



POLICY BRIEFING · ARTIFICIAL INTELLIGENCE ACT (REGULATION (EU) 2024/1689)

The Employment Blind Spot

How Article 27 of the EU AI Act Exempts the Private-Sector
Systems Most Likely to Discriminate Against Religious Minorities
from Fundamental Rights Review

Goran Kolouch

Independent Researcher · Founder, EqualFaith Europe

July 2026

EQUALITY · DIGNITY · JUSTICE

Executive Summary

Article 27 of the AI Act creates the Regulation's only pre-deployment check on how a high-risk AI system will affect a person's fundamental rights — the Fundamental Rights Impact Assessment ("FRIA"). But Article 27(1) confines that duty to **public bodies, private entities providing public services, and deployers of credit-scoring or insurance-pricing systems**. It does not name Annex III, point 4 — the category covering recruitment, CV-screening, candidate ranking, and performance and termination decisions. The result: an ordinary private employer's AI hiring tool is classified as *high-risk* under the Act, carries the full documentation and human-oversight burden of that classification, and yet is the one deployer category **statutorily exempt from ever having to assess, in advance, whether the system will discriminate**.

This is not a hypothetical risk. A substantial and consistent body of field-experimental research across Germany, the Netherlands, Norway, Spain and the United Kingdom shows employers penalising applicants whose CVs signal Muslim identity — through name, headscarf photographs, or faith-based volunteer work [10, 11] — and a 2025 socio-legal study specifically confirms that AI hiring systems reproduce exactly this bias against headscarf-wearing and minority-name candidates [12]. The Court of Justice's case law on workplace religious symbols (*Achbita*, *Bougnouli*) was built to test a single, legible policy applied by a human decision-maker [13, 14] — it has no purchase on an opaque ranking algorithm that never announces a policy at all.

The Digital Omnibus, formally adopted by the European Parliament (16 June 2026) and the Council (29 June 2026), has just deferred Annex III high-risk obligations — including employment — from 2 August 2026 to **2 December 2027** [16, 17]. This briefing sets out the legal reasoning behind the gap, its consequences, and a concrete legislative and interpretive remedy — and argues that the eighteen months this delay has opened is the narrow window in which to secure it.

1. Two Rights, One Missing Link

The AI Act gives affected individuals two things where a high-risk system is used against them: a right to complain to a market surveillance authority (Article 85) and a right to an explanation of the decision (Article 86) [3, 4]. Both are reactive — they operate only once a person has already been rejected, downgraded, or dismissed. The Act's only *preventive* mechanism, capable of catching a discriminatory design choice before it ever reaches a candidate, is the Fundamental Rights Impact Assessment under Article 27 [1]. Whether that mechanism reaches employment AI at all is therefore not a technical footnote — it determines whether the entire regime is preventive or purely remedial for the labour market.

2. The Legal Architecture of Article 27

2.1 What Article 27 Requires

Before first use of a high-risk system, a covered deployer must document its intended use, the categories of people likely to be affected, the specific risks of harm to those people, the human-oversight measures in place, and the steps it will take if risks materialise — and must notify the market surveillance authority of the result [16]. It is the single most substantive fundamental-rights obligation the Act places on deployers, as distinct from providers.

2.2 Who It Binds — and Who It Doesn't

Article 27(1) limits the duty to "deployers that are bodies governed by public law, or are private entities providing public services, and deployers of high-risk AI systems referred to in points 5(b) and (c) of Annex III" — namely creditworthiness assessment and life/health insurance risk-pricing [1, 15]. "Private entities providing public services" is read narrowly, by reference to Recital 96, as covering services "linked to tasks in the public interest" such as education, healthcare, social services, housing and the administration of justice [7]. An ordinary private employer — a retailer, a logistics firm, a consultancy — is neither a public body nor, on any plausible reading, a provider of a public-interest service.

2.3 The Annex III, Point 4 Carve-Out

Annex III, point 4 classifies as high-risk any AI system "intended to be used for recruitment or selection," including targeted job advertising, CV filtering and candidate ranking, and any system used to make decisions on promotion, termination, task allocation, or performance evaluation [8, 9]. This classification applies *regardless of the size of the organisation* and regardless of whether the tool is bought off-the-shelf or built in-house [9]. Yet point 4 is conspicuously absent from the Article 27(1) carve-in list that names points 5(b) and (c). Two independent sources confirm this is not an oversight of interpretation but the plain text: a specialist compliance guide states flatly that "Article 27 FRIA is **not** triggered for a private-sector HR deployer (§4 is not listed in Article 27)" [8], and a 2026 academic paper on regulating AI agents lists "employment management systems" among the uses that "do not trigger a FRIA obligation under the Act," concluding that this "limits the mechanism's capacity to address unfair decision-making" precisely where it is most needed [9].

The asymmetry is stark: a bank's AI credit-scoring tool must undergo a FRIA before it ever scores a customer. A private employer's AI tool that screens, ranks, and eliminates candidates for a job — arguably a more consequential and more proxy-laden decision — need not.

3. The Redress Deficit: Articles 85 and 86

Even where a FRIA is not required, one might expect the Act's redress provisions to compensate for the missing prevention. They do not. Article 85 allows any person to lodge a complaint with a market surveillance authority, but the authority's response is discretionary market monitoring, not adjudication of the individual's claim [3]; Article 86 entitles a person to an explanation of a decision, but only after it has been made, and it obliges the deployer to explain, not to justify or remedy [4]. Neither article creates a right to compensation. One legal commentator has already described Article 85's complaint right as "an insignificant consolidation" of what individual claimants actually need [18], and a 2025 review of the Act's redress architecture concludes plainly that "the Act does not provide redress mechanisms for individuals" and "does not guarantee affected persons compensation for harms" [19].

Contrast this with the General Data Protection Regulation, whose Article 82 gives any person a directly enforceable right to compensation for

material *and* non-material damage arising from an infringement — a right the Court of Justice has since interpreted expansively, extending it to dignitary and emotional harms that fall well short of financial loss. The AI Act, despite governing decisions with materially greater consequences for a person’s livelihood, transplanted none of this. A rejected candidate can ask *why* an algorithm ranked them last; under the current text, she cannot claim compensation *under the AI Act itself* for having been ranked last unlawfully in the first place.

4. Why This Gap Falls Hardest on Religious Minorities

4.1 The Empirical Record

The evidence that hiring discrimination against religious minorities is systemic, cross-national, and legible to a machine-learning model is unusually strong for a labour-market bias claim. A harmonised field experiment across Germany, the Netherlands, Norway, Spain and the United Kingdom found consistent penalties against applicants signalled as Muslim through CV and cover-letter details, holding qualifications constant [10]. German correspondence testing found that applicants with Turkish names received markedly fewer callbacks than identical applicants with German names, and that a headscarf photograph compounded the penalty further [11, 20]. Critically, a 2025 empirical study of real AI hiring systems — not laboratory proxies — found that deployed CV-parsing, assessment, and video-interview tools “discriminate against applicants who wear a headscarf or have a Black-sounding name,” and traces this to the well-documented case of an AI recruiting tool that taught itself to downgrade applications correlated with being a woman [12]. In other words: the very features (name, address, faith-linked activities, photograph) that decades of correspondence-testing literature identify as triggers for human bias are the same categorical features an automated ranking model learns to weight — often without anyone instructing it to.

4.2 A Doctrinal Blind Spot

EU non-discrimination doctrine on religious symbols in employment was built for a different fact pattern. In *Achbita v G4S* and *Bouagnaoui v Micropole*, the Court of Justice tested whether a company’s stated neutrality policy, applied by a named manager to a known employee, amounted to unlawful direct or indirect discrimination [13, 14]. That framework needs a visible policy and a traceable decision-maker to bite. An AI ranking system that quietly down-weights CVs correlated with a Muslim-sounding name or a majority-Muslim postcode produces the same exclusionary outcome without ever generating a policy statement a court, or the applicant, can point to. Absent a mandatory FRIA to force that reasoning into the open *before* deployment, the doctrine has nothing left to test.

5. Consequences of Inaction

- **Scale without scrutiny:** Annex III, point 4 applies to every private employer in the EU using a recruitment, ranking, or performance-management AI tool, “regardless of the size of the organisation” [9] — the exemption is therefore economy-wide, not a marginal case.
- **A one-way ratchet:** because deployers face no obligation to look for discriminatory impact in advance, harm can only surface after the fact, through litigation an individual candidate would have to fund and prove without the documentary trail a FRIA would have generated.
- **Undermined trust in the Act’s own architecture:** the AI Act presents itself as a fundamental-rights-centred, risk-based regulation; a hole at precisely the intersection of “high-risk” and “protected characteristic” invites the criticism that the risk-based approach protects processes, not people [19].
- **A widening compliance-fairness gap:** employers are already told that a FRIA is “strongly recommended best practice” even where not legally required [9] — meaning the market itself recognises the risk, while the law declines to mandate the safeguard.

6. The Policy Window: The Digital Omnibus

The Digital Omnibus — the Commission’s package amending the AI Act, endorsed by the European Parliament on 16 June 2026 and given final Council approval on 29 June 2026 — defers the application of Annex III high-risk obligations, including employment, from 2 August 2026 to **2 December 2027**; publication in the Official Journal was expected during July 2026 [16, 17]. The Omnibus was adopted to give industry and regulators more time to prepare for standards that were not ready — it did not touch the underlying scope question of who Article 27 binds [17]. That is precisely what makes this the right moment: the substantive obligations employers will eventually face are not yet in force, so a scope amendment or interpretive clarification can still be secured *before* the current gap hardens into settled practice and vested compliance expectations.

7. Proposed Solution

7.1 Legislative fix — amend Article 27(1)

Add Annex III, point 4 to the list of carve-ins in Article 27(1), placing employment alongside credit scoring and insurance pricing. This is a narrow, single-clause amendment, not a reopening of the Act’s structure, and it can ride the same simplification vehicle that produced the Digital Omnibus, or a future targeted amendment ahead of the December 2027 application date.

7.2 Interpretive fix — clarify Recital 96

Failing legislative action, the AI Office or the Commission can issue guidance recognising that large-scale intermediation of access to employment — a right protected under Article 15 of the EU Charter of Fundamental Rights — constitutes a “task in the public interest” within the meaning of Recital 96 when performed by dominant recruitment platforms or applicant-tracking providers used at scale. This route is faster and does not require reopening the Regulation, though it rests on a more contestable reading of the existing text.

7.3 Redress fix — a compensation right modelled on GDPR Article 82

Insert into Article 85, or a new Article 85a, an explicit right to compensation for material and non-material damage caused by an infringement of the Act — mirroring GDPR Article 82 and the Court of Justice’s now-settled jurisprudence recognising non-material, dignitary harm as

compensable. This closes the redress deficit identified in Section 3 regardless of whether the FRIA scope itself is widened, and gives Article 86’s right to explanation a remedy worth invoking.

8. Recommended Actions for EqualFaith Europe

- Draft a formal submission to DG CONNECT and the AI Office, in the format of the existing Article 28/34(1)(b) DSA briefing, proposing the Article 27(1) amendment and the Article 85a compensation right set out above.
- Time the submission to the pre-December-2027 window: the earlier this reaches the relevant committees, the more plausible it is that a scope fix can be attached to implementing guidance or a further legislative vehicle before enforcement begins.
- Commission or conduct a comparative study, drawing on EqualFaith’s existing 27-country research base, documenting AI recruitment-tool penetration and its differential impact on religious minorities — giving the submission original empirical weight beyond the secondary literature cited here.
- Coordinate messaging with the existing DSA Article 40(12) data-access strategy: both arguments rest on the same underlying claim — that rights written for civil society and individuals are being absorbed by carve-outs that favour the platforms and employers the rights were meant to constrain.

Respectfully submitted,

Goran Kolouch *Independent Researcher; Founder, EqualFaith Europe*

References

[1] Regulation (EU) 2024/1689 (AI Act), Article 27 (Fundamental rights impact assessment for high-risk AI systems).

[2] Regulation (EU) 2024/1689, Article 6(2) and Annex III (classification of high-risk AI systems).

[3] Regulation (EU) 2024/1689, Article 85 (Right to lodge a complaint with a market surveillance authority).

[4] Regulation (EU) 2024/1689, Article 86 (Right to explanation of individual decision-making).

[5] Regulation (EU) 2016/679 (GDPR), Article 82 (Right to compensation and liability).

[6] activeMind.legal, ‘Article 27 — Fundamental rights impact assessment for high-risk AI systems,’ AI Act (Nov. 2024).

[7] Center for Democracy & Technology, ‘EU AI Act Brief — Pt. 4, AI at Work’ (Mar. 2026).

[8] Sprinkling Act, ‘AI Act for HR: Recruitment, screening, performance evaluation’ (2026).

[9] ‘Regulating AI Agents,’ arXiv:2603.23471 (2026).

[10] Di Stasio et al., ‘Muslim by default or religious discrimination? Results from a cross-national field experiment on hiring discrimination,’ *Journal of Ethnic and Migration Studies*.

[11] Weichselbaumer, ‘Discrimination Against Female Migrants Wearing Headscarves,’ *IZA Journal of Labor Economics* (2018).

[12] Sheard, ‘Algorithm-facilitated discrimination: a socio-legal study of the use by employers of artificial intelligence hiring systems,’ *Journal of Law and Society* (2025).

[13] CJEU, Case C-157/15, *Achbita v G4S Secure Solutions*.

[14] CJEU, Case C-188/15, *Bouagnaoui and ADDH v Micropole SA*.

[15] Legalithm, ‘Article 27 FRIA’ guide (Apr. 2026).

[16] Jones Walker LLP, ‘Yes, August 2 Still Matters: The EU Approved a High-Risk AI Delay’ (Jul. 2026).

[17] Gibson Dunn, ‘EU AI Act Omnibus Agreement — Postponed High-Risk Deadlines and Other Key Changes’ (May 2026).

[18] MediaLaws, ‘The AI Act’s Right to Explanation: A Plea for an Integrated Remedy’ (Oct. 2024).

[19] ‘The EU AI Act in Development Practice: A Pro-justice Approach,’ arXiv:2504.20075 (2025).

[20] Antidiskriminierungsstelle des Bundes, ‘Wearing a headscarf in the workplace.’