International Journal of Science, Technology and Society
2026, Vol. 14, No. 5, pp. 222–237
https://doi.org/10.11648/j.ijsts.20261405.13

SciencePG
Science Publishing Group

Research Article

Regulating the Visible: the EU AI Act, Labour Governance, and the Limits of a Product-Risk Framework

Achi Iseko*

Independent Researcher, Surrey, United Kingdom

Abstract

The European Union Artificial Intelligence Act (Regulation (EU) 2024/1689) establishes a comprehensive legal framework for artificial intelligence (AI) governance. Organised around a risk-based classification of AI systems and their deployment contexts, the Act makes workers visible principally as persons affected by AI decisions, not as the distributed workforce whose labour, judgement and exposure produce AI systems in the first place. This paper presents a systematic regulatory audit of the Act using the Ethical Visibility Framework (EVF), a three-dimensional analytical framework for analysing material, epistemic and narrative dimensions of labour invisibility in AI production. Applying EVF indicators to the Act's operative articles, recitals, annexes and current implementing instruments, we find that the Act weakly recognises AI labour relations across all three dimensions: it governs the data artefact produced by annotation work without specifying the conditions under which that work is performed (material); it positions human oversight as deployer control over system outputs rather than as worker participation in system governance (epistemic); and its operative vocabulary recognises data-production activities without defining corresponding production-worker entitlements (narrative). We identify and examine the Act's implicit rival logic: that regulation of AI labour is intentionally deferred to existing EU labour law and the Platform Work Directive. We examine why this division of regulatory functions can leave gaps for transnational AI production work, illustrated through secondary evidence concerning Sub-Saharan Africa and comparative material on the Arab Gulf. We term this the enforcement paradox: a regulatory architecture with significant sanctioning powers but no independent, general duty governing production-labour conditions. We conclude with targeted proposals for labour-inclusive amendments to the Act's system and model governance frameworks, arguing that a product-risk regime cannot adequately guarantee trustworthy AI while giving insufficient attention to the labour conditions that underpin it.

Keywords

EU AI Act, AI Labour, Labour Invisibility, Ethical Visibility Framework, Regulatory Audit, Data Annotation, Worker Participation, Supply Chain Transparency

1. Introduction

Since Turing [1] first asked whether machines can think, the governance of artificial intelligence (AI) has been preoccupied with the capacities of the machine. Who built it, under what conditions, and at what human cost have remained secondary questions, when they have been asked at all. The European Union (EU) Artificial Intelligence Act, published in the Official Journal on 12 July 2024 and now subject to staged application [2], represents a landmark attempt to answer the first order of

*Correspondence: Achi Iseko (a.iseko@yahoo.com)

Received: 6 October 2026; **Accepted:** 16 October 2026; **Published:** 31 October 2026

Copyright: © The Author(s), 2026. Published by Science Publishing Group. This is an **Open Access** article, distributed under the terms of the Creative Commons Attribution 4.0 License (http://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution and reproduction in any medium, provided the original work is properly cited.

questions comprehensively. It classifies AI systems by risk, mandates conformity assessments, establishes transparency obligations for providers, and creates an enforcement architecture capable of restricting market access and imposing substantial financial penalties on non-compliant actors. It is, by any measure, a significant regulatory achievement.

Yet the Act makes a foundational choice that this paper examines and contests. It governs artificial intelligence as a product deployed in specific contexts, not as an output of a labour process involving specific workers. The distinction matters because the two frames produce entirely different regulatory subjects. A product-risk framework asks: what does this system do to people, and is it safe, accurate and transparent enough? A labour process framework asks: who made this system, under what conditions, and do those conditions affect what the system is? The EU AI Act is principally organised around the former, with limited direct engagement with the latter. In doing so, it makes workers visible principally as persons affected by AI deployment decisions, not as the distributed workforce whose labour, judgement and exposure produce AI systems in the first place.

This paper examines that structural choice through a systematic regulatory audit using the Ethical Visibility Framework (EVF), an analytical instrument developed to identify how labour invisibility operates across material, epistemic and narrative dimensions in AI production systems [3]. EVF was constructed empirically through 80 semi-structured interviews with data annotators, content moderators and AI ethics practitioners and policy consultants across eight countries in Sub-Saharan Africa and the Arab Gulf. Those workers constitute part of the transnational production workforce whose labour underpins the AI systems now subject to EU regulation. Applying EVF as a regulatory audit tool, we pose the following research question: to what extent and in what ways does the EU AI Act recognise the labour relations and worker knowledge involved in AI production?

The present article draws on that published framework and its reported findings as secondary evidence; it does not collect new interviews or reanalyse the earlier study's transcripts.

Our answer, developed through systematic analysis of the Act's operative articles, recitals, annexes and current implementing instruments, is that the Act weakly recognises AI labour across all three EVF dimensions. This weak recognition is not uniform negligence: Article 10 comes closer to the production process than any other provision, and Article 25 acknowledges that responsibilities exist across an AI value chain. But neither provision constitutes a labour supply-chain transparency regime, and neither establishes a general obligation concerning the conditions of the workers whose knowledge and exposure make AI systems possible. The gap is structural rather than incidental, rooted in a product-risk regulatory logic that cannot easily accommodate labour governance without targeted amendment.

The paper makes four contributions. First, it extends EVF from an empirical diagnostic tool to a normative instrument for regulatory analysis, setting out an explicit interpretive audit method applicable to other AI governance frameworks. Second, it presents a provision-level account of what the EU AI Act does and does not say about AI labour, mapping specific provisions against explicit EVF indicators. Third, it identifies and critically examines the Act's implicit rival logic, its reliance on existing EU labour law as the appropriate instrument for worker protection, and demonstrates why this deferral produces a specific governance gap for transnational AI production work. Fourth, it proposes targeted amendments that could close this gap within the Act's existing architecture, without requiring revision of its foundational risk-based approach.

The paper proceeds as follows. Section 2 provides a structured overview of the Act and situates our contribution within the existing literature. Section 3 presents the audit methodology, including the EVF indicators and corpus definition. Section 4 reports the audit findings across all three EVF dimensions. Section 5 examines the rival explanation and the enforcement paradox. Section 6 proposes amendments. Section 7 concludes.

2. Background: The EU AI Act and the Question of Labour

2.1. Architecture and Scope of the Act

The EU AI Act entered into force on 1 August 2024 following a legislative process initiated by the European Commission's proposal in April 2021. Its architecture rests on a risk-based classification of AI systems often described through four risk categories, alongside a distinct general-purpose model regime. Unacceptable risk practices are prohibited under Article 5, including social scoring, certain real-time remote biometric identification uses in publicly accessible spaces and certain emotion recognition applications, subject to the conditions and exceptions specified in that article. High-risk systems, identified through specified product-related rules and Annex III domains including employment and workers' management, critical infrastructure, law enforcement and migration, are subject to the Act's most demanding requirements. Limited-risk systems face transparency obligations. Systems outside these categories are not subject to the high-risk requirements solely because they use AI; horizontal and other applicable legal duties may still matter [2].

For high-risk AI systems, providers bear a set of interlocking pre-market obligations. Article 9 mandates a continuous risk management system. Article 10 requires that training, validation and testing datasets meet quality criteria. Article 11 requires technical documentation. Article 13 mandates transparency toward deployers. Article 14 requires that systems be designed to allow effective human oversight by natural persons. Article 15 addresses accuracy, robustness and cyberse-

curity. Article 25 establishes that responsibilities may be distributed across actors along the AI value chain, including importers and distributors. These requirements are enforced through conformity assessments, European conformity (CE) marking, market surveillance authorities, and penalties reaching €35 million or 7% of global annual turnover for specified serious violations, with lower maxima and different routes for other infringements. Application remains subject to the statutory timetable [2].

General-purpose artificial intelligence (GPAI) models are subject to a separate regime under Chapter V. Article 53 obliges GPAI providers to maintain technical documentation, make documentation available to downstream system providers, and publish a summary of training content under Article 53(1)(d). These obligations began applying in August 2025, subject to transitional arrangements. The voluntary Code of Practice, published in July 2025 and assessed by the Commission and AI Board, provides a compliance pathway. The Commission reports that the AI Omnibus entered into force on 27 July 2026, with Annex III high-risk rules applying from 2 December 2027 and relevant product-embedded high-risk rules from 2 August 2028. The timetable should therefore be distinguished from the design of the obligations discussed here [4, 5].

2.2. What the Act Governs and What It Assumes

Beneath the Act's detailed provisions lies a set of foundational assumptions about where AI risk resides. The Act principally frames AI risk through systems, models and their effects on affected persons. Providers place systems on the market; deployers put them into service; affected persons experience their outputs. This downstream orientation is explicit in the Act's definitions, which specify economic operators but do not define annotators, moderators or reinforcement learning from human feedback (RLHF) workers as a distinct production workforce. Article 10 (2) (c) nevertheless expressly recognises annotation and labelling as data-preparation activities [2].

Article 10, the Act's most substantive engagement with AI production, governs data governance and management practices for training, validation and testing datasets. Its requirements concern data collection processes, design choices, preparation, bias detection and mitigation. These requirements address the data artefact, asking whether it is representative, accurate and appropriate for its intended purpose. They do not establish a general labour-welfare duty concerning the conditions under which human data work is performed, the contractual arrangements through which annotators are engaged, the geographic location of that work, or the pay, hours and psychological conditions under which labelling, moderation and evaluation tasks are completed. The worker is present in the production process that Article 10 governs; the worker's conditions are not independently specified as a general compliance requirement.

Article 25, which distributes responsibilities along the AI value chain, similarly operates within an economic and accountability logic rather than a labour supply-chain transparency logic. It determines which actor bears which conformity obligation when a provider relies on a third-party component or model. It does not require disclosure of the human labour arrangements that constitute that value chain at the production level.

Recital 9 makes the Act's labour scope explicit. It states that the Act "should not affect Union law on social policy and national labour law" and "should not affect the exercise of fundamental rights as recognised in the Member States," including the right to take collective action. This recital preserves existing labour rights and presents the Act as complementary to them; it is not an operative exclusion of every labour-related concern. Section 5 examines whether that deferral is adequate. Article 2 (11) also preserves more favourable worker protections concerning employers' use of AI. The issue is whether these complementary routes adequately connect AI production conditions to responsibility and remedy [2].

2.3. Existing Critiques and the Labour Gap

The Act has generated a substantial critical literature across several distinct lines. Earlier work on ethical AI and algorithmic accountability questioned whether technical and organisational assurance adequately addresses power and inequality. Labour-law scholarship has also directly examined the regulation of workplace AI. Cefaliello and Kullmann, for example, identify risks to workers' fundamental rights in the draft Act [6]. Their analysis predates the adopted instrument and is used here as a conceptual interlocutor rather than as an account of current obligations. For GPAI, the statutory training-content summary and the Commission's template provide transparency about training sources without constituting a labour supply-chain disclosure regime [7].

A separate body of literature, the AI labour studies tradition, has documented the conditions of the workforce that produces AI systems. This literature has established the structural precarity of platform-based annotation and moderation work [8, 9], the epistemic dynamics through which workers' domain knowledge is harvested while their status as knowledge contributors is denied [3, 10, 11], and the opportunities and constraints of cross-border digital work in selected Global South settings [12, 13]. These sources do not establish the global geographic distribution of the entire AI production workforce.

These two literatures offer a basis for closer engagement on the specific question of production labour. Much AI governance scholarship examines what the Act does and does not regulate within the deployment chain. AI labour scholars document what the Act does not see, the production workforce and its conditions. This paper bridges that gap through a systematic audit that brings EVF, developed from the AI labour

tradition, to bear on the Act's regulatory architecture. Our contribution is not to add another critique of the Act's deployment-side provisions, but to examine what a framework built around worker testimony reveals about the Act's treatment of those workers.

This matters for more than completeness. The Act's foundational claim is that it governs "trustworthy AI." If the conditions under which AI systems are produced are materially unjust, epistemically extractive and narratively invisible, the trustworthiness of those systems is compromised at its source. A product-risk regime should therefore be explicit about what it verifies concerning production and which questions remain for labour-governance institutions.

3. Methodology: EVF as a Regulatory Audit Instrument

3.1. Analytical Approach

This paper employs a systematic regulatory audit methodology, combining critical legal analysis with the interpretive tools of the Ethical Visibility Framework. Regulatory audit as a method involves the systematic application of an established analytical framework to a defined legal corpus, with the aim of identifying what a regulatory instrument addresses, what it partially recognises, and what it structurally excludes [14, 15]. Its emphasis differs from an exclusively internal doctrinal analysis in that its evaluative criteria are drawn from outside the legal text itself, in this case from a framework developed through empirical research with the workers whose conditions the text fails to address. This positioning is deliberate: we audit the Act not on its own terms but against the material realities documented in the AI labour literature and grounded in the EVF's empirical foundations.

The methodological choice reflects a broader argument about regulatory knowledge. One question for legal analysis is to evaluate a regulation by asking whether it achieves its stated objectives. We ask a prior question: whose experiences and conditions shaped the objectives in the first place, and whose were excluded. Applying a worker-derived analytical framework to a regulator-produced text is not an external imposition but an epistemically justified move: if the Act claims to govern trustworthy AI, the conditions of those who produce it are directly relevant to that claim, regardless of whether the Act's drafters recognised them as such.

3.2. Corpus Definition

The audit corpus comprises three overlapping layers of material, selected to capture both the Act as adopted and the implementing context that determines how its provisions operate in practice.

The primary corpus consists of the operative text of Regulation EU 2024/1689 as published in the Official Journal on 12 July 2024, read alongside the Commission AI Act Service Desk presentation of the consolidated text dated 27 July 2026. For the purposes of the audit, particular attention was paid to the provisions governing high-risk AI system requirements (Articles 8 to 27), GPAI obligations (Articles 51 to 56), the definitional framework (Article 3), and scope and labour-law interactions (Article 2 and Recital 9).

The secondary corpus consists of current implementing instruments, comprising the GPAI Code of Practice published in July 2025 and assessed by the Commission and AI Board; the European Commission's Explanatory Notice and Template for the Public Summary of Training Content under Article 53 (1)(d), adopted 24 July 2025; and the application timetable following the 2026 AI Omnibus as confirmed by the European Commission's current overview. These materials help clarify how the Act's abstract provisions operate in practice and were included because any audit of a regulation in force must account for the gap between text as adopted and implementation context [4, 5, 7].

The tertiary corpus consists of the Annex III domain classifications and the conformity assessment procedures under Articles 43 to 49, which determine which systems are subject to mandatory requirements and on what basis. These materials were analysed to establish the Act's definitional boundaries and the actors it recognises as regulatory subjects.

Materials excluded from the corpus include national implementation plans, which vary across Member States and fall outside the Act's uniform text; a comprehensive review of enforcement decisions; and parliamentary preparatory materials, which inform legislative intent but do not constitute binding legal text.

3.3. The Ethical Visibility Framework: Indicators for Regulatory Audit

EVF was originally developed as an empirical diagnostic tool for analysing labour invisibility in AI production systems, grounded in 80 semi-structured interviews across Sub-Saharan Africa and the Arab Gulf [3]. In its original application, EVF identified three mutually reinforcing dimensions through which workers are rendered invisible: material structures that obscure labour through supply chain fragmentation and contractual displacement; epistemic hierarchies that deny workers recognition as knowledge producers; and narrative framings that discursively erase human contributions through automation mythology and product-centred language.

For regulatory audit purposes, each dimension requires operationalisation as a set of explicit indicators that can be applied to legal text. Table 1 sets out the full indicator set used in this audit. Each indicator represents a question the Act either answers, partially answers, or does not address. The coding frame assigns one of three values: presence (P), meaning the Act's provisions directly and substantively address the indicator; partial recognition (PR), meaning the Act touches on the indicator but incompletely or indirectly; and absence (A),

meaning that no dedicated provision answering the specified production-labour question was identified. A does not mean absence of all protection under EU or national law. PR also includes general participation and complaint routes that can include production workers without conferring production-specific entitlements.

Coding is a single-author interpretive exercise, not a validated numerical measure of legal effectiveness. Each judgement considers the provision's duty holder, beneficiary and connection to production work, alongside counterexamples such as Articles 10 (2)(c), 26 (7), 67 and 85. Table 2 records the legal basis and residual gap so that these judgements can be assessed independently.

Table 1. EVF Regulatory Audit Indicators.

Dimension	Indicator	Code
Material	Supply chain disclosure: Are providers required to disclose the human labour arrangements in their AI production supply chains?	A
Material	Worker classification: Does the Act define or address the employment status of AI production workers?	A
Material	Geographic transparency: Are providers required to disclose the geographic distribution of their production workforce?	A
Material	Pay and conditions: Does the Act require disclosure of, or set standards for, pay or working conditions in AI production?	A
Material	Intermediary accountability: Does the Act address the role of subcontractors and intermediaries in AI production?	PR
Epistemic	Worker knowledge: Does the Act recognise AI production workers as contributors of knowledge or expertise?	A
Epistemic	Worker participation: Does the Act require or enable worker participation in AI governance or standards development?	PR
Epistemic	Human oversight by workers: Does the Act's human oversight framework extend to oversight exercised by production workers?	PR
Epistemic	Expertise attribution: Are annotation, moderation or evaluation workers recognised as skilled contributors in technical documentation?	A
Narrative	Production vocabulary: Does the Act name the production workforce or recognise its data-production activities?	PR
Narrative	Labour framing: Does the Act frame AI systems as outputs of a labour process, or exclusively as technical products?	A
Narrative	Worker as subject: Can AI production workers appear as persons with procedural rights or interests under the Act?	PR

3.4. The Rival Explanation

A rigorous audit must address the strongest counter-argument to its central claim. In this case, the rival explanation is that the Act's silence on AI production labour is not a governance gap but a rational division of regulatory labour: AI labour conditions are properly regulated by existing EU instruments, principally the Platform Work Directive (EU 2024/2831), the Information and Consultation Directive (2002/14/EC), the General Data Protection Regulation (GDPR), and national employment law transposing EU minimums. On this reading, the Act is not failing to regulate AI labour; it is correctly leaving that regulation to the instruments designed for it, in keeping with the principle of subsidiarity and the integrity of EU employment law as a coherent domain.

We take this rival explanation seriously and engage it directly in Section 5. Here we note three preliminary observations that shape the audit. First, general employment and safety rules can address AI work without naming each task, but their application must be established for the relevant relationship. Second, workers in transnational production chains may be located outside the territorial scope of particular EU labour instruments; this cannot be inferred from nationality or geography alone. Third, the Act promotes trustworthy AI, while its conformity requirements do not amount to a comprehensive fair-work certification. The question is whether the combined framework provides an effective connection between production conditions, responsible entities and remedies.

3.5. Scope and Limitations

This audit examines the Act as a legal text and implementing instrument. It does not assess enforcement outcomes, nor does it make predictions about how national competent authorities will interpret ambiguous provisions. The audit's findings describe the Act's regulatory architecture as designed, not as it may evolve through enforcement practice or judicial interpretation. We note also that the Act is a living instrument: Article 7 empowers the Commission to amend Annex III, and the GPAI Code of Practice is subject to revision. Our findings concern the Act in its current form as checked on 15 October 2026.

The audit is not a comprehensive legal commentary on the Act. Our focus is precisely delimited to the EVF indicator set. Provisions that fall outside this indicator set, such as the Act's conformity assessment procedures, its sandbox provisions, or its enforcement architecture, are discussed only to the extent that they bear on the labour visibility question.

Selected adjacent EU instruments and institutional reports are used comparatively, rather than treated as additional AI Act provisions. Regional sources differ in population, method and task coverage and are not pooled into a prevalence estimate. The analysis cannot establish actual compliance, enforcement outcomes or the effectiveness of the proposed amendments.

4. Findings: An EVF Audit of the EU AI Act

This section presents the findings of the regulatory audit, organised by EVF dimension. Each sub-section examines the relevant Act provisions against the indicator set established in Table 1, assigns a presence, partial recognition or absence coding, and explains the reasoning. Table 2 at the close of the section summarises the full coding. Throughout, we distinguish between what the Act directly addresses, what it approaches but does not reach, and what remains outside the dedicated duties examined here.

4.1. Material Dimension: Governing the Artefact, Not the Labour

4.1.1. Supply Chain Disclosure (A)

The Act does not require providers of high-risk AI systems or GPAI models to disclose the human labour arrangements in their AI production supply chains. Article 25, which is the Act's most direct engagement with supply chain responsibilities, distributes conformity obligations across providers, importers, distributors and other actors in the AI value chain. Its logic concerns allocation of AI compliance responsibilities: it determines which actor bears which regulatory obligation when a provider integrates a third-party component or relies on an upstream model. It does not require any actor to disclose the identity, location, pay or working conditions of the human workers who produced training data, performed annotation, or conducted reinforcement learning evaluations at any point in that chain. These omissions do not establish that a particular provider satisfies Article 25 or other applicable law; they identify information that the article does not generally require as labour disclosure [2].

Article 53 (1) (d) requires GPAI providers to publish a summary of training content, and the Commission's implementing template of July 2025 supports disclosure of training sources and processing information. This is a content transparency requirement directed at rightsholders and the public, not a labour supply chain disclosure requirement. The summary template contains no field for the human labour arrangements through which training content was produced, annotated or curated. The worker who labelled the data is not the subject of a dedicated labour-disclosure field in that summary [7].

4.1.2. Worker Classification (A)

The Act does not define or address the employment status of AI production workers. Its definitional article, Article 3, defines provider, deployer, authorised representative, importer, distributor and product manufacturer. It does not define annotator, content moderator, data labeller, RLHF contractor or any equivalent term. The production workforce is not defined as a distinct occupational category. This matters for the specificity of obligations, but does not remove general procedural rights: Article 85 permits complaints by natural or legal persons concerning alleged infringements. Section 4.3.3 distinguishes this route from a production-labour entitlement [2].

4.1.3. Geographic Transparency (A)

No provision of the Act requires providers to disclose the geographic distribution of their production workforce. Article 10 requires documentation of data collection processes and design choices, and the GPAI technical documentation template under Annex XI requires disclosure of training data sources and methods. Neither instrument requires identification of the countries in which annotation, moderation or evaluation work was performed, the proportion of work contracted to third-party platforms in specific jurisdictions, or the regulatory context governing workers in those jurisdictions. The location of the workforce is therefore distinct from data provenance and the geographic context of intended use. The absence classification concerns workforce location, not every geographic feature of data governance.

4.1.4. Pay and Working Conditions (A)

The Act sets no standards for, and requires no disclosure of, pay or working conditions in AI production. This indicator codes as a straightforward absence. Recital 9 preserves rather than excludes applicable labour law. Even where the Act's provisions come close to the production process, as in Article 10's data governance requirements, they do not establish an

independent, general standard for production workers' pay and occupational conditions [2].

4.1.5. Intermediary Accountability (PR)

This is the indicator on which the Act comes closest to the material dimension of AI labour visibility, and it merits close analysis. Article 25 addresses circumstances in which a distributor, importer, deployer or other third party becomes the provider of a high-risk system; it also requires specified written arrangements for supplier information and assistance needed for compliance. This provision does hold providers accountable for the regulatory compliance of systems they did not entirely build themselves. In that limited sense, it acknowledges that AI systems are produced through distributed arrangements involving multiple actors.

However, Article 25's accountability logic concerns product conformity, not labour conditions. The question it answers is: if this system is non-compliant, which actor is liable? It does not ask: if the workers who contributed to this system's production were exploited, who bears responsibility? The intermediary chain it governs is a chain of technical and legal accountability for system outputs, not a labour supply chain in which worker rights, pay and conditions are tracked and disclosed. We code this as partial recognition because Article 25 acknowledges distributed production without constituting a labour supply chain transparency regime.

4.2. Epistemic Dimension: Human Oversight Without Worker Voice

4.2.1. Worker Knowledge (A)

The Act does not expressly recognise AI production workers as contributors of knowledge or expertise. The Act's epistemic framework is organised around a different set of human actors: providers who possess technical expertise about their systems, deployers who must be informed enough to oversee system outputs, and affected persons whose interests must be protected. Annotation workers, content moderators and RLHF evaluators, whose domain knowledge, cultural competence and contextual judgement directly shape what AI systems learn to do, have no dedicated entitlement to epistemic recognition in this framework. The Act's transparency provisions, including the technical documentation requirements of Article 11 and the transparency obligations of Article 13, concern what providers must disclose about their systems. Article 10 (2) (c) nevertheless recognises annotation and labelling. The A classification concerns recognition of workers as knowledge contributors, rather than absence of every reference to their activities [2].

This absence is particularly significant given what the AI labour literature has established: that annotation is a skilled epistemic practice, that the quality of AI system outputs is directly dependent on the judgement of those who label training data, and that systematic denial of workers' epistemic status is a mechanism of economic extraction, not simply a gap in recognition [10, 11]. The distinction between recognising an activity and recognising its practitioners is therefore consequential.

4.2.2. Worker Participation (PR)

The Act establishes several mechanisms for stakeholder participation in AI governance. The AI Board, established under Article 65, includes representatives of Member States. The Advisory Forum under Article 67 includes industry, civil society and academia, and Recital 150 expressly includes social partners. The scientific panel under Article 68 brings in independent experts. Article 56 permits relevant stakeholders to support GPAI code development, while Recital 165 expressly envisages trade-union participation in voluntary inclusive-development approaches. Article 26 (7) also requires employer-deployers to inform workers' representatives and affected workers before workplace use of a high-risk system. These provisions justify partial recognition rather than absence. They do not guarantee dedicated representation for transnational production workers, but the text alone cannot establish their actual exclusion from every forum or standardisation process [2].

4.2.3. Human Oversight by Workers (PR)

Article 14 is the Act's central provision on human oversight, and it warrants careful reading. Article 14 (1) requires that high-risk AI systems "be designed and developed in such a way that they can be effectively overseen by natural persons during the period in which the AI system is in use." Article 14 (4) specifies that human oversight measures must enable persons to "correctly interpret the high-risk AI system's output" and to "decide, in any particular situation, to refrain from using the high-risk AI system."

The natural persons contemplated by Article 14 include staff assigned oversight by deployers: the doctor who reviews a diagnostic recommendation, the credit officer who assesses an algorithmic credit score, the border official who evaluates a biometric match. The oversight the Act mandates is oversight of system outputs in deployment contexts, exercised by those who use the system, not those who built it. Annotation workers, content moderators and RLHF evaluators, who exercise a form of continuous oversight over the system's learning process during production, do not gain a general production-phase oversight right under Article 14 merely by performing these tasks. We code this as partial recognition not because the Act recognises production-phase human oversight but because Article 14 does acknowledge the general principle that human judgement must be available to assess and correct AI system outputs. That principle, applied to the production phase rather than the deployment phase, would constitute exactly the kind of worker oversight the Act does not require. The gap between principle and application is the analytical finding.

4.2.4. Expertise Attribution (A)

Technical documentation under Article 11 and Annex IV must include information about the AI system's design, development, testing and performance. It must describe the data used for training and the measures taken to examine and address biases. It does not require attribution of, or reference to, the human workers whose knowledge produced that data or whose evaluations shaped the system's behaviour. For GPAI models, the separate documentation route is Article 53 and Annexes XI and XII; it likewise does not establish a dedicated occupational-attribution requirement. The expertise of annotation workers is legally invisible in the documents that regulatory authorities use to assess system compliance.

4.3. Narrative Dimension: The Vocabulary of Invisible Production

4.3.1. Production Vocabulary (PR)

The Act's operative vocabulary does not define annotators, content moderators or RLHF workers as a distinct production workforce. However, Article 10 (2)(c) expressly names annotation and labelling among data-preparation operations. The finding is therefore partial narrative recognition: production activity is visible, while its practitioners are not assigned a corresponding occupational identity or production-specific entitlement. Definitions are relevant to this visibility, but procedural rights can also apply through general categories, as Article 85 demonstrates. The claim is not that every term related to data work is absent, but that recognition of the task does not itself establish recognition of the worker [2].

4.3.2. Labour Framing (A)

The Act principally frames AI systems as technical products moving through a market chain from provider to deployer to affected person. It does not frame AI systems as outputs of a labour process in which workers contribute knowledge, bear costs and exercise judgement. This framing is visible in the Act's risk classification logic, which identifies risk by the domain in which a system is deployed, not by the conditions under which it was produced; in its conformity assessment procedures, which evaluate system performance and documentation, not supply chain labour conditions; and in its transparency obligations, which concern what systems do and how they work, not who built them and under what arrangements. An independent framework for production-labour conditions is absent from the duties examined here; adding one would require a substantive decision about the scope of compliance.

4.3.3. Worker as Subject (PR)

AI production workers can have procedural interests under the Act without being defined as a distinct occupation. Article 85 allows natural or legal persons with grounds to consider that an infringement has occurred to complain to the relevant authority. Production workers may also be affected by workplace AI and benefit from applicable protections in that capacity. This justifies partial recognition. The remaining gap concerns the substantive basis of a grievance: unpaid remuneration or harmful content exposure does not by itself become an AI Act infringement solely because the work contributed to an AI system. Other law may provide remedies, and the same facts may reveal an applicable AI compliance breach. Standing to complain and a right to remedy production-labour conditions are different questions [2].

4.4. Summary of Audit Findings

Table 2 presents the complete coding of EVF indicators against the Act's provisions.

Table 2. *EVF Regulatory Audit Results.*

Dimension	Indicator	Provisions and recitals	Code	Recognition and residual gap
Material	Supply chain disclosure	Arts 11, 25, 53; Annexes IV, XI, XII; Recital 107	A	Technical and training-content transparency do not require a general labour supply-chain statement.
Material	Worker classification	Art 3; Art 2 (11); Recital 9	A	No dedicated determination of production workers' employment status; applicable labour law remains relevant.
Material	Geographic transparency	Art 10 (2)(b), (4); Annexes IV, XI	A	Data provenance and intended-use context differ from workforce location.
Material	Pay and conditions	Arts 9–11; Recital 9	A	No independent general duty on production pay or welfare; complementary labour law is preserved.
Material	Intermediary accountability	Art 25 (1)–(4); Recitals 83–90	PR	Supplier responsibility and cooperation concern AI compliance rather than a general labour remedy.
Epistemic	Worker knowledge	Arts 10 (2) (c), 11, 13; Recital 67	A	Annotation is named, but workers are not expressly recognised as knowledge contributors with an entitlement to that recognition.

Dimension	Indicator	Provisions and recitals	Code	Recognition and residual gap
Epistemic	Worker participation	Arts 26 (7), 56 (3), 67; Recitals 150, 165	PR	Stakeholder and social-partner routes exist without dedicated transnational production-worker representation.
Epistemic	Human oversight by workers	Arts 14, 26 (2); Recital 73	PR	Human judgement is recognised during use; a general production-phase authority is not established.
Epistemic	Expertise attribution	Art 11; Art 53; Annexes IV, XI, XII	A	Documentation does not require occupational or collective attribution to production workers.
Narrative	Production vocabulary	Art 10 (2)(c); Recital 67	PR	Annotation and labelling are named activities without a distinct production-workforce category.
Narrative	Labour framing	Arts 1, 9–11, 25, 53; Recitals 1, 9	A	No independent general production-labour compliance framework is identified.
Narrative	Worker as subject	Arts 2 (1)(g), 26 (7), 85; Recital 170	PR	General complaint routes can include production workers; the grievance still requires an applicable substantive duty.

Seven of twelve indicators code as absent and five as partial recognition. Partial recognition covers intermediary accountability, participation, human oversight, production vocabulary and general procedural interests. No indicator codes as present. The pattern is selective visibility rather than complete absence; the classifications concern the specified indicators and are not a quantitative measure of legal effectiveness.

The partial recognitions are analytically important precisely because they reveal where the Act comes closest to labour visibility without achieving it. Article 25 acknowledges distributed production without requiring labour supply chain transparency. Article 14 acknowledges the necessity of human judgement without recognising the human judgement embedded in the production process. These near-misses are not coincidental: they mark the boundary of the product-risk frame. Within that frame, the Act acknowledges distributed accountability, human judgement and general participation and complaint routes. These connections do not add up to an independent, general duty concerning the working conditions under which systems and models are produced.

5. Discussion: The Rival Explanation and the Enforcement Paradox

5.1. The Rival Explanation: Intentional Division of Regulatory Labour

The most credible challenge to the findings presented in Section 4 is not that the audit is wrong but that it misidentifies a gap where there is in fact a deliberate design choice. On this reading, the Act's silence on AI production labour is not negligence but rational regulatory architecture: AI labour conditions are properly governed by existing EU instruments, and duplicating those protections within the AI Act would produce regulatory overlap, legal uncertainty and unnecessary compliance burdens on providers. Recital 9 makes this logic explicit, preserving social policy and national labour law as complementary protections. The Act governs AI systems; employment law governs workers; each instrument does its job.

This argument deserves to be taken seriously because it reflects a genuine principle of EU regulatory design. The Platform Work Directive, which entered into force in 2024 and has a transposition deadline of 2 December 2026, addresses algorithmic management of platform workers. The Information and Consultation Directive provides for worker representation in undertakings above a certain size. The General Data Protection Regulation governs the use of personal data, including data generated by workers. The Working Time Directive, the Temporary Agency Work Directive, and national minimum wage legislation provide a floor of protection for workers in the EU. If these instruments collectively govern the conditions of AI production workers, a division of regulatory functions may be appropriate [16-20].

We engage this rival explanation across three dimensions: jurisdictional reach, substantive adequacy, and regulatory interaction.

5.2. Jurisdictional Reach: Where the Gap is Largest

The rival explanation faces an important qualification at the level of jurisdiction. The EU labour instruments identified above require the relevant work, employer, undertaking or other connecting factor to fall within their scope. Work performed in Sub-Saharan Africa, the Arab Gulf, Southeast Asia or Latin America is not automatically covered simply because its output enters an EU-regulated AI product. The published EVF study included participants across Kenya, Nigeria, South Africa, Ghana, Tanzania, the United Arab Emirates (UAE), Qatar and Bahrain [3]. This does not establish that every par-

ticipant lacked every EU remedy, or that local labour law offered no protection. Those questions require analysis of the actual work, corporate relationships, jurisdiction and applicable law. The global proportion of the AI production workforce in these regions is not established by that qualitative sample.

The EU AI Act, by contrast, includes market and output-based connections beyond EU establishment. Article 2 (1) (c) makes clear that the Act applies to providers and deployers located outside the EU when the output of their AI system is used within the EU. Article 2 (1) (a) separately covers providers placing AI systems or GPAI models on the Union market irrespective of establishment. In an illustrative California–Nairobi–Munich chain, the applicable duties depend on whether the entity supplies a model or a system and on the relevant classification; GPAI models do not automatically follow the high-risk system conformity route. The same Act that reaches into California to regulate the system's deployment does not reach into Nairobi to ask about the conditions under which it was produced. The jurisdictional asymmetry is precise and consequential: extraterritorial reach for product governance, no automatic transmission of employment entitlements to every contributing worker [2].

This jurisdictional gap is not addressable by pointing to existing EU labour instruments. It is a gap that a production-side provision of the Act, coordinated labour enforcement or other supply-chain measures, such as a supply chain disclosure requirement applied to any provider placing systems on the EU market, could close. Supply chain disclosure requirements in the EU context, most notably the Corporate Sustainability Due Diligence Directive (CSDDD) adopted in 2024, illustrate legislative approaches to supply-chain responsibility beyond EU borders, although their scope, connecting factors and implementation differ from the AI Act. The CSDDD must be read with its subsequent amendments, including Directives (EU) 2025/794 and (EU) 2026/470; its narrowed coverage and phased timetable preclude treating it as a universal current protection [21].

5.3. Substantive Adequacy: Platform Work and Its Limits

Even where EU labour instruments have jurisdictional reach, their substantive fit with AI production work requires examination. The Platform Work Directive is particularly relevant where tasks are organised through a digital labour platform. It addresses correct employment-status determination and algorithmic management, including specified protections for persons performing platform work without an employment relationship. An independent-contractor label therefore does not automatically remove a person from all protection. The Directive concerns platform work performed in the Union, and adoption must be distinguished from completion of national transposition. Outsourcing through a business-process firm may instead engage ordinary employment and other applicable rules [16].

Furthermore, protections against algorithmic management and protections concerning production conditions overlap without being identical. Occupational safety and health rules, including the preventive framework of Directive 89/391/EEC, can address psychosocial risks such as harmful-content exposure where their scope and national implementation apply [22]. Subcontracting does not itself remove employer duties; it may complicate the identification of responsible entities and access to evidence. GDPR can protect workers' personal data where its material and territorial conditions are met, including processing in the context of an EU establishment [18]. The remaining concern is whether these separate routes can be connected effectively to the provider's production chain and to a remedy for the worker.

5.4. Regulatory Interaction: Trustworthiness Without Production Standards

The deepest problem with the rival explanation concerns the relationship between the Act's claim and its scope. The Act does not claim merely to regulate the deployment of AI systems. It claims to govern trustworthy AI: Article 9's risk management requirements, Article 10's data quality standards, and the Act's overarching framing in Recital 1 contribute to the objective that AI systems placed on the EU market meet standards that justify public trust. The CE marking that a compliant high-risk AI system may bear indicates conformity with applicable requirements; it is not a comprehensive ethical or fair-work certification.

The question this claim raises is what trustworthiness requires at the production level. If AI systems are produced through labour conditions that are materially exploitative, epistemically extractive and narratively invisible, does the product's conformity with technical and deployment requirements make it trustworthy in any meaningful sense? We argue it does not, and that the Act's failure to engage production conditions therefore undermines the trustworthiness claim it makes. The concern is a possible gap between wider claims of trustworthy AI and what the compliance process actually verifies about production labour. Trustworthiness as a regulatory standard cannot be limited to the product if the product's characteristics are shaped by the conditions of its production.

5.5. The Enforcement Paradox

The rival explanation, even where it partially holds, cannot address what we term the enforcement paradox. The Act's enforcement powers apply through a differentiated timetable and allocation of responsibilities [4, 23]. National competent authorities can investigate non-compliant providers. The AI Office can act against GPAI providers in breach of Chapter V obligations. Penalties reach €35 million or 7% of global annual turnover for the most serious violations, and market access can be restricted for non-compliant systems. These are significant powers. They represent a regulatory capacity that

most previous AI governance initiatives entirely lacked.

These powers do not establish an independent, general duty governing the labour conditions of AI production. Consider an illustrative case in which a contractor fails to pay for completed annotation tasks. The workers may have claims under applicable labour or contract law; the facts may also reveal an AI Act infringement, for example inaccurate compliance documentation or deficient data governance. However, unpaid remuneration alone is not a general production-labour infringement under the Act. Compliance with Article 10 does not establish compliance with every other obligation. The enforcement paradox is therefore a mismatch between substantial compliance powers and the absence of a dedicated production-labour duty, rather than an assertion that labour evidence can never trigger action.

This paradox is not a defect in the enforcement architecture. It is a direct consequence of the product-risk frame. Within that frame, non-compliance means a system that is inaccurate, non-transparent, or inadequately overseen in deployment. Exploitative production conditions do not independently establish an AI Act infringement unless they engage an applicable duty. Article 85 remains available for complaints concerning such infringements. Closing the enforcement paradox requires not merely adding a provision but introducing a new regulatory logic into the Act: one that treats production conditions as a dimension of product compliance rather than as an external matter properly left to labour law.

5.6. The Global South Dimension

The enforcement paradox warrants particular attention in transnational production chains involving the Global South. Published digital-labour studies in selected Sub-Saharan African and Asian settings document bargaining constraints, intermediary relationships and the coexistence of autonomy with algorithmic control [12, 13]. The earlier EVF study provides bounded qualitative evidence concerning production workers in five African and three Gulf countries [3]. These are secondary sources for the present article: no new interviews were conducted and no transcripts were reanalysed. Their findings should not be extended into claims of uniform regional wages, universal absence of support or the global majority of AI workers.

The International Labour Organization (ILO) global platform-work report and the World Bank's Working Without Borders provide wider evidence about earnings, working time and social protection in digital work [24, 25]. Fairwork's 2025 cloudwork assessment adds an independent account of selected platforms against pay, conditions, contracts, management and representation criteria [26]. These sources cover mixed kinds of work; their aggregate figures are not counts of AI annotators. For the Arab Gulf, the ILO's report on digital and AI transformation in Arab labour markets provides comparative policy context [27], but occupational projections do not establish annotation-work incidence or Gulf-wide labour conditions. Evidence for the two regions is therefore uneven, and the Gulf discussion should be read as a bounded comparison rather than equivalent regional coverage.

These conditions are not merely unjust in themselves. They may have implications for the quality and character of the AI systems those workers produce. Training data produced under conditions of time pressure and inadequate psychological support may reflect those conditions in ways that affect system outputs. Annotation guidelines interpreted without cultural context or adequate compensation for expertise may produce mislabelled data that propagates errors through the system. The relationship between production conditions and product quality is a plausible technical as well as ethical concern, but this documentary audit does not establish causation for a particular system. Yet Article 10 reaches the data artefact without reaching the conditions of its production, treating the output as separable from the labour process that generated it.

The Global South dimension also exposes the limits of the Act's jurisdictional logic in a different register. The Act's extraterritorial reach rests on a market access rationale: providers who want access to the EU market must comply with EU product standards. That same rationale could in principle be applied to labour standards: providers who want access to the EU market must disclose and meet minimum standards for the labour conditions in their production supply chains. The Corporate Sustainability Due Diligence Directive provides a qualified precedent for considering such responsibility within EU regulatory architecture. The design of any analogous AI obligation would require an assessment of legal competence, scope, cost and access to remedy, alongside coordination with applicable local law.

6. Towards Labour-Inclusive AI Regulation: Targeted Amendments

6.1. Principles of Amendment

The findings of the EVF audit establish that the Act's product-risk logic produces limited and uneven production-labour visibility across all three EVF dimensions. This section proposes targeted amendments to address that absence. The amendments are designed around four principles drawn from the audit findings and from the rival explanation examined in Section 5.

The first principle is architectural compatibility. The amendments proposed here work within the Act's existing risk-based architecture rather than replacing it. They do not require the Act to become a labour regulation. They require the Act to treat certain labour conditions as dimensions of product compliance, consistent with the Act's own trustworthiness claim and analogous to precedents already established in EU supply chain regulation.

The second principle is jurisdictional reach. Each amend-

ment must be capable of applying to the transnational AI production workforce, not only to workers within EU jurisdiction. This means attaching requirements to providers seeking EU market access, rather than to workers or their employers directly, using the same extraterritorial market access logic the Act already employs for technical compliance.

The third principle is proportionality. The amendments target providers of high-risk AI systems and GPAI models, where the Act's existing requirements are most demanding and where the market access rationale is strongest. They do not impose requirements on providers of minimal or limited risk systems, where proportionality concerns would be legitimate.

The fourth principle is operational realism. Drawing on existing regulatory precedents and current compliance practice, each amendment is designed to be implementable within the Act's conformity assessment and documentation framework while identifying the need for legal authority, supervisory resources and proportionate compliance burdens. Where analogues exist in EU law, those analogues are identified.

6.2. Amendment 1: AI Labour Supply Chain Disclosure as a Conformity Requirement

What the Act currently does: Article 11 requires technical documentation covering system design, development, testing and data governance. Annex IV specifies the content of that documentation. Neither provision requires disclosure of the human labour arrangements through which training data was produced, annotated or curated.

Proposed amendment: a legislative requirement, reflected in Article 11 and Annex IV for high-risk systems and Article 53 and Annexes XI and XII for GPAI models, should require relevant providers to include in their technical documentation an AI Labour Supply Chain Statement. This statement should specify: the categories of human labour used in training data production, annotation and evaluation; the geographic distribution of that workforce at the country level; the contractual basis on which workers were engaged, distinguishing direct employment from platform-mediated contracting; the identity of primary intermediary contractors where platform-mediated arrangements were used; and the pay structures applied, expressed as ranges by task category and geographic region.

Analogue in EU law: The Corporate Sustainability Due Diligence Directive (CSDDD, EU 2024/1760) establishes a framework under which companies within its amended scope are to conduct due diligence on human rights and environmental conditions in their supply chains, including upstream supply chains outside the EU. The proposed AI Labour Supply Chain Statement would extend a similar disclosure logic to AI providers, using the Act's existing technical documentation framework as the vehicle. It does not require providers to guarantee specific labour conditions; it requires them to disclose the conditions that exist. Disclosure is both a transparency measure and an enforcement enabler: national competent authorities and the AI Office cannot investigate conditions they cannot see.

Enforcement: Non-submission of the AI Labour Supply Chain Statement, or submission of a materially incomplete statement, should constitute a documentation deficiency under Article 11, triggering the same corrective action provisions as any other documentation failure. A mandatory labour-disclosure duty would require an adequate legislative basis; an annex amendment cannot be assumed sufficient to create it. Enforcement for GPAI models would need a separate connection to Chapter V obligations.

6.3 Amendment 2: Labour Conditions as a Dimension of Data Quality Under Article 10

What the Act currently does: Article 10 requires that training, validation and testing datasets meet quality criteria including relevance, representativeness, freedom from errors, and completeness. Data governance practices must address data collection processes, design choices, and bias detection. The quality criteria concern the data artefact; they do not address the conditions under which human data work was performed.

Proposed amendment: Article 10 (2) should be amended to add a new sub-paragraph requiring providers to document the conditions of human data production as a component of data governance. Specifically, where training data was produced through human annotation, content moderation, or reinforcement learning feedback, providers must demonstrate that data collection processes include provisions for worker welfare, defined as access to psychological support for workers exposed to harmful content, task completion conditions that do not create systematic time pressure incompatible with accurate annotation, and rates of pay that are disclosed and not below applicable minimum wage standards in the country of engagement.

This amendment does not require providers to pay EU minimum wages to workers in Kenya or Nigeria. It requires them to pay at least the applicable minimum in the jurisdiction of engagement, to disclose what that rate is, and to demonstrate that production conditions do not systematically compromise data quality through worker welfare failures. The quality rationale is integral: systematic time pressure, psychological harm and inadequate compensation for expertise are not only ethical concerns; as Section 5.6 suggests, they are production conditions with potential implications for training data quality and therefore for the system outputs Article 10 is designed to protect.

Worker welfare also matters independently of measurable system performance. A proposed labour duty should not make protection conditional on proving a reduction in model accuracy. Consultation would need to address work without a statutory minimum wage, unpaid waiting and rejected tasks, and the risks of exposing individual workers in public disclosures.

Analogue in EU law: Article 10's existing requirement that

data collection processes be documented and that bias be detected and mitigated already assumes that the human choices involved in data production affect system quality. The amendment extends that assumption consistently: if human choices in annotation affect quality, human conditions in annotation affect quality, and both should be within the governance scope Article 10 establishes.

6.4 Amendment 3: Worker Participation in GPAI Governance

What the Act currently does: Article 67 establishes an Advisory Forum for the AI Board, comprising representatives of industry, civil society, academia and other stakeholders. Article 56 provides for GPAI providers to develop codes of practice in consultation with stakeholders. Article 67, read with Recital 150, includes social-partner participation, while Article 56 permits relevant stakeholder input. Neither guarantees dedicated production-worker representation.

Proposed amendment: Article 67 (2) should be amended to require that the Advisory Forum include at least two seats designated for representatives of AI production worker organisations, selected through a process with clearly established legal authority and consultation with worker organisations and the International Labour Organization (ILO). Where no recognised production worker organisations exist with sufficient representational reach, a supported consultation process should enable workers to choose representatives; researchers may assist access but should not automatically substitute for worker voice.

Article 56 should be amended to require that the development and revision of GPAI codes of practice include a structured consultation phase with production worker representatives, documented in the code's published methodology. The current Code of Practice published in July 2025 demonstrates that codes can be developed with multi-stakeholder input; the amendment would make production-worker consultation a specified requirement within the existing stakeholder route.

Rationale: The epistemic audit finding, that workers whose knowledge shapes AI systems lack dedicated production-specific participation guarantees in AI governance, is the dimension of labour invisibility most directly addressable through institutional design rather than through documentation requirements. Worker participation in governance is not merely a representational aspiration; it is an epistemic resource. Workers who identify annotation errors, cultural misinterpretations and systematic biases in training guidelines, as participants in the EVF empirical base described doing, possess knowledge that governance structures may otherwise struggle to access. Making that knowledge available to the AI Office and the Advisory Forum is consistent with the Act's own commitment to trustworthy AI.

Analogue in EU law: The Information and Consultation Directive (2002/14/EC) and the European Works Council Directive [28] already establish worker participation rights in undertakings above specified thresholds. The amendment proposes a governance-level analogue at the EU regulatory body level, not a works council requirement on individual providers.

6.5 Amendment 4: Extending the Fundamental Rights Impact Assessment

What the Act currently does: Article 27 requires deployers of certain high-risk AI systems to conduct a fundamental rights impact assessment (FRIA) before deploying the system. The assessment must consider the risks the system poses to fundamental rights of affected persons. It is a deployment-side instrument.

Proposed amendment: Article 27 should be extended to require providers of high-risk AI systems and GPAI models to conduct a production-phase fundamental rights assessment covering the labour conditions of the production workforce. This assessment should consider: freedom from forced or compulsory labour; the right to fair and just working conditions; freedom of association and the right to collective bargaining; and the right to dignity at work, including protection from exposure to harmful content without adequate support.

The proposed production-phase assessment should be coordinated with the applicable high-risk conformity and supervisory route, and with a separate Chapter V duty for GPAI models. A new duty should specify review and referral responsibilities rather than presume individual regulatory approval of every assessment. It need not be made fully public, given legitimate concerns about commercial confidentiality, but a summary should be included in the AI Labour Supply Chain Statement proposed under Amendment 1.

Rationale: Fundamental rights impact assessment is already established in the Act's architecture as a mechanism for identifying and addressing harms that technical conformity requirements alone cannot capture. Extending it to production conditions is consistent with the Act's own logic: if fundamental rights of affected persons are a regulatory concern for AI deployment, the fundamental rights of the workers who produce those systems warrant equivalent consideration. The extension would introduce a substantive new provider-side duty, even though it builds on an existing assessment approach. Legislative authority, proportionality and supervisory capacity would need to be established.

6.6 The Amendment Package as a System

The four amendments proposed above address each EVF dimension directly. Amendment 1 addresses material invisibility by requiring supply chain disclosure. Amendment 2 addresses material and epistemic invisibility by incorporating production conditions into data quality governance. Amendment 3 addresses epistemic invisibility by creating worker participation rights in governance. Amendment 4 addresses all

three dimensions by requiring fundamental rights consideration at the production level.

Together, the amendments do not transform the Act into a labour regulation. They introduce a production accountability logic alongside the Act's existing product accountability logic, treating the conditions of AI production as a dimension of product trustworthiness rather than as an external matter. They use the Act's existing enforcement architecture, documentation requirements, conformity assessment procedures, and Advisory Forum structure, while requiring assessment of additional legal powers, resources and cooperation with labour authorities. And they apply the Act's own extraterritorial market access rationale to labour conditions, extending to the production workforce the jurisdictional reach the Act already exercises over technical compliance.

The amendments are individually meaningful but mutually reinforcing. Supply chain disclosure without worker participation produces data that governance bodies cannot act on with informed expertise. Worker participation without fundamental rights assessment produces representation without the evidentiary base to exercise it effectively. The package is designed to be adopted together, though each amendment has independent value if the others are not adopted.

Implementation should specify how workers can contest inaccurate records and obtain referrals to competent labour or judicial bodies, with confidentiality and protection against retaliation. Disclosure, participation and market restrictions are not substitutes for payment or compensation where these are owed. These proposals require institutional capacity and should be assessed for effects on small suppliers and the risk of abrupt contract termination.

7. Conclusion

This paper set out to examine what the EU AI Act has to say about the workers who produce AI systems. The answer, established through a systematic regulatory audit using the Ethical Visibility Framework, is that the Act weakly recognises AI labour relations across all three dimensions of visibility. Seven of twelve EVF indicators code as absent and five as partial recognition. The partial recognitions mark the boundary of the Act's product-risk frame rather than a comprehensive production-labour entitlement: Article 25 acknowledges distributed production without requiring labour supply chain transparency, and Article 14 acknowledges the necessity of human judgement without recognising the human judgement embedded in the production process. Participation, production vocabulary and general complaint rights provide further partial connections.

We examined and qualified the strongest rival explanation: that the Act's silence reflects a rational division of regulatory labour, with AI labour conditions properly governed by existing EU instruments. That explanation requires qualification on three grounds. Jurisdictionally, product-market reach does not automatically transmit employment entitlements throughout a transnational production chain. Substantively, existing EU labour instruments can address AI work, but their coverage and interaction must be established for particular production arrangements: contracting and employment status, exposure to psychologically harmful content, cross-border responsibility and recognition of workers' domain knowledge. And at the level of regulatory coherence, an Act that claims to govern trustworthy AI but does not independently govern the labour conditions of its production cannot fully deliver what that claim requires.

The enforcement paradox captures the consequential form of this failure. The Act now operates with significant sanctioning powers. Fines reach €35 million or 7% of global annual turnover. Market access can be restricted. The AI Office has investigative authority over GPAI providers. These powers attach to specified infringements; they do not create a general production-labour duty. Labour evidence may reveal an AI Act breach and support an Article 85 complaint, while other grievances require routes under applicable labour or contract law. The gap concerns the substantive object of enforcement, rather than the complete exclusion of workers from every procedure.

The four amendments proposed in Section 6 address this gap within the Act's existing architecture. An AI Labour Supply Chain Statement as part of technical documentation introduces supply chain transparency subject to an adequate legislative basis and separate system and GPAI documentation routes. An extension of Article 10's data quality provisions to include worker welfare conditions connects production conditions to the quality rationale the Act already applies. Worker participation in the Advisory Forum and GPAI code development introduces the epistemic resources that production workers possess into governance processes in which dedicated production-worker participation is not guaranteed. And a production-phase fundamental rights assessment applies an instrument already in the Act through a substantive extension to production conditions.

These amendments are intended to be targeted and proportionate; their operational feasibility requires consultation and impact assessment. They do not ask the Act to become something it is not. They ask it to take seriously what it already claims: that AI governance is about trustworthiness, and that trustworthiness cannot be assessed at the point of deployment without attending to the conditions of production.

The broader implication of this audit extends beyond the EU AI Act. AI governance frameworks are proliferating globally, across national and international policy settings. All of them face the same foundational choice: whether to govern AI as a product that enters a market or as an output of a labour process that involves specific workers bearing specific costs. The EU AI Act has made the first choice comprehensively and the second only in limited and indirect ways. Other frameworks in development now have the opportunity to make a different choice, or to learn from the EU's foundational gap.

The workers behind the AI are not incidental to its governance. They are its precondition. Making them visible, not as persons affected by AI decisions but as the distributed workforce whose labour makes those decisions possible, is not an add-on to trustworthy AI regulation. It is its unfinished foundation.

Abbreviations

AI	Artificial Intelligence
EU	European Union
EVF	Ethical Visibility Framework
GPAI	General-purpose Artificial Intelligence
RLHF	Reinforcement Learning from Human Feedback
CE	European Conformity Marking
GDPR	General Data Protection Regulation
CSDDD	Corporate Sustainability Due Diligence Directive
ILO	International Labour Organization
UAE	United Arab Emirates
FRIA	Fundamental Rights Impact Assessment
P	Presence
PR	Partial Recognition
A	Absence of a Dedicated Provision.

Author Contributions

Achi Iseko: Conceptualization, Formal Analysis, Investigation, Methodology, Writing – original draft, Writing – review & editing

Data Availability Statement

This article analyses public legal documents and published research; no new participant data were collected or reanalysed. The legal sources, indicators and classifications are identified in the references and Tables 1 and 2.

Conflicts of Interest

The author is employed by Microsoft and conducts this research independently. The views expressed are the author's own. Author to confirm: completeness of this disclosure and any additional relevant interests.

References

[1] Turing, A. M. Computing machinery and intelligence. Mind. 1950, 59(236), 433–460. https://doi.org/10.1093/mind/LIX.236.433

[2] European Parliament and Council of the European Union. Regulation (EU) 2024/1689 of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). Official Journal of the European Union. 2024, L 2024/1689. https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng

[3] Iseko, A. The Ghost Workforce of Generative AI: Reclaiming Ethical Visibility in AI Labour. International Journal of Science, Technology and Society. 2026, 14(2), 88–101. https://doi.org/10.11648/j.ijsts.20261402.15

[4] European Commission. AI Act. Shaping Europe's Digital Future. https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai Accessed 15 October 2026.

[5] European Commission. The General-Purpose AI Code of Practice. 2025; webpage updated 7 October 2026. https://digital-strategy.ec.europa.eu/en/policies/contents-code-gpai Accessed 15 October 2026.

[6] Cefaliello, A.; Kullmann, M. Offering false security: How the draft artificial intelligence act undermines fundamental workers rights. European Labour Law Journal. 2022, 13(4), 542–562. https://doi.org/10.1177/20319525221114474

[7] European Commission. Explanatory Notice and Template for the Public Summary of Training Content for general-purpose AI models. 24 July 2025. https://digital-strategy.ec.europa.eu/en/library/explanatory-notice-and-template-public-summary-training-content-general-purpose-ai-models Accessed 15 October 2026.

[8] Gray, M. L.; Suri, S. Ghost Work: How to Stop Silicon Valley from Building a New Global Underclass. Houghton Mifflin Harcourt, Boston, 2019.

[9] Tubaro, P.; Casilli, A. A.; Coville, M. The trainer, the verifier, the imitator: Three ways in which human platform workers support artificial intelligence. Big Data & Society. 2020, 7(1). https://doi.org/10.1177/2053951720919776

[10] Miceli, M.; Posada, J. The Data-Production Dispositif. Proceedings of the ACM on Human-Computer Interaction. 2022, 6(CSCW2), Article 460, 1–37. https://doi.org/10.1145/3555561

[11] Sambasivan, N.; Kapania, S.; Highfill, H.; Akrong, D.; Paritosh, P.; Aroyo, L. M. "Everyone wants to do the model work, not the data work": Data Cascades in High-Stakes AI. Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems. 2021, Article 39, 1–15. https://doi.org/10.1145/3411764.3445518

[12] Graham, M.; Hjorth, I.; Lehdonvirta, V. Digital labour and development: Impacts of global digital labour platforms and the gig economy on worker livelihoods. Transfer: European Review of Labour and Research. 2017, 23(2), 135–162. https://doi.org/10.1177/1024258916687250

[13] Wood, A. J.; Graham, M.; Lehdonvirta, V.; Hjorth, I. Good Gig, Bad Gig: Autonomy and Algorithmic Control in the Global Gig Economy. Work, Employment and Society. 2019, 33(1), 56–75. https://doi.org/10.1177/0950017018785616

[14] Baldwin, R.; Cave, M.; Lodge, M. Understanding Regulation: Theory, Strategy, and Practice. 2nd ed. Oxford University Press, Oxford, 2012. https://doi.org/10.1093/acprof:osobl/9780199576081.001.0001

[15] Black, J. Critical reflections on regulation. Australian Journal of Legal Philosophy. 2002, 27, 1–35.

[16] European Parliament and Council of the European Union. Directive (EU) 2024/2831 of 23 October 2024 on improving working conditions in platform work. Official Journal of the European Union. 2024, L 2024/2831. https://eur-lex.europa.eu/eli/dir/2024/2831/oj/eng

[17] European Parliament and Council of the European Union. Directive 2002/14/EC of 11 March 2002 establishing a general framework for informing and consulting employees in the European Community. Official Journal of the European Communities. 2002, L 80, 29–34. https://eur-lex.europa.eu/eli/dir/2002/14/oj

[18] European Parliament and Council of the European Union. Regulation (EU) 2016/679 of 27 April 2016 (General Data Protection Regulation). Official Journal of the European Union. 2016, L 119, 1–88. https://eur-lex.europa.eu/eli/reg/2016/679/oj

[19] European Parliament and Council of the European Union. Directive 2003/88/EC of 4 November 2003 concerning certain aspects of the organisation of working time. Official Journal of the European Union. 2003, L 299, 9–19. https://eur-lex.europa.eu/eli/dir/2003/88/oj

[20] European Parliament and Council of the European Union. Directive 2008/104/EC of 19 November 2008 on temporary agency work. Official Journal of the European Union. 2008, L 327, 9–14. https://eur-lex.europa.eu/eli/dir/2008/104/oj

[21] European Commission. Corporate sustainability due diligence. Overview of Directive (EU) 2024/1760 as amended by Directives (EU) 2025/794 and (EU) 2026/470. https://commission.europa.eu/topics/business-and-industry/company-law-and-corporate-governance/corporate-sustainability-due-diligence_en Accessed 15 October 2026.

[22] Council of the European Communities. Directive 89/391/EEC of 12 June 1989 on the introduction of measures to encourage improvements in the safety and health of workers at work. Official Journal of the European Communities. 1989, L 183, 1–8. https://eur-lex.europa.eu/eli/dir/1989/391/oj

[23] European Commission. The enforcement framework of the AI Act. Shaping Europe's Digital Future. https://digital-strategy.ec.europa.eu/en/policies/enforcement-ai-act Accessed 15 October 2026.

[24] International Labour Organization. World Employment and Social Outlook 2021: The role of digital labour platforms in transforming the world of work. ILO, Geneva, 2021. https://researchrepository.ilo.org/esploro/outputs/report/995218610802676 Accessed 15 October 2026.

[25] World Bank. Working Without Borders: The Promise and Peril of Online Gig Work. World Bank, Washington, DC, 2023. https://doi.org/10.1596/40066 Accessed 15 October 2026.

[26] Fairwork. Fairwork Cloudwork Ratings 2025: Advancing Standards in Digital Labour and AI Supply Chain Governance. 2025. https://fair.work/wp-content/uploads/sites/17/2025/05/Fairwork-Cloudwork-Report-2025-FINAL.pdf Accessed 15 October 2026.

[27] International Labour Organization. Navigating the digital and artificial intelligence revolution in Arab labour markets: Trends, challenges and opportunities. ILO Regional Office for the Arab States, 2025. https://doi.org/10.54394/OYOA8722 Accessed 15 October 2026.

[28] European Parliament and Council of the European Union. Directive 2009/38/EC of 6 May 2009 on the establishment of a European Works Council or a procedure in Community-scale undertakings and Community-scale groups of undertakings for the purposes of informing and consulting employees (Recast). Official Journal of the European Union. 2009, L 122, 28–44. https://eur-lex.europa.eu/eli/dir/2009/38/oj