

**INTERNATIONAL JOURNAL OF LAW
MANAGEMENT & HUMANITIES**
[ISSN 2581-5369]

Volume 9 | Issue 5

2026

© 2026 *International Journal of Law Management & Humanities*

Follow this and additional works at: <https://www.ijlmh.com/>

Under the aegis of VidhiAagaz – Inking Your Brain (<https://www.vidhiaagaz.com/>)

This article is brought to you for free and open access by the International Journal of Law Management & Humanities at VidhiAagaz. It has been accepted for inclusion in the International Journal of Law Management & Humanities after due review.

In case of **any suggestions or complaints**, kindly contact support@vidhiaagaz.com.

To submit your Manuscript for Publication in the **International Journal of Law Management & Humanities**, kindly email your Manuscript to submission@ijlmh.com.

Review Capacity and Contestability: Procedural Justice Beyond the EU AI Act in Technologically Dependent States

BUDWIN ISHARA SIRIWARDENA* 

ABSTRACT

This article examines whether the procedural safeguards of the European Union Artificial Intelligence Act can operate effectively in technologically dependent states. Using doctrinal analysis of the Act and comparative public-law materials, it argues that transparency, documentation, logging, human oversight, impact assessment and monitoring presuppose a prior institutional condition: review capacity. Where public authorities depend on external vendors to design, operate, maintain or audit AI systems, legal responsibility may remain with the state while practical control lies elsewhere. The resulting accountability gap is examined through the concept of review dependency and through litigation concerning automated welfare administration and a hypothetical digital-identity example. The article argues that inadequate review capacity may produce a de facto ex parte form of administration, in which affected persons cannot meaningfully answer the technical assessment underlying a decision. Drawing on audi alteram partem, it proposes a risk-calibrated framework of notice, explanation, competent human review, contestability and institutional supervision. Informed, withdrawable consent has a role where participation is genuinely voluntary; mandatory public functions require effective procedural safeguards independent of individual consent. The central conclusion is that adopting EU-style safeguards requires corresponding legal, technical and institutional capacity to make review effective.

Keywords: EU AI Act, review capacity, technological dependency, procedural justice, public sector AI, contestability

I. INTRODUCTION

AI systems are increasingly used in administrative decision-making, particularly in areas where automated or AI-assisted decisions may affect fundamental rights, access to public services, legal status or individual entitlements. The European Union Artificial Intelligence Act is rapidly becoming the central reference point for contemporary global discussions of AI governance, a

* Author is a Ph.D. Candidate in International Law at School of Law, Xiamen University, Xiamen, China. ORCID: <https://orcid.org/0009-0009-0679-6656>

development that may be partly explained by the ‘Brussels Effect’.¹ Two features are especially important for the purposes of this article: first, the Act’s risk-based architecture, which classifies AI systems according to different levels of potential harm, and second, its procedural form of governance, which relies on transparency, documentation, logging, human oversight, and post-market monitoring.

Although the ‘Brussels Effect’ may help explain why EU-style AI rules travel beyond the Union, it does not follow that the institutional and technical conditions needed to make those rules effective travel with them.² Where public authorities depend on external vendors for the design, operation, maintenance or updating of AI systems, procedural safeguards depend not only on formal legal rules but also on the state’s practical review capacity.³ The risk, developed below, is that where review capacity is weak, AI-assisted decisions may operate as a *de facto ex parte* form of administration: public authorities may act on technical assessments that affected persons cannot know, understand or answer.

The assumption that procedural safeguards can be made effective becomes fragile where governments remain legally responsible for AI-assisted decisions without practical control over the systems behind them. This concern is particularly acute in high-risk settings, including judicial administration, policing, migration, welfare and digital identity, where decisions may affect fundamental rights, legal status or access to essential services.⁴

Public institutions in technologically dependent states, as well as affected citizens, may lack the practical means to examine, test or challenge AI-assisted decisions. The difficulty is not merely financial or technical. It may also arise from weak audit institutions, limited disclosure powers, dependence on vendor-controlled infrastructure, or the absence of domestic mechanisms capable of reviewing the system behind the decision.⁵ The AI Act itself does not make consent the general basis for lawful AI deployment. Instead, ‘consent’ appears in a limited and specific form, most clearly in relation to real-world testing of high-risk AI systems, while ordinary deployment is governed primarily through risk-based procedural safeguards.⁶ This limited role

¹ Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act), 2024 O.J. (L 2024/1689) [hereinafter AI Act]; ANU BRADFORD, THE BRUSSELS EFFECT: HOW THE EUROPEAN UNION RULES THE WORLD 1–6 (2020).

² See Bradford, *supra* note 1, at 1–6, 25–65.

³ See Cary Coglianese & Erik Lampmann, *Contracting for Algorithmic Accountability*, 6 ADMIN. L. REV. ACCORD 175, 179–80, 184–86, 198–99 (2021); DILLON REISMAN ET AL., ALGORITHMIC IMPACT ASSESSMENTS: A PRACTICAL FRAMEWORK FOR PUBLIC AGENCY ACCOUNTABILITY 3–4, 13–14, 20 (AI Now Inst. 2018); Danielle Keats Citron, *Technological Due Process*, 85 WASH. U. L. REV. 1249, 1253–54, 1298–1301, 1305–07 (2008).

⁴ AI Act, art. 6(2) and Annex III, points 1, 5–8.

⁵ See Coglianese & Lampmann, *supra* note 3, at 175, 184–86, 192–95, 198–99; Citron, *supra* note 3, at 1253–54, 1298–1301, 1305–07; Reisman et al., *supra* note 3, at 3–4, 13–14, 17–20.

⁶ AI Act, arts. 9–15, 26–27, 60–61 and 72–73.

of consent is significant because it shows that the EU model relies less on individual permission than on institutional capacity, procedural control and reviewability.⁷ Where those conditions are absent or weak, consent, notice and contestability become more important as safeguards against a *de facto ex parte* form of AI-assisted administration.

That limited role of consent creates the starting point for this article's claim, which is not that consent should operate as a universal veto over all public-sector AI use. Such a position would be unrealistic in public governance, where taxation, welfare, immigration, policing, digital identity and judicial processes cannot ordinarily depend on individual opt-in.⁸ The argument is that, in technologically dependent states, consent should be understood as part of a broader risk-calibrated procedural framework. Where AI systems are used in high-risk public functions, prior notice, meaningful consent where participation is genuinely voluntary, human review and contestability should operate as preconditions to legitimate deployment.⁹ In this sense, consent is not treated as a private act of individual permission alone but as one element of procedural justice. It functions as part of a wider mechanism through which affected persons are informed of AI involvement, given a meaningful opportunity to understand its role, and enabled to challenge or seek human reconsideration of AI-assisted outcomes.¹⁰ The article further argues that risk-calibrated consent is best understood as a modern procedural expression of *audi alteram partem*, ensuring that persons affected by AI-assisted public decisions are not governed by technical assessments that they cannot know, understand or answer.

This argument has important limits. It does not treat 'technological dependency' as a fixed or uniform condition but as a spectrum of practical control over system design, data processing, documentation, auditability, maintenance and exit capacity. The concern is not the use of external vendors as such, but the deployment of high-risk AI systems by states without sufficient mechanisms to explain, audit, contest, suspend or correct AI-assisted decisions. Nor does this article argue that consent and contestability are relevant only in technologically dependent states. Rather, it argues that they assume greater importance where institutional review capacity is weak.

⁷ For institutional capacity and reviewability, see AI Act, arts. 70(3), 74(12)–(13), 77, 85–86; Jennifer Cobbe, Michelle Seng Ah Lee & Jatinder Singh, *Reviewable Automated Decision-Making: A Framework for Accountable Algorithmic Systems*, in PROCEEDINGS OF THE 2021 ACM CONFERENCE ON FAIRNESS, ACCOUNTABILITY, AND TRANSPARENCY 598, 598–602 (2021), <https://doi.org/10.1145/3442188.3445921>.

⁸ For the limits of consent in public-authority contexts, see Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 (General Data Protection Regulation), recital 43, 2016 O.J. (L 119) 1 [hereinafter GDPR]; EUROPEAN DATA PROTECTION BOARD, *Guidelines 05/2020 on Consent under Regulation 2016/679* ¶¶ 13, 16–18 (Version 1.1, May 4, 2020), https://www.edpb.europa.eu/system/files/documents/files/file1/edpb_guidelines_202005_consent_en.pdf [hereinafter EDPB Consent Guidelines].

⁹ AI Act, art. 6(2) and Annex III; AI Act, arts. 13–14, 26, 50, 60–61, 85–86.

¹⁰ Citron, *supra* note 3, at 1298–1301, 1305–07; Cobbe et al., *supra* note 7, at 598–602.

The article uses doctrinal and conceptual analysis, drawing on the AI Act, public-law principles and selected comparative materials. It does not offer an empirical assessment of any state's capacity. Part II examines the Act as a model of risk-based procedural governance. Part III develops review capacity and distinguishes technological dependency from ordinary interdependence. Part IV examines the gap between formal safeguards and practical control. Part V proposes a risk-calibrated framework for high-risk public-sector AI. Part VI concludes.

II. THE EU AI ACT AS RISK-BASED PROCEDURAL GOVERNANCE

The EU AI Act entered into force on 1 August 2024 and applies in phases. Prohibited practices and AI literacy obligations began applying on 2 February 2025, followed by governance rules and general-purpose AI obligations on 2 August 2025.¹¹ On 29 June 2026, the Council gave its final approval to the Digital Omnibus on AI, since published as Regulation (EU) 2026/1744, which provides for the high-risk obligations to apply from 2 December 2027 for stand-alone high-risk systems listed in Annex III and from 2 August 2028 for high-risk systems embedded in regulated products covered by Annex I.¹²

This Part examines the Act as a model of risk-based procedural governance, focusing on the relationship between risk classification and transparency, documentation, logging, human oversight and post-market monitoring.¹³ The Act excludes AI systems placed on the market, put into service or used exclusively for military, defence or national security purposes.¹⁴

A. Risk classification and differentiated legal obligations

Article 5 prohibits specified AI practices, including certain manipulative or exploitative uses, social scoring and biometric practices, subject to the conditions and exceptions in that provision.¹⁵ For example, social scoring based on social behaviour or personal characteristics is prohibited where it leads to detrimental or unfavourable treatment in unrelated social contexts, or to treatment that is unjustified or disproportionate to the behaviour or its gravity. Denial of services or benefits must therefore be assessed against those statutory conditions; social scoring is not prohibited merely because it informs an adverse decision.¹⁶

¹¹ AI Act, art. 113; AI Act, recital 179.

¹² COUNCIL OF THE EUROPEAN UNION, Press Release, *Artificial Intelligence: Council Gives Final Green Light to Simplify and Streamline Rules* (June 29, 2026), <https://www.consilium.europa.eu/en/press/press-releases/2026/06/29/artificial-intelligence-council-gives-final-green-light-to-simplify-and-streamline-rules/>; Regulation (EU) 2026/1744 of the European Parliament and of the Council of 8 July 2026 Amending Regulations (EU) 2024/1689, (EU) 2018/1139 and (EU) 2023/1230 as Regards the Simplification of the Implementation of Harmonised Rules on Artificial Intelligence (Digital Omnibus on AI), art. 1(40), 2026 O.J. (L 2026/1744) (amending AI Act, art. 113).

¹³ AI Act, arts. 5–7, 9–15, 26–27, 72–73 and Annex III.

¹⁴ AI Act, art. 2(3); *see also* AI Act, recital 24.

¹⁵ AI Act, art. 5(1)(a)–(h).

¹⁶ AI Act, art. 5(1)(c).

The second category consists of high-risk AI systems. Unlike Article 5 systems, these are not prohibited, but they are permitted only subject to strict legal and procedural obligations.¹⁷ Article 6 identifies high-risk systems in two principal ways: first, where AI is linked to certain regulated products subject to third-party conformity assessment, and secondly, where the system falls within the use cases listed in Annex III.¹⁸ Annex III is particularly important for public administration because it includes AI systems used in areas such as biometrics, education, employment, access to essential public and private services, law enforcement, migration and asylum, and the administration of justice.¹⁹

Article 6(3) excludes certain Annex III systems that do not pose a significant risk of harm to health, safety or fundamental rights, including where they do not materially influence a decision, provided the specified conditions are satisfied. Systems that profile natural persons remain high-risk. Material influence is therefore an important, but not exclusive, marker of regulatory concern.²⁰ In judicial administration, a tool confined to file organisation may remain ancillary, while a system assisting the interpretation of facts or law, or the application of law to a concrete dispute, falls within the relevant Annex III use case.²¹

The Act also imposes targeted transparency duties, including in specified interactions with AI and uses of generated or manipulated content. These duties can apply independently of high-risk classification. Other systems may fall outside the high-risk regime, although general law, data-protection and consumer-protection rules may still apply, and Article 95 provides for voluntary codes of conduct.²²

B. Procedural safeguards for high-risk AI systems

High-risk AI systems are regulated through a set of *ex ante* and ongoing obligations designed to make them safe, transparent, traceable and subject to human control.²³ Articles 9 and 10 show that high-risk AI governance begins before deployment and continues through the system lifecycle, through risk management and data governance obligations aimed at identifying and reducing foreseeable harm. Articles 11, 12 and 13 then require technical documentation, record-keeping and transparency information. These obligations establish procedures for accountability: they create the information trail through which AI systems can later be

¹⁷ AI Act, arts. 6, 9–15, 26–27 and 72–73.

¹⁸ AI Act, art. 6(1)–(2) and Annexes I and III.

¹⁹ AI Act, Annex III, points 1, 3–8.

²⁰ AI Act, art. 6(3).

²¹ AI Act, Annex III, point 8(a); AI Act, recital 61.

²² AI Act, arts. 50, 95 and recitals 165–66.

²³ AI Act, arts. 9–15.

inspected, understood and challenged.²⁴ Articles 14 and 15 complete this framework by requiring human oversight, accuracy, robustness and cybersecurity. Human oversight, in this context, is not merely the formal presence of a human in the process; it requires a human actor capable of understanding, questioning and, where necessary, overriding the system.²⁵

The Act also imposes obligations on deployers of high-risk AI systems, not only on providers. Article 26 requires deployers to use such systems in accordance with instructions, ensure appropriate human oversight and monitor their operation.²⁶ Article 27 is especially important for public authorities, as it requires a fundamental-rights impact assessment before certain high-risk systems are deployed.²⁷ Articles 72 and 73 further show that compliance is not a one-time approval exercise: high-risk AI systems remain subject to post-market monitoring and serious-incident reporting after deployment.²⁸ The significance of these safeguards is that they presuppose review capacity. Documentation, logs, human oversight and post-market monitoring are meaningful only where public institutions, regulators, reviewing bodies and affected persons can access, interpret and act upon them.²⁹

C. The limited role of consent under the EU AI Act

Articles 60 and 61 illustrate the limited but important role of consent under the EU AI Act. Article 60 permits real-world testing of certain high-risk AI systems only under a controlled testing plan, authority supervision, safeguards, time limits and reversibility requirements.³⁰ Subject to the limited law-enforcement exception in Article 60(4)(i), Article 61 requires freely given informed consent from testing subjects before participation, after they have received concise, clear, relevant and understandable information about the nature and objectives of the testing and its possible inconvenience, the conditions and expected duration of their participation, their rights, including the right to refuse to participate and to withdraw at any time, and the arrangements for requesting the reversal or disregarding of the system's outputs.³¹ Consent, therefore, appears most clearly in the Act as a safeguard for experimental real-world testing, rather than as the general foundation for lawful AI deployment.

²⁴ AI Act, arts. 11–13; Cobbe et al., *supra* note 7, at 598–602.

²⁵ AI Act, arts. 14–15, especially art. 14(4)(a)–(e); Lena Enqvist, “Human Oversight” in the EU Artificial Intelligence Act: What, When and by Whom?, 15 L., INNOVATION & TECH. 508, 511–14, 528–31 (2023), <https://doi.org/10.1080/17579961.2023.2245683>.

²⁶ AI Act, art. 26.

²⁷ AI Act, art. 27.

²⁸ AI Act, arts. 72–73.

²⁹ AI Act, arts. 70(3), 71, 74(12)–(13), 77 and 85–86; Cobbe et al., *supra* note 7, at 598–602.

³⁰ AI Act, art. 60.

³¹ AI Act, art. 61(1)–(2).

Recital 141 supports this limited but important role of consent. It explains that real-world testing of high-risk AI systems may be permitted to support innovation, but only with appropriate safeguards, including informed consent of testing subjects, oversight by competent authorities, a testing plan, time limits, protection for vulnerable groups and the ability to reverse or disregard AI outputs.³²

III. TECHNOLOGICAL DEPENDENCY AND REVIEW CAPACITY

This inquiry is necessary because the EU AI Act is likely to function beyond the European Union as a regulatory reference point. Through the ‘Brussels Effect’, the Act may influence the design of AI governance frameworks in other jurisdictions, including states that rely on EU-style classifications, transparency duties, documentation requirements and human-oversight obligations when regulating AI in public administration.³³ Such influence matters because AI regulation in public governance does not operate in the abstract; it shapes how states use automated or AI-assisted systems. The central question is therefore whether states adopting EU-style procedural safeguards possess the technological, institutional and legal capacity to make those safeguards meaningful.³⁴ This question brings technological dependency and review capacity to the centre of the analysis.

A. Technological dependency as formal authority without practical control

Technological dependency should be distinguished from ordinary technological interdependence. No contemporary state is fully technologically self-sufficient, and absolute technological sovereignty is unrealistic in a world of global software ecosystems, cloud infrastructure, semiconductor supply chains, cybersecurity dependencies and cross-border data architectures. The concern is not foreign technology alone, which all states use to varying degrees, but a lack of practical control over the systems through which public power is exercised.

Technological dependency should not be equated with labels such as developed/developing states, Global North/Global South or post-colonial economic dependency. Those labels remain useful in identifying historical inequality, colonial legacies and structural economic asymmetries, but they do not always show where practical technological control lies.³⁵ China

³² AI Act, recital 141.

³³ See Bradford, *supra* note 1, at 1–6, 25–65.

³⁴ For the AI Act’s procedural safeguards and institutional conditions of review, see AI Act, arts. 9–15, 26–27, 70(3), 72–73, 74(12)–(13), 77 and 85–86; Cobbe et al., *supra* note 7, at 598–602.

³⁵ See generally M. SORNARAJAH, THE INTERNATIONAL LAW ON FOREIGN INVESTMENT ch. 1, § 1.2 (5th ed. 2021); Sumudu Atapattu & Carmen G. Gonzalez, *The North–South Divide in International Environmental Law: Framing the Issues*, in INTERNATIONAL ENVIRONMENTAL LAW AND THE GLOBAL SOUTH 1 (Shawkat Alam et al. eds., 2015).

and India, for example, are often located outside the traditional Global North, yet both possess significant state-backed capacity in advanced technologies.³⁶ National strategies in the United Kingdom and Canada and public support for quantum computing in Australia further illustrate that capability is actively cultivated through differing institutional arrangements. These examples contextualise technological capacity; they do not establish a state's ability to review public-sector AI.³⁷

For the purposes of this article, technological dependency refers to a condition in which a state retains formal legal authority over AI deployment but depends on external actors for the design, operation, maintenance, auditability or modification of the systems through which public decisions are made.³⁸ This definition encompasses several related dimensions, including design dependency, maintenance dependency, review dependency, regulatory dependency and exit dependency.

Design dependency may arise where the state cannot meaningfully determine or inspect the architecture, data flows or system logic of a technology-based system. Maintenance dependency may arise where updates, cybersecurity, repairs or model adjustments remain controlled by an external provider. Review dependency may arise where audit tools, logs, technical documentation or explanations are limited, unavailable or dependent on external-provider cooperation, making judicial or administrative review practically ineffective. Regulatory dependency may arise where domestic regulators lack the expertise, resources or institutional power to supervise the system effectively. Exit dependency may arise where the state cannot suspend, replace or migrate away from the system without disrupting essential public functions. This article focuses particularly on review dependency because the procedural safeguards associated with the EU AI Act are meaningful only where public institutions, regulators, courts and affected persons can access, understand and act upon the information needed to review AI-assisted public decisions.³⁹

³⁶ See DEP'T OF SCI. & TECH., GOV'T OF INDIA, *National Quantum Mission (NQM)* (Sept. 24, 2025), <https://dst.gov.in/national-quantum-mission-nqm>; NAT'L DEV. & REFORM COMM'N, *The 14th Five-Year Plan, Chapter 4: Boosting China's Strategic Science and Technology Capabilities* 23–24 (Apr. 25, 2022), <https://en.ndrc.gov.cn/policies/202204/P020220426646044910072.pdf>.

³⁷ DEP'T FOR SCI., INNOVATION & TECH., *National Quantum Strategy* 11–12 (Mar. 15, 2023), <https://www.gov.uk/government/publications/national-quantum-strategy>; INNOVATION, SCI. & ECON. DEV. CAN., *Canada's National Quantum Strategy* (Jan. 13, 2023), <https://ised-isde.canada.ca/site/national-quantum-strategy/en/canadas-national-quantum-strategy>; DEP'T OF INDUS., SCI. & RES., *Leading Quantum Company Chooses Australia as Site for Its Groundbreaking Utility Scale Quantum Computer* (Apr. 30, 2024), <https://www.industry.gov.au/news/leading-quantum-company-chooses-australia-site-its-groundbreaking-utility-scale-quantum-computer>.

³⁸ Coglianese & Lampmann, *supra* note 3, at 184–86, 192–95, 198–99; Citron, *supra* note 3, at 1253–54, 1298–1301, 1305–07; Reisman et al., *supra* note 3, at 3–4, 13–14, 17–20.

³⁹ AI Act, arts. 11–14, 26–27, 70(3), 72–73, 74(12)–(13), 77 and 85–86; Cobbe et al., *supra* note 7, at 598–602.

B. Review capacity in technologically dependent states

Existing scholarship on human oversight has warned that human involvement may create only an appearance of accountability where the human reviewer lacks the information, competence, authority or institutional conditions necessary to evaluate and depart from the system's output. Review capacity should therefore be understood as having both technical and institutional dimensions. Technically, it depends on the ability to access and understand the hardware, software, data infrastructures and documentation through which public AI systems operate.⁴⁰ Institutionally, it concerns all three organs of the state, although in different ways: the executive must procure, operate and supervise systems; the legislature must set the legal conditions for transparency, audit and exit; and courts must be able to review AI-assisted decisions. Where a state's review capacity over the design, maintenance and operation of its digital infrastructure is weakened, dependency may become entrenched. The state may then find it difficult to regulate, modify, suspend or exit unfavourable technological arrangements, with consequences for the practical conditions under which sovereignty is exercised.⁴¹

This makes review capacity closely connected to meaningful human oversight. Human oversight is not effective merely because a human official is formally present in the decision-making process. Its effectiveness depends on what the human reviewer is able to examine, when intervention is possible, and whether the reviewer has the competence, training, authority and institutional support needed to question or depart from the system's output. A public official who cannot access relevant logs, understand system limitations, identify automation bias, interpret the output, or disregard, override or suspend the system cannot meaningfully perform oversight. In such circumstances, human involvement risks becoming a formal safeguard rather than a substantive check on AI-assisted public power.

Weak review capacity affects the judiciary by making judicial review of administrative decisions formally available but practically ineffective. For example, a retired public servant may be denied a pension entitlement, gratuity payment or welfare-related benefit on the basis of an AI-assisted administrative system designed, hosted or maintained by an external provider. Even where domestic law provides a formal right of review, that right may be hollow if the affected person cannot obtain meaningful reasons, the court cannot access relevant logs or documentation, and the public authority itself cannot explain how the system influenced the decision.⁴² Enforcement may also be complicated where the relevant data, infrastructure or

⁴⁰ Cobbe et al., *supra* note 7, at 598–602; AI Act, arts. 11–14, 26–27, 70(3), 74(12)–(13), 77 and 85–86.

⁴¹ See Sornarajah, *supra* note 35, ch. 1; Coglianese & Lampmann, *supra* note 3, at 184–86, 198–99.

⁴² Citron, *supra* note 3, at 1298–1301; Cobbe et al., *supra* note 7, at 598–602; AI Act, arts. 11–14, 74(12)–(13), 77 and 85–86.

provider is subject to legal regimes outside the reviewing state.⁴³ This concern is intensified where systems generate opaque, probabilistic or unsupported outputs that cannot be adequately verified, including but not limited to generative AI hallucinations.⁴⁴

The consequences extend beyond judicial review. For the executive, weak review capacity may impair cybersecurity supervision, data governance, administrative enforcement, procurement oversight and policy autonomy.⁴⁵ Public authorities may remain formally responsible for AI systems that they cannot adequately monitor, correct, suspend or control. For the legislature, the problem is constitutional as well as regulatory. Weak review capacity may generate sovereignty concerns, weaken the protection of fundamental rights and produce blind law-making: legislation for digital governance that assumes technical capacities, such as access, auditability, interoperability, explainability and exit, which the state does not in fact possess.⁴⁶

C. The accountability gap

The combination of technological dependency and weak review capacity produces an accountability gap. Public authorities remain legally responsible for decisions made in their name, but the technical conditions necessary to explain, audit, correct or defend those decisions may depend on external providers.⁴⁷ In such circumstances, accountability becomes divided: legal responsibility remains with the state, while practical control over system design, documentation, logs, updates, auditability or infrastructure may lie elsewhere. The result is not merely administrative inconvenience but a deeper mismatch between public-law responsibility and technological control.

This gap is especially serious where AI-assisted systems are used in high-risk public functions.⁴⁸ A state may formally promise transparency, human oversight, reasons, review or remedies, yet lack the practical capacity to make those safeguards effective. The problem is therefore not only whether legal safeguards exist, but also whether the state can operationalise them against the technical systems through which public power is exercised. This is the point at which technological dependency becomes directly relevant to risk-calibrated consent: where the state cannot fully explain, audit or control the system, legitimacy cannot rest on formal

⁴³ Coglianese & Lampmann, *supra* note 3, at 184–86, 192–95, 198–99; Sornarajah, *supra* note 35, ch. 1.

⁴⁴ NATIONAL INSTITUTE OF STANDARDS AND TECHNOLOGY, *Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile* 4–5, 44–45 (NIST AI 600-1, 2024).

⁴⁵ AI Act, arts. 9–10, 15, 26 and 72–73; Coglianese & Lampmann, *supra* note 3, at 184–86, 192–95, 198–99; Reisman et al., *supra* note 3, at 13–14, 17–20.

⁴⁶ AI Act, arts. 11–14, 27, 70(3), 74(12)–(13), 77 and 85–86; Sornarajah, *supra* note 35, ch. 1; Cobbe et al., *supra* note 7, at 598–602.

⁴⁷ Coglianese & Lampmann, *supra* note 3, at 184–86, 192–95, 198–99; Citron, *supra* note 3, at 1298–1301, 1305–07; Cobbe et al., *supra* note 7, at 598–602; AI Act, arts. 11–14, 26, 72–73 and 85–86.

⁴⁸ AI Act, art. 6(2) and Annex III, especially points 1, 5–8; *see also* arts. 26–27.

deployment authority alone. It requires, at minimum, prior notice; meaningful human review; contestability; and dependency-sensitive disclosure and, where participation is genuinely voluntary or experimental, informed and withdrawable consent.⁴⁹

IV. THE LIMITS OF PROCEDURAL SAFEGUARDS IN TECHNOLOGICALLY DEPENDENT STATES

Once technological dependency and review capacity are identified, the next question is what they reveal about the EU AI Act's procedural model. The issue is not the absence of safeguards. The Act contains a detailed procedural architecture built around risk management, transparency, documentation, logging, human oversight, deployer duties, explanation rights and post-market monitoring.⁵⁰ The harder question is whether those safeguards can operate effectively where the institutions expected to use them lack the technical, financial or legal capacity to access, understand and act upon the relevant information. This Part examines the institutional assumptions behind those safeguards and the risk that technological dependency may reduce them to formal indicators of compliance.⁵¹

The consequence is not merely defective compliance. Where review capacity is weak, the accountability gap becomes procedural. Affected persons may face an administrative outcome shaped by a data match, classification, risk score, model limitation or system-generated recommendation that neither they nor the public authority can adequately explain. The result is not only opacity but also one-sidedness: public power is exercised on the basis of a technical assessment that the affected person cannot meaningfully answer. Part V treats this as a problem of *audi alteram partem* and develops a risk-calibrated response based on consent, notice, human review and contestability.

A. Review capacity as an implied assumption of the EU AI Act

The EU AI Act does not expressly describe its safeguards in terms of review capacity. Nevertheless, the Act's procedural architecture depends on it. Duties of technical documentation, logging, transparency, human oversight, post-market monitoring and explanation assume that relevant actors can access, understand and use system information.⁵²

⁴⁹ AI Act, arts. 13–14, 26–27, 60–61 and 85–86; Citron, *supra* note 3, at 1298–1301, 1305–07; Cobbe et al., *supra* note 7, at 598–602.

⁵⁰ AI Act, arts. 9–15, 26–27, 72–73 and 85–86.

⁵¹ Cobbe et al., *supra* note 7, at 598–602; Citron, *supra* note 3, at 1298–1301, 1305–07; Coglianese & Lampmann, *supra* note 3, at 184–86, 198–99.

⁵² AI Act, arts. 11–14, 26, 72–73 and 85–86; *see also* arts. 70(3), 74(12)–(13) and 77.

The Act therefore presupposes not merely legal duties on providers and deployers but institutional conditions under which those duties can be made effective.⁵³

This assumption is most visible in relation to high-risk AI systems. Technical documentation matters only if regulators, deployers and reviewing bodies can understand and rely on it. Logging supports accountability only if logs are accessible, reliable and usable in review. Human oversight is meaningful only if the human reviewer has the authority, competence and information needed to question or depart from the system's output. A right to explanation is valuable only if the explanation is specific enough to show how the system contributed to the decision affecting the person. The AI Act's safeguards are therefore procedural, but they are not self-executing.⁵⁴

The assumption becomes fragile when EU-style safeguards are adopted in settings where public authorities lack practical control over the systems they deploy. The 'Brussels Effect' may help explain why states borrow the legal form of the AI Act, but it does not reproduce the legal, institutional and technical conditions on which that form depends. A formal right to complain or obtain an explanation may become hollow where the evidence needed to make that right meaningful remains beyond effective legal control, or where affected persons and institutions lack the financial and technical capacity to obtain expert assistance.⁵⁵

B. Transparency, documentation and logging

The EU AI Act's procedural model combines transparency with documentation and logging safeguards. Articles 11 and 12 require technical documentation and record-keeping for high-risk AI systems, while Article 13 requires such systems to be sufficiently transparent to enable deployers to interpret outputs and use the system appropriately.⁵⁶ The Act's transparency model is not limited to information provided to deployers. Article 26 imposes additional information duties on deployers, including duties to inform affected workers in workplace contexts and natural persons subject to certain high-risk AI-assisted decision-making.⁵⁷ Article 50 imposes transparency duties toward natural persons in specific AI interactions, including systems that interact directly with natural persons and AI-generated or manipulated content, some of which

⁵³ Cobbe et al., *supra* note 7, at 598–602; Citron, *supra* note 3, at 1298–1301, 1305–07; Coglianese & Lampmann, *supra* note 3, at 184–86, 198–99.

⁵⁴ AI Act, arts. 70(3), 72–73, 74(12)–(13), 77 and 85–86; Cobbe et al., *supra* note 7, at 598–602; Citron, *supra* note 3, at 1298–1301, 1305–07.

⁵⁵ AI Act, arts. 85–86; Citron, *supra* note 3, at 1298–1301, 1305–07; Cobbe et al., *supra* note 7, at 598–602; Reisman et al., *supra* note 3, at 17–20.

⁵⁶ AI Act, arts. 11–13.

⁵⁷ AI Act, art. 26(7) and (11).

are not necessarily high-risk.⁵⁸ Article 86 further recognises, within its limited scope, a right to obtain clear and meaningful explanations of the role of a high-risk AI system in certain individual decisions producing legal or similarly significant adverse effects.⁵⁹

These provisions are important, but they also reveal the limits of transparency as a review safeguard. Article 26(11), for example, may be read as a limited notice provision rather than a full contestability mechanism. It requires deployers of Annex III high-risk AI systems that make, or assist in making, decisions concerning natural persons to inform those persons that they are subject to the use of the high-risk AI system.⁶⁰ It does not, however, itself require that notice to explain the person's rights, including the right to seek an explanation under Article 86, the right to lodge a complaint under Article 85, the availability of domestic review or challenge procedures, or the technical material needed to contest the AI-assisted decision.⁶¹

Documentation and logging raise a related problem. Articles 11 and 12 require technical documentation and record-keeping for high-risk AI systems, while Article 13 requires information enabling deployers to understand the system's characteristics, capabilities and limitations. These obligations create an information trail. Yet an information trail does not automatically produce accountability. In technologically dependent states, technical documentation may remain with external providers; logs may be inaccessible, incomplete or stored in foreign infrastructure; audit records may be protected by contract or trade secrecy; and public authorities may lack the expertise to interpret the materials even where they are available.⁶²

C. Human oversight, contestability and competent authority

Article 14 of the EU AI Act is one of the stronger procedural safeguards in the regulation. It does not treat human oversight as the mere formal presence of a human decision-maker. Article 14(4)(a)–(e) requires oversight measures that enable the human overseer to understand the capacities and limitations of the high-risk AI system, remain aware of automation bias, correctly interpret the system's output, decide not to use or to disregard, override or reverse that output, and intervene in or interrupt the system where necessary. This is reinforced by Article 26(2), which requires deployers to assign human oversight to natural persons with the necessary

⁵⁸ AI Act, art. 50(1)–(4).

⁵⁹ AI Act, art. 86.

⁶⁰ AI Act, art. 26(11).

⁶¹ AI Act, arts. 85–86; *see also* Citron, *supra* note 3, at 1298–1301, 1305–07; Cobbe et al., *supra* note 7, at 598–602.

⁶² AI Act, arts. 11–13; Coglianese & Lampmann, *supra* note 3, at 184–86, 192–95, 198–99; Cobbe et al., *supra* note 7, at 598–602; Reisman et al., *supra* note 3, at 17–20.

competence, training, authority and support.⁶³ The difficulty therefore arises not from the design of the AI Act itself, but from the external conditions required to make human oversight effective.

In technologically dependent states, these oversight functions may be difficult to perform where public authorities lack the technical understanding, institutional authority, access to system information or practical control necessary to supervise the system in practice. The Act assumes that the institutions responsible for checking illegality, unfairness or rights-based harm will have the capacity to understand, inspect and verify the basis of the AI system's output. That assumption may not hold where the system is designed, hosted, updated or maintained by an external provider. In such cases, the legal inquiry becomes more complex because relevant information may be controlled by private contracts, foreign infrastructure, trade secrecy, data-protection rules or competing legal regimes.⁶⁴ Human oversight may then become formal rather than substantive: a human official remains legally present in the decision-making chain but is unable to meaningfully question, test or depart from the AI system's output.

The AI Act also recognises procedural routes for contestation. Article 85 allows any natural or legal person to submit a complaint to the relevant market surveillance authority where they consider that the regulation has been infringed, while expressly preserving other administrative or judicial remedies. This complaint mechanism is supported by Article 70, which requires Member States to establish or designate national competent authorities, including at least one market surveillance authority. Article 70(3) further requires those authorities to be provided with adequate technical, financial and human resources, as well as infrastructure, to fulfil their tasks effectively under the Regulation.⁶⁵

For present purposes, the significance of this institutional structure is not that technologically dependent states should reproduce the EU model exactly. The point is that the AI Act's complaint and oversight mechanisms assume the existence of a technically competent public authority with adequate financial, human and technical resources. Without such an institution, rights of complaint, explanation and review may remain formally available but practically weak. Courts, individual complainants and ordinary public authorities may not possess the technical capacity to test high-risk AI systems on their own. The institutional lesson of the AI Act is therefore that contestability requires more than individual rights on paper. It requires a public body capable of obtaining technical material, examining system behaviour, supervising

⁶³ AI Act, arts. 14(4)(a)–(e) and 26(2).

⁶⁴ Coglianese & Lampmann, *supra* note 3, at 184–86, 192–95, 198–99; Cobbe et al., *supra* note 7, at 598–602.

⁶⁵ AI Act, arts. 70(1), 70(3) and 85.

providers and requiring corrective action where high-risk AI systems affect rights, legal status or essential interests.⁶⁶

The consequences of weak contestability are illustrated by the Michigan Integrated Data Automated System, or MiDAS, used in the administration of unemployment benefits in the United States. In litigation arising from the system, claimants alleged that automated fraud determinations led to the immediate termination of benefits, severe monetary penalties and enforcement action, while many affected persons did not receive meaningful notice until the time for appeal had expired. The example is not important because MiDAS was an advanced AI system but because it shows how automated public decision-making can become procedurally one-sided when affected persons lack timely notice, meaningful reasons and an accessible opportunity to contest the system's output.⁶⁷

A modest hypothetical based on Sri Lanka's proposed Unique Digital Identity project illustrates how the same problem may arise in a technologically dependent state. The June 2025 draft procurement contract for SL-UDI provides that arbitration under the contract 'shall be heard in New Delhi, India', although the same clause also states that the place of arbitration is to be neutral and decided mutually by the Master System Integrator and the Government of Sri Lanka.⁶⁸ That contractual arrangement would not, by itself, require an ordinary retired public servant to arbitrate abroad. The problem is more practical. If a retired public servant is denied a pension entitlement, gratuity payment or related public benefit because an AI-assisted identity-verification or administrative-matching process produces a data mismatch, fraud flag or eligibility error, the person's domestic remedy may depend on logs, audit material or system explanations controlled by the external Master System Integrator or embedded in the state-provider contract. The public authority remains legally responsible for the administrative decision, but the technical material needed to explain or correct that decision may lie outside the affected person's practical reach because of contractual confidentiality, vendor control, cybersecurity restrictions or conflict-of-laws issues arising from dispute resolution conducted outside Sri Lanka.⁶⁹

⁶⁶ AI Act, arts. 70(1), 70(3), 74(12)–(13), 77 and 85–86; Cobbe et al., *supra* note 7, at 598–602; Citron, *supra* note 3, at 1298–1301, 1305–07.

⁶⁷ *Cahoo v. SAS Analytics Inc.*, 912 F.3d 887, 893–95, 900–04 (6th Cir. 2019); *see also Bauserman v. Unemployment Ins. Agency*, 983 N.W.2d 855 (Mich. 2022).

⁶⁸ NATIONAL INSTITUTE FOR SMART GOVERNMENT, *Bidding Document – Draft Contract: Appointment of a Master System Integrator for Development, Implementation and Maintenance of the Unique Digital Identity (SL-UDI) Project of Government of Sri Lanka*, vol. 3, PC cl. 51.2(c) (Invitation for Bids No. NISG/SLUDI-2025, June 2025), https://www.nisg.org/files/ugd/ce4487_672ee6cc86844a9ab77420d6071359c1.pdf [hereinafter SL-UDI Draft Contract].

⁶⁹ SL-UDI Draft Contract, *supra* note 68, GCC cls. 1.12, 1.24, 1.51–1.52, 7.1, 15–17, 21.3, 25, 52, PC cls. 1.1.12, 1.1.62, 5.1, 12.2, 15.3–15.4, 16.1 (cited for contractual provisions concerning practical control, confidentiality,

V. RISK-CALIBRATED CONSENT, NOTICE AND CONTESTABILITY

The preceding Parts argued that the EU AI Act provides a sophisticated procedural model for high-risk AI but that its effectiveness depends on review capacity. In technologically dependent states, transparency, documentation, logging, human oversight and complaint mechanisms may exist in legal form while remaining difficult to operationalise in practice. This Part develops a response to that accountability gap grounded in *audi alteram partem*, in which consent, notice, human review, contestability and institutional supervision are adjusted to the seriousness of the decision and the degree of technological dependency.

A. *Audi alteram partem* and one-sided AI administration

The principle of *audi alteram partem* reflects a basic requirement of procedural justice: a person affected by public power should know the substance of the case against them and should be given a meaningful opportunity to answer it.⁷⁰ Its roots lie deep in the common-law tradition of natural justice and fair hearing, a tradition often associated with the broader legacy of Magna Carta and the idea that public power must be exercised according to law.⁷¹ For EU institutions, bodies, offices and agencies, Article 41 of the Charter of Fundamental Rights links good administration to the right to be heard, access to the file and the duty to give reasons. AI-assisted public decision-making should not be treated as an exception to that tradition.⁷² The fact that a decision is supported by a technical system does not reduce the need for notice, reasons or an opportunity to respond where the decision affects rights, legal status or essential interests.⁷³

The difficulty is that AI-assisted decision-making may alter the practical conditions under which a person can be heard. In ordinary administration, the affected person may respond to an allegation, document, factual finding or official assessment. In AI-assisted administration, however, the material basis of the decision may take the form of a data match, risk score, classification, fraud flag or system-generated recommendation. The affected person may be invited to challenge the final decision while being unable to know, understand or answer the

system access, documentation, cybersecurity and dispute resolution); Coglianese & Lampmann, *supra* note 3, at 184–86, 192–95, 198–99; Cobbe et al., *supra* note 7, at 598–602.

⁷⁰ *Cooper v. Wandsworth Board of Works* (1863) 14 C.B. (N.S.) 180, 190 (Eng.); *Ridge v. Baldwin* [1964] A.C. 40 (H.L.) 64–65 (appeal taken from Eng.); *Lloyd v. McMahon* [1987] A.C. 625 (H.L.) 702 (appeal taken from Eng.); *R v. Secretary of State for the Home Department, ex parte Doody* [1994] 1 A.C. 531 (H.L.) 560 (appeal taken from Eng.).

⁷¹ Magna Carta ch. 39 (1215); see generally J.C. HOLT, *MAGNA CARTA* (3d ed. 2015); TOM BINGHAM, *THE RULE OF LAW* chs. 1–2 (2011).

⁷² Charter of Fundamental Rights of the European Union art. 41(2)(a)–(c), 2012 O.J. (C 326) 391; Citron, *supra* note 3, at 1298–1301, 1305–07; Enqvist, *supra* note 25, at 511–14.

⁷³ Citron, *supra* note 3, at 1253–54, 1298–1301; Emily M. Weitzenboeck, *Simplification of Administrative Procedures through Fully Automated Decision-Making: The Case of Norway*, 11 ADMIN. SCIS., no. 4, art. 149, §§ 1, 3.2 (2021), <https://doi.org/10.3390/admsci11040149>.

technical assessment that materially shaped it. In this sense, AI-assisted decision-making may become procedurally one-sided.⁷⁴

This risk is heightened because reliance on AI may reduce human interaction and contextual judgement in public administration. Administrative decisions often require attention to personal circumstances, local context, vulnerability, explanation and discretion. If public officials treat AI outputs as authoritative, neutral or technically superior, they may become mere rubber stamps rather than responsible decision-makers.⁷⁵ The danger is not only that AI may produce inaccurate or biased outputs but also that public power may be exercised through mechanical assessments that affected persons cannot meaningfully contest. *Audi alteram partem* therefore requires more than the formal availability of an appeal after the event. It requires notice, explanation, human review and contestability where AI materially influences decisions affecting rights, status, benefits, liberty or access to essential services.⁷⁶

B. The limits of consent in mandatory public functions

Where AI is used in genuinely voluntary or experimental settings, consent should be express, informed and withdrawable.⁷⁷ Where AI is used in mandatory public functions, however, consent in the narrow sense may be fictional, because the affected person cannot realistically refuse the service, benefit or legal process merely because AI is used.⁷⁸ Nor does the mere existence of human supervision cure this problem. Human supervision matters only if the reviewer has the competence, information, authority and institutional support needed to understand, evaluate and, where necessary, depart from the system's output. In such cases, the emphasis should shift from consent as individual permission to safeguards that make the AI-assisted decision knowable, reviewable and contestable.⁷⁹

The calibration should become stricter as the consequences of the decision become more serious. If AI materially influences decisions affecting rights, legal status, liberty, public benefits, identity, migration, policing, education, health or judicial outcomes, the affected person should be informed of the AI's role, the human authority responsible for the final

⁷⁴ Citron, *supra* note 3, at 1253–54, 1298–1301; Weitzenboeck, *supra* note 73, §§ 3.1.1, 3.2, 4.4.

⁷⁵ Enqvist, *supra* note 25, at 511–14, 528–31; see generally Ben Green, *The Flaws of Policies Requiring Human Oversight of Government Algorithms*, 45 COMPUT. L. & SEC. REV. 105681 (2022), <https://doi.org/10.1016/j.clsr.2022.105681>; Kiel Brennan-Marquez, Karen Levy & Daniel Susser, *Strange Loops: Apparent versus Actual Human Involvement in Automated Decision Making*, 34 BERKELEY TECH. L.J. 745 (2019).

⁷⁶ AI Act, arts. 13–14 and 85–86; Citron, *supra* note 3, at 1298–1301; Enqvist, *supra* note 25, at 519–24.

⁷⁷ AI Act, arts. 60–61; AI Act, recital 141.

⁷⁸ GDPR, recital 43; EDPB Consent Guidelines, *supra* note 8, ¶¶ 13, 16–18.

⁷⁹ AI Act, arts. 14(4)(a)–(e), 26(2) and 85–86; Enqvist, *supra* note 25, at 511–14, 528–31; Citron, *supra* note 3, at 1298–1301, 1305–07; Cobbe et al., *supra* note 7, at 598–602.

decision, and the procedure available for review.⁸⁰ Where the state is technologically dependent on an external provider, additional safeguards should apply, including disclosure of vendor involvement, local access to logs and documentation, independent audit capacity and the ability to suspend or disregard unreliable outputs.⁸¹

This distinction is especially important in adjudication and other high-risk public functions. The EU AI Act partly reflects this concern by treating certain AI systems used by or on behalf of judicial authorities to assist in researching and interpreting facts or law, or applying the law to a concrete set of facts, as high-risk.⁸² Where AI materially influences the reasoning or outcome in such settings, affected persons should have access to meaningful human reconsideration by a person with authority to depart from the AI-assisted output. Human review should be treated as a fairness safeguard, not as a premium service available only to those who can afford it.⁸³

C. A risk-calibrated procedural model

A risk-calibrated model should not impose identical procedural requirements on every use of AI in public administration. The appropriate safeguard should depend on three factors: the seriousness of the decision, the practical possibility of refusal, and the degree of technological dependency involved. Where AI is used for low-risk internal functions, such as document organisation, scheduling, translation for internal comprehension or administrative triage, and where the output does not materially influence a decision affecting a person, full individual consent may be unnecessary, provided that the output is professionally verified before it is relied upon. Where AI is used in genuinely voluntary or experimental settings, consent should be express, informed and withdrawable.⁸⁴ Where AI is used in mandatory public functions, however, consent in the narrow sense may be fictional, because the affected person cannot realistically refuse the service or legal process. In such cases, the relevant safeguards should be prior notice, meaningful explanation, access to human review and a practical right to contest the AI-assisted outcome.⁸⁵

⁸⁰ AI Act, art. 6(2) and Annex III, especially points 1, 5–8; AI Act, arts. 13–14, 26 and 85–86; Enqvist, *supra* note 25, at 511–14, 528–31.

⁸¹ AI Act, arts. 11–14, 26, 70(3), 74(12)–(13), 77 and 85–86; Coglianese & Lampmann, *supra* note 3, at 184–86, 192–95, 198–99; Cobbe et al., *supra* note 7, at 598–602.

⁸² AI Act, Annex III, point 8(a); AI Act, recital 61.

⁸³ AI Act, arts. 14(4)(a)–(e), 26(2) and 85–86; Enqvist, *supra* note 25, at 511–14, 528–31; Citron, *supra* note 3, at 1298–1301, 1305–07.

⁸⁴ AI Act, arts. 60–61.

⁸⁵ GDPR, recital 43; EDPB Consent Guidelines, *supra* note 8, ¶¶ 13, 16–18; AI Act, arts. 13–14, 26, 85–86.

The proposed framework therefore escalates with both decisional risk and review dependency. Internal assistance requires proportionate professional supervision; material influence over rights or essential services requires notice, an identifiable responsible authority and accessible review. Where vendors control the evidence needed for review, procurement and supervision should secure access to relevant logs and documentation, independent audit and the ability to suspend or disregard unreliable outputs. These are proposed institutional safeguards, drawing on the Act's procedural architecture, rather than a claim that the Act imposes identical duties on every public use of AI.⁸⁶

D. Minimum safeguards for high-risk public-sector AI

Minimum safeguards are necessary because AI can make public administration faster by reducing human involvement. That efficiency is not inherently undesirable. AI may help public authorities process information, reduce delay, organise records and improve access to services. The danger is that efficiency may be pursued by treating human participation as friction to be removed, rather than as a source of legality, accountability and contextual judgement. Public-sector AI governance should therefore proceed from the premise that AI may assist administration but should not replace responsible public decision-making where decisions affect rights, legal status or essential interests.⁸⁷

Minimum safeguards must also be dependency-sensitive. AI systems may carry assumptions drawn from the legal, administrative and technical traditions of technologically capable states. Those assumptions may not fit the legal culture, administrative practice or social realities of technologically dependent states. Public officials may also lack the technical training needed to identify when an AI system is applying unfamiliar logic or producing an output inconsistent with domestic legal doctrine. The proper safeguard is therefore not blind trust in either the human or the machine but explainable AI, reviewable outputs and competent human responsibility.⁸⁸

These safeguards require institutional support. A technologically dependent state cannot rely only on individual complaints, ordinary judicial review or internal administrative supervision to control high-risk AI systems. Courts, affected persons and ordinary public officials may not have the technical capacity to obtain logs, test system behaviour, evaluate model limitations or

⁸⁶ AI Act, arts. 11–14, 26, 50 and Annex III; Coglianese & Lampmann, *supra* note 3, at 184–86, 192–95, 198–99; Cobbe et al., *supra* note 7, at 598–602.

⁸⁷ Enqvist, *supra* note 25, at 511–14, 528–31; *see generally* Green, *supra* note 75; Brennan-Marquez et al., *supra* note 75.

⁸⁸ Citron, *supra* note 3, at 1298–1301, 1305–07; Cobbe et al., *supra* note 7, at 598–602; Coglianese & Lampmann, *supra* note 3, at 184–86, 192–95, 198–99.

determine whether a vendor-controlled system is operating lawfully. A competent public authority should therefore be empowered to supervise high-risk AI systems used in public administration, receive complaints, obtain relevant documentation and logs, conduct or require audits, investigate serious incidents and require corrective action.⁸⁹

Such an authority should be technically competent but not purely technocratic. It should be connected to democratic oversight and should report periodically to the legislature on the use of high-risk AI systems in public administration, complaints received, corrective measures ordered, dependency risks arising from foreign or private control over critical AI infrastructure, and the adequacy of domestic technical and human resources. There should also be an accessible review pathway capable of obtaining technical material, requiring explanations, hearing affected persons and recommending or ordering corrective action. This approach is consistent with Council of Europe guidance on algorithmic systems, which emphasises human rights impact assessment, audit, monitoring, accountability, redress and effective transparency in public-sector algorithmic decision-making.⁹⁰ In this sense, risk-calibrated consent is also dependency-sensitive governance: it requires AI law to match both the risks of the technology and the real capacities of the state in which it operates.

VI. CONCLUSION

The EU AI Act offers an important procedural model for AI governance. Its risk-based structure, transparency duties, documentation, logging, human oversight, fundamental-rights impact assessment and post-market monitoring show that high-risk AI should be reviewable, contestable and institutionally supervised. Its central lesson is that AI governance cannot depend on risk classification alone; it also requires procedures capable of making technical systems accountable.

This article has argued that those procedures are not self-executing. They depend on review capacity. In technologically dependent states, public authorities may remain legally responsible for AI-assisted decisions while lacking practical control over the systems that shape them. Where system design, maintenance, documentation, logs, audit tools or explanations are controlled by external providers, safeguards may exist in law but fail in practice. The result is an accountability gap between public-law responsibility and technological control.

⁸⁹ AI Act, arts. 11–14, 26–27, 70(3), 74(12)–(13), 77, 85–86; Coglianese & Lampmann, *supra* note 3, at 184–86, 192–95, 198–99; Reisman et al., *supra* note 3, at 13–14, 17–20.

⁹⁰ COUNCIL OF EUR., COMM. OF MINISTERS, *Recommendation CM/Rec(2020)1 to Member States on the Human Rights Impacts of Algorithmic Systems*, app. ¶¶ A.1, A.3, B.1, B.3.3, B.4, B.5.2–5.3, C.1–C.3 (Apr. 8, 2020).

That gap has a procedural consequence. AI-assisted public decision-making may become a *de facto ex parte* form of administration where affected persons are governed by technical assessments that they cannot know, understand or answer. This risk is especially serious in welfare, migration, policing, digital identity and judicial administration, where decisions may affect rights, legal status, liberty, public benefits or access to essential services. In such settings, *audi alteram partem* requires more than an appeal after the event. It requires notice, explanation, human reconsideration and a practical opportunity to contest the technical basis of the decision.

The answer is not to reject AI in public administration or to treat consent as a universal veto. AI may improve efficiency, reduce delay and expand access to services, while many public functions cannot realistically depend on individual opt-in. The better approach is risk-calibrated consent, understood as part of a broader procedural framework. In voluntary or experimental settings, consent should be express, informed and withdrawable. In mandatory public functions, the emphasis should shift to prior notice, meaningful explanation, competent human review, accessible contestability and institutional supervision.

For technologically dependent states, the challenge is not merely to adopt EU-style safeguards but to build the review capacity needed to make them real. Public officials should not be reduced to rubber stamps for technical outputs they cannot properly examine, explain or override. Nor should affected persons be left to challenge decisions without access to the material needed to understand how AI shaped the outcome. These safeguards are therefore necessary to preserve procedural justice where public power is exercised through technical systems.
