principle of purpose limitation requires that data be collected for a specified, explicit, and legitimate purpose. Machine learning systems are typically trained on data collected for entirely different purposes, and the inferences they generate go far beyond what the original data collection contemplated. The Data Protection Act's consent framework does not address this adequately, and the AI Bill does not amend or supplement the Data Protection Act to close the gap. Section 28 of the AI Bill grants users the right to explanation and the right to human review of automated decisions that produce significant legal effects, which is welcome. However, the right to explanation under Section 26(1)(c) applies only to providers and deployers of high-risk systems. A credit scoring system that denies a Kenyan citizen a mobile loan may or may not meet the classification threshold for high risk depending on how the Commissioner eventually defines that category in regulations under Section 36(2)(a). Until those regulations exist, the right to explanation is contingent on a classification decision that has not yet been made.

- Second, is the Computer Misuse and Cybercrimes Act 2018. This Act governs unauthorised access to computer systems, data interference, and cybercrimes. AI systems introduce entirely new attack surfaces that the Cybercrimes Act was not designed to address. Adversarial inputs can manipulate an AI system's outputs without constituting unauthorised access in the traditional sense. Model inversion attacks can extract training data from a deployed model without accessing the underlying system directly. Prompt injection attacks against large language models can cause those systems to act in ways their operators did not authorise. The AI Bill mentions robustness and cybersecurity in Section 26(1)(f) but creates no specific framework for these AI-specific threats and does not connect to the Cybercrimes Act in any way that would address them.

- Third, is the Kenya Information and Communications Act, Cap 411A. This Act vests regulatory authority over communications infrastructure in the Communications Authority of Kenya (CAK). AI systems deployed in telecommunications networks, content recommendation systems, and digital infrastructure fall within CAK's sectoral mandate. Section 14 of the AI Bill allows the Commissioner to delegate powers to any regulator established by an Act of Parliament, which creates a theoretical pathway for CAK to exercise AI governance functions in the communications sector. But theoretical is the operative word. The bill establishes no coordination mechanism, no protocol for shared jurisdiction, and no procedure for resolving conflicts between the AI Commissioner and CAK. The same problem applies to the Central Bank of Kenya's mandate under the Central Bank of Kenya Act, the Insurance Regulatory Authority's mandate under the Insurance Act, and the Capital Markets Authority's mandate under the Capital Markets Act Cap 485A. All of these sector regulators already govern AI deployment in their respective sectors to varying degrees of adequacy, and all of them will encounter the AI Bill's Commissioner as either a partner or a competitor in exercising that authority. The bill does not decide which.

- Fourth, is the Science, Technology and Innovation Act 2013. This Act establishes the National Commission for Science, Technology and Innovation (NACOSTI) and vests in it responsibility

for promoting science, technology, and innovation in Kenya. NACOSTI appears in the AI Bill only as a member of the Advisory Committee under Section 17(1)(d). The relationship between the AI Commissioner's mandate to develop AI governance policy and NACOSTI's existing mandate over technology development and standardisation is never resolved. This is the kind of legislative overlap that generates institutional turf disputes that ultimately harm the regulatory coherence the bill is trying to create.

3. The Comparative Framework: What South Korea and the EU Have Built

The AI Bill is not the first of its kind. The EU and South Korea already have their AI Acts.

The EU Artificial Intelligence Act, which entered into force in August 2024, shares the Kenya bill's risk-based architecture and its concern with high-risk AI in critical sectors. What it adds that the Kenya bill lacks:

- Is a governance framework for general-purpose AI models and foundation models. Under the EU Act, providers of general-purpose AI models are required to prepare and maintain technical documentation, publish summaries of training data in compliance with EU copyright law, implement policies to respect intellectual property, and report serious incidents to the European AI Office.
- Providers of general-purpose AI models posing systemic risk, defined by reference to training computation thresholds, are subject to additional requirements including adversarial testing, cybersecurity obligations, and energy efficiency reporting. This matters for Kenya because every significant AI application deployed in Kenya that is built on a foreign foundation model is built on infrastructure that the EU now regulates at source.
- Kenya's bill does not regulate that infrastructure at all. A Kenyan bank deploying a credit scoring system built on a foreign large language model is subject to Section 26's obligations as a deployer of a high-risk system, but the foundation model provider faces no obligation under Kenyan law whatsoever.

South Korea's Framework Act on the Development of Artificial Intelligence and Establishment of Trust 2025, which came into force in January 2026, takes a different but instructive approach.

- Rather than building a separate AI regulatory institution, South Korea positions AI governance as a whole-of-government responsibility coordinated through the Prime Minister's office, with the Ministry of Science and ICT as the lead ministry. More importantly for Kenya's purposes, the South Korean framework includes specific provisions on high-impact AI services, defined as AI services used in areas with significant effects on human life, physical safety, or fundamental rights, and places disclosure and accountability obligations on operators of such services including foreign operators whose services are accessed by Korean users. The principle that foreign providers of AI infrastructure bear accountability obligations towards the jurisdictions in which their systems produce effects is fundamental to making AI governance meaningful in countries that consume rather than produce foundational AI technology. Kenya's bill does not engage with this principle.

- The South Korean framework also includes specific provisions on data quality for AI training, requiring that high-impact AI systems be trained on data that is representative, accurate, and free from discriminatory bias to the extent technically feasible. This addresses the root cause of most AI discrimination: biased outputs derive from biased training data, and a governance framework that only scrutinises the deployment layer while leaving the training data layer unexamined is treating the symptom rather than the cause. The Kenya bill at Section 26(1)(d) requires providers to maintain records of training datasets. It does not require that those datasets meet any quality, diversity, or representativeness standard before a high-risk system is deployed.

4. What the Bill Must Address to Be Fit for Purpose

- *A civil liability framework*

Sections 35 imposes criminal penalties for specified offences including deploying unacceptable-risk systems, deploying high-risk systems without conducting required assessments, failing to comply with transparency obligations, and generating harmful synthetic media without consent. Criminal liability is an important but entirely insufficient accountability mechanism for AI-related harm. The overwhelming majority of AI-related injuries, including discriminatory credit decisions, erroneous health risk classifications, biased employment screening, and inaccurate criminal justice risk scores, do not constitute criminal offences. They are civil wrongs. A Kenyan citizen who is denied a bank loan because an AI system scored them as high credit risk on the basis of their postal address, which correlates with race and income in Nairobi's geography, has a procedural right to explanation under Section 28 but no statutory cause of action in damages. The EU AI Act connects AI obligations to civil liability through the AI Liability Directive and revised product liability rules. The Kenya bill creates obligations without creating remedies, which means that enforcement depends entirely on the Commissioner's capacity and willingness to act, and injured persons have no independent legal avenue.

- *A provision on governing foundation models and large language models (LLMs)*

The absence of any provision governing foundation models and large language models is the single most significant structural gap in the bill. The foundational AI infrastructure on which Kenya's digital economy increasingly depends is not built in Kenya. It is built and controlled by technology corporations domiciled in the United States. The bill governs Kenyan deployers of AI systems. It does not govern the providers of the foundational systems on which those deployers build. This means that the most consequential AI governance questions, including the quality and representativeness of training data, the alignment of model values with Kenyan constitutional values, the energy costs of model inference, and the profit flows generated by AI deployment in Kenya, are all outside the bill's scope. An AI governance framework that cannot reach the infrastructure layer is a framework that will always be governing at one remove from where the real decisions are made.

- *A provision on enforcing environmental sustainability*

The environmental dimension of AI governance is stated as an ethical guideline under Section 30(2)(d) rather than as a mandatory legal obligation. The AI Bill requires the Commissioner to include environmental sustainability, specifically assessments of energy consumption and carbon footprint, in the ethical guidelines it publishes for AI system development and deployment. This formulation reduces what should be a binding disclosure and compliance obligation to advisory guidance. Data centres in Kenya, which provide the computational infrastructure for AI training and inference, consume electricity and water at scales that have direct consequences for Kenya's energy security and water resources. The bill mentions environmental sustainability but creates no mechanism for enforcing it, no connection to Kenya's environmental law framework under the Environmental Management and Coordination Act 1999, and no relationship to the fiscal governance of data centre investment incentives. This is a gap that needs to be closed through primary legislation rather than guidelines.

- *A provision requiring public entities to publish registers of AI systems in use*

The public sector AI provisions under Section 34 simply require public entities to ensure compliance with the Act. This is a compliance statement, not a governance framework. It does not require government agencies to publish registers of AI systems in use, does not mandate impact assessments before AI is procured or deployed in public services, does not address the accountability gap created when algorithmic systems make or recommend decisions about access to public resources, land registration, taxation, or welfare. A mandatory government AI register, modelled on the approach taken

by Canada's Directive on Automated Decision-Making, would be a minimum requirement for meaningful public sector AI governance.

- *Include a representative from the financial sector*

The bill's Advisory Committee under Section 17 does not include a representative of any financial sector regulator. Given that Kenya's financial sector is the most AI-intensive sector in the economy and that algorithmic credit scoring, fraud detection, and insurance underwriting affect millions of Kenyan citizens, the absence of the Central Bank of Kenya, the Insurance Regulatory Authority, and the Capital Markets Authority from the Advisory Committee's composition is a design error that should be corrected.

- *Funding*

The funding architecture under Section 21 is entirely inadequate for the mandate the bill creates. The Office is expected to conduct risk assessments of AI systems across the entire economy, maintain a public register of high-risk systems, establish and manage regulatory sandboxes, perform post-market surveillance including conformity audits under Section 10(1)(c), implement national AI literacy programmes, engage internationally, and receive and investigate complaints about AI harms including bias and discrimination. These are not light-touch regulatory functions. The Central Bank of Kenya, which performs comparable technical oversight in its sector, operates with an annual budget measured in billions of shillings. A funding mechanism that relies on National Assembly allocations and voluntary donations will not sustain the technical capacity this mandate requires.

5. The Domestic Product Problem: Why AI Cannot Be Ringfenced

The bill governs AI as if it were a domestic product. This is problematic for the following reasons:

- The bill's operative provisions, most significantly the obligations under Section 26 on providers and deployers of high-risk AI systems, are addressed to natural or legal persons who develop or deploy AI systems in Kenya. The definitions in Section 2 define a provider as a person who develops an AI system and places it on the market or puts it into service under their own name or trademark, and a deployer as a person who puts an AI system into service or uses it under their authority. Read carefully, these definitions do assume a degree of domestic presence or domestic legal personality. So, the enforcement powers under Section 11, including entering premises, requiring production of records, and issuing enforcement notices, are all premised on a physical or legal subject within Kenya's jurisdiction against whom those powers can be exercised.

- The practical reality of how AI systems operate makes that assumption increasingly fictitious. When a Kenyan hospital deploys a diagnostic AI tool, the model may have been trained on servers in the United States, fine-tuned on servers in Ireland, deployed via an API hosted by a cloud provider whose legal entity is incorporated in the Netherlands, accessed by a Kenyan technology company that is the deployer for the purposes of the bill, and used on a device held by a patient in Nairobi. The Kenyan deployer is subject to Section 26. The entity that built the model, made the foundational architectural choices that determine what the model can and cannot detect, and controls the training data that determines the model's performance characteristics across demographic groups, is not subject to any obligation under this bill whatsoever. When the model produces a diagnostic error that injures a patient, the Kenyan deployer is liable under the bill's framework. The model provider, whose technical decisions created the conditions for that error, is not.

- This is not an abstract jurisdictional point. It is the structure through which AI-related harm will predominantly manifest in Kenya over the next decade. The foundational AI systems on which

Kenyan applications are built are not Kenyan products. They are American and, to a lesser extent, Chinese products. The bill cannot exercise territorial jurisdiction over their providers in the traditional sense, and it makes no attempt to do so.

Therefore, the question is whether there are legal mechanisms through which Kenya can extend governance obligations to foreign AI providers whose systems produce effects within Kenya, without claiming territorial sovereignty over activities conducted entirely abroad. There are.

- There is the EU AI Act's approach. The EU can require foreign providers to comply with its framework as a condition of market access because the EU market is large enough to make exclusion from it commercially unacceptable. Kenya does not have that leverage individually. However, the principle of effects-based jurisdiction, which is well established in competition law, data protection law, and international tax law, provides a legal basis for asserting that a foreign provider whose AI system produces effects in Kenya, including by generating revenue from Kenyan users or causing harm to Kenyan persons, is subject to Kenyan regulatory obligations in respect of those effects. This is not an unprecedented claim. Kenya's own Digital Service Tax, that was introduced through the Finance Act 2020, asserted exactly this principle in the fiscal context: a foreign digital services provider with no physical presence in Kenya is subject to Kenyan tax on revenue derived from Kenyan users. The AI Bill could apply the same logic to governance obligations, requiring foreign providers of AI systems used in high-risk contexts by Kenyan deployers to register with the Commissioner, provide technical documentation on their systems, and submit to conformity assessments in respect of their Kenya-facing deployments.

CONCLUSION AND SUGGESTIONS

The Artificial Intelligence Bill 2026 is a serious legislative document. It establishes an independent governance institution with a meaningful mandate, it creates accountability obligations for the highest-risk AI deployments, it contains a workforce impact assessment requirement that reflects real political economy thinking, and it establishes a regulatory sandbox architecture that is consistent with modern governance practice. These contributions deserve acknowledgement. However, a bill that governs AI deployment without governing AI infrastructure, that creates obligations without creating remedies, that acknowledges environmental sustainability as a guideline while data centres consume megawatts without disclosure, that leaves financial sector AI entirely to existing sector regulators without coordination architecture, and that has no framework for the foundation models on which Kenya's AI economy depends, is a bill that addresses the present without preparing Kenya for what AI governance will need to do over the next decade.

I would propose the following insertions into the bill to inform good debate and enactment into law:

a) A new Part on foundation model governance imposing disclosure and accountability obligations on foreign providers of general-purpose AI models whose systems produce effects for Kenyan users;

b) A civil liability provision creating a statutory cause of action for persons harmed by AI systems that fail to meet the bill's obligations, with a reverse burden of proof on the provider or deployer in high-risk contexts;

c) A mandatory government AI register requiring all public entities to publish inventories of AI systems in operational use and impact assessments conducted before deployment;

d) A statutory coordination mechanism between the AI Commissioner and all sector regulators with AI jurisdiction, replacing the discretionary delegation under Section 14 with a mandatory inter-agency protocol;

e) Mandatory energy and water disclosure requirements for AI infrastructure operators as a condition of the AI Commissioner's approval of high-risk AI system registration, connecting the AI governance framework to the environmental accountability deficit in digital infrastructure governance;

f) A minimum funding guarantee expressed as a percentage of the national ICT sector revenue to ensure the Office is resourced in proportion to its mandate; and

g) An amendment to the Advisory Committee composition under Section 17 to include representatives of the Central Bank of Kenya, the Insurance Regulatory Authority, and the Capital Markets Authority, whose sectors account for the largest volume of high-risk AI deployment in Kenya.

These proposals are not exhaustive. They are the minimum additions without which the bill, as enacted, will not achieve the objectives its architect has stated. The Senate's deliberations on this bill are an opportunity to build a governance framework that is equal to what AI is actually doing to Kenya's economy and to Kenyan citizens. That opportunity should not be lost to the usual pressures of legislative haste.

About the Committee on Fiscal Studies

We are dedicated to the reform and strengthening of tax systems in the global south generating research needed for knowledge transfer, policy advocacy, capacity building and training. We are based at the University of Nairobi, Faculty of Law (Parklands Campus).