waveImpact
RESPONSIBLE AI CONSULTANCY

Working scaffold · Path 1 · Banks and insurers

Fundamental Rights Impact Assessment, Article 27

The six required contents, mirrored against the model file — with an evaluation matrix

Regulation (EU) 2024/1689 is accommodating towards regulated financial institutions in many places. The risk management system under Article 9 may form part of existing procedures (paragraph 10); for providers that are financial institutions, the quality management system is largely deemed fulfilled through internal governance (Article 17(4)); the deployer's monitoring obligation likewise (Article 26(5)); and the logs are simply part of the documentation already kept (Article 26(6)).

Article 27 contains no such clause. The fundamental rights impact assessment is the one artefact for which the model file offers no template — even where material from other parts of the institution feeds into it. The only bridge the provision itself offers points not towards model validation but towards the data protection impact assessment; the Digital Omnibus widened precisely that bridge in July 2026.

Scope of use	A working scaffold for the first pass. Step 0 sets the perimeter, sections 1 to 6 work through the six required contents of Article 27(1), section 7 covers evaluation and the residual risk decision, and sections 8 to 10 cover notification, updating and signature.
What it is not	Not a template for the AI Office questionnaire — that does not yet exist. Not legal advice in an individual case. And deliberately not a procedure for adjusting the model: for internal models that is a supervisory event of its own kind and does not belong in a review scaffold.
On the rows	The row “Still to evidence” names the methodologically appropriate evidence for each required content, not a statutory catalogue of metrics. Article 27 fixes the subject of the assessment; the institution chooses the method — and justifies it.
Legal position	Regulation (EU) 2024/1689 as amended by Regulation (EU) 2026/1744 (Digital Omnibus on AI), in force since 27 July 2026. Article 27 sits in Chapter III, Section 3 and applies to Annex III systems from 2 December 2027. For systems already in use before then, the transitional rules in Article 111 apply, turning in substance on significant changes — any substantial model overhaul puts the question afresh.

On sequence: points (c) and (d) require a data basis that can often only be reconstructed incompletely after the fact. An institution that stands up group-differentiated measurement only at the application date has a cross-section and not a time series — and therefore nothing to say about whether a measured difference is entrenching or eroding. The effort of the first pass is manageable; the value accrues in repetition.

System / model	Institution / business area	Working status as at

Step 0 Set the perimeter, rather than assume it

Article 27(1) reaches deployers of high-risk AI systems under Annex III, point 5(b) and (c) — regardless of legal form or ownership. The reach is narrower than it looks at first glance, and where it does bite it is not negotiable.

Inside the perimeter

- ☐ Evaluation of the creditworthiness of **natural persons** or establishing their credit score (Annex III, point 5(b)).
- ☐ Risk assessment or pricing in relation to **natural persons** in **life or health insurance** (Annex III, point 5(c)).

Outside the perimeter — record the exclusion expressly

- ☐ Systems used to **detect financial fraud** (carve-out in point 5(b)). The carve-out only holds where fraud detection is genuinely the intended purpose.
- ☐ Creditworthiness assessment of **legal persons**.
- ☐ Insurance lines **outside life and health** — motor, property and liability are not covered by point 5(c).

Check profiling	Where the system performs profiling of natural persons it is always high-risk, and the filter in Article 6(3) is closed. Credit scoring and individualised risk or premium assessment will regularly meet the profiling definition; whether they do in a given case turns on how the system actually works and which attributes it processes.
Check the role	The obligation attaches to the deployer, not the provider. An institution using a model from a credit bureau, a shared service centre, a group IT provider or a specialist scoring or software vendor cannot transfer responsibility for the assessment to that supplier. The provider information under Article 13 is an essential input; the assessment itself is owed by the institution that uses the system.

Systems in scope (number / designation)	Excluded, with reasons

Sections 1–6

The six required contents of Article 27(1)

The six contents fall into two groups. Points (a), (e) and (f) can be assembled from existing material — with re-cutting, cross-referencing and one addition at (f). Points (b), (c) and (d) are produced from scratch.

Cross-referencing	Since the Digital Omnibus, Article 27(4) expressly permits cross-references to the relevant sections of the data protection impact assessment and the incorporation of parts of it. For a scoring system such an assessment will usually already exist. The first working session should therefore bring risk control, the business line and data protection to the same table.

Point (a)

Description of the processes in which the system is used as intended

What is required	The deployer's processes in which the system is used in line with its intended purpose.
Usually available	Work instructions, credit or underwriting policy, organisational and process documentation.
Still to evidence	The re-cut around the specific system: where in the process the output takes effect, which decision is based on it, which alternative paths exist. Existing material is cut along processes, not systems.

Reference inside the institution	Owner	Status / date

Point (b)

Period and frequency of the intended use

What is required	A description of the period within which, and the frequency with which, the system is intended to be used.
Usually available	Usually nothing that fits. The model file records validation validity periods — a different thing from intended deployment duration.
Still to evidence	An express determination: start of use, intended duration, decision volume per period, planned review points. A usage decision rather than a model property — and therefore not something the model side can supply.

Reference inside the institution | Owner | Status / date

Point (c)

Categories of natural persons and groups affected

What is required	The categories of natural persons and groups of persons likely to be affected by the use in the specific context.
Usually available	Portfolio segmentation, scope of application, target population.
Still to evidence	Risk classes are not affected groups. What is required is a reasoned selection of the categories and groups affected in the actual deployment context. The starting point is Article 21 of the Charter of Fundamental Rights together with the Union equal treatment directives and their national implementing law — in particular Directive 2004/113/EC for insurance pricing and Directive 2000/43/EC for access to goods and services — but not as a closed list. Context-specific and particularly vulnerable groups belong in it too, such as people with thin credit files, atypical employment histories, or in the role of co-borrower or guarantor, as do indirect, proxy and intersectional effects. Determine at the same time which groups are measured empirically and which are assessed qualitatively; where no reliable data basis exists, record the limit of inference expressly.

Reference inside the institution | Owner | Status / date

Point (d)

Specific risks of harm to the groups determined under (c)

What is required	The specific risks of harm to those groups, taking into account the information provided by the provider under Article 13.
Usually available	Validation report covering discriminatory power, calibration and stability; override statistics. Direction of view: the institution's risk.
Still to evidence	The direction of view inverts — what is assessed is the risk to the affected person. Methodologically appropriate and regularly necessary: group-differentiated outcome and error distributions, refusal and loading rates by group, development over time, and a documented methodology including case numbers, uncertainties and limits of inference. This is the most data- and method-intensive part of the assessment.

Reference inside the institution | Owner | Status / date

On the legal basis	The blanket objection that special categories of personal data may not be processed for a bias analysis under any circumstances no longer holds since the introduction of Article 4a by the Digital Omnibus: paragraph 2 expressly extends to deployers. The provision does not, however, open a general permission to collect protected attributes. It is a narrowly construed exception subject to cumulative conditions and the standard of strict necessity, and on the recitals it is to be read as a specification of substantial public interest within the meaning of Article 9(2)(g) GDPR rather than as a free-standing legal basis. Purpose, absence of alternatives, data scope, safeguards, retention, access rights and deletion are to be documented beforehand, and the data protection impact assessment will regularly need updating. The provision sits in Chapter I and has applied since 27 July 2026 — it does not follow the deferral of the high-risk obligations. It expressly creates no obligation to perform bias detection.

Point (e) Implementation of human oversight measures

What is required	A description of the implementation of human oversight measures in accordance with the instructions for use.
Usually available	Authority framework, dual-approval and second-opinion rules, four-eyes principle. The content on which regulated institutions are best placed.
Still to evidence	What has to be evidenced is not the rule but its practical effectiveness. Article 14(4) names the test points itself: comprehensibility of the output, awareness of the tendency to rely automatically on it, and the authority to depart from it. Collect qualification, powers, information available and time available to the people exercising oversight, together with the analysed override and escalation cases and their distribution across cases, staff and organisational units. A low override rate on its own evidences nothing: it may indicate a well-functioning system or an oversight function that is effectively inert.

Reference inside the institution	Owner	Status / date

Point (f) Measures on materialisation, internal governance and complaint mechanisms

What is required	The measures to be taken should those risks materialise, including the arrangements for internal governance and complaint mechanisms.
Usually available	Complaint handling, escalation routes, ombudsman and ADR schemes.
Still to evidence	The missing return path: a defined route from complaints, overrides and other anomalies, via examination of the individual case, back into model monitoring and governance. Determine the attribution of the complaint to the system actually used, the individual redress — for solely automated decisions supplemented by Article 22(3) GDPR, and for consumer credit by the human-intervention and explanation rights under Directive (EU) 2023/2225 — the trigger threshold for a systemic investigation, the escalation route and the owner of corrective action. In practice the complaint usually ends in customer service.

Reference inside the institution	Owner	Status / date

Section 7 Evaluation and the residual risk decision

Points (c) and (d) establish whom the system affects and to what extent. An impact assessment is not finished at that point: it calls for an evaluation. Otherwise it remains open which measured difference is objectively explicable, which remains legally or ethically problematic, which residual risk is borne — and who takes that decision.

The matrix below records that step. It stays deliberately at deployer level: controls and decisions concern the process, human oversight and the complaint route — not adjustment of the model. Adjusting an internal model is a supervisory model change event and belongs in the procedure established for it, not in a review scaffold.

Affected group	Observed difference	Objective explanation examined	Severity and reach	Existing control (process level)	Residual risk	Decision, decision-maker, date

How to use it

The matrix answers four questions that regularly stay open in practice. Is the observed difference explained by an objective, articulable reason — and is that examination documented? How severe and how far-reaching is it for the group concerned? Which existing control at process level actually mitigates it? And who decides that the remaining residual risk is borne, by name and date. The fourth is the question on which assessments fail in review.

Section 8 Notification: for whom the results are prepared

Article 27(3) requires the deployer to notify the market surveillance authority of the results of the assessment, submitting the completed questionnaire referred to in paragraph 5. That is not the same as sending the full internal report — though it does mean the internal report has to carry the results notified.

Who receives the notification is settled by Article 74(6): for high-risk AI systems used by financial institutions regulated under Union financial services law, the market surveillance authority is the national authority responsible for their financial supervision, in so far as the use is in direct connection with the provision of those financial services. The identity of that authority changes with the jurisdiction — BaFin, ACPR, DNB, IVASS, the Central Bank of Ireland — but the pattern does not. Where an institution sits outside the perimeter of the principal supervisor, the competent authority is a different one; establish it before notification rather than assuming it.

The quality standard follows. The results are evaluated in a prudential setting and have to connect to the level of evidence customary there. The documentation standard of a validation report is the natural benchmark: named method, named data basis, named limits of inference. That it is literally the same body that assesses the models does not hold throughout, however — internal model approval for SSM-significant institutions sits with the European Central Bank, not with the national supervisor.

On the questionnaire

The questionnaire under Article 27(5) does not yet exist. That is no reason to wait: the six required contents are in paragraph 1 and are independent of format. The questionnaire is the transmission form, not the substance.

Section 9 Define the update triggers

Article 27(2) attaches the obligation to first use, permits reliance on earlier assessments in similar cases, and requires updating as soon as an element is no longer current. Without defined triggers the obligation cannot be managed. The most economical place to anchor them is the existing validation cycle — with the caveat that the fundamental rights impact assessment must not be absorbed into it, because it answers a different question.

Trigger	Threshold / definition adopted	Owner	Last reviewed
Model change (re-estimation, feature change, threshold adjustment)			
Material change to the data basis (new source, new proxy variable, changed coverage)			
Change to the credit or underwriting policy			
Portfolio or population shift beyond a defined threshold			
Threshold breach in ongoing monitoring			
Technical or organisational incident (malfunction, data quality issue, faulty interface)			
New provider information under Article 13 (new version, changed instructions for use)			
Finding from the complaint return path (point (f))			
New legal or factual findings (case law, supervisory guidance, research)			
Periodic review, linked to the validation cycle			

Section 10

Signature: separate preparation, review and accountability

Article 27 prescribes no particular organisational signing order. For regulated financial institutions, however, a separation of specialist preparation, independent review, advisory involvement and final accountability suggests itself — not least because three functions compete for ownership in practice: risk control, compliance and data protection. Allocating the file to any one of them alone produces a document with a known weakness: data protection lacks the model empirics, compliance lacks the measurement basis, risk control lacks the fundamental rights perspective.

Role	Typical function	Name	Date
Specialist preparation	Model ownership and the responsible business line		
Methodological review and countersignature	Independent validation or an equivalent control function		
Legal and fundamental rights review	Compliance or the legal function		
Data protection involvement	Data protection officer or data protection function (paragraph 4)		
Accountability and release	The responsible management body as deployer		

waveImpact's recommendation: a fundamental rights impact assessment that independent validation does not countersign is a document about a model for whose statements nobody stands. Whether release sits with the full management body or a delegated function follows the institution's own authority framework.

If you would like to discuss how this cuts

The point at which this scaffold stalls in practice is regularly point (d) — the group-differentiated measurement that does not fall out of validation. waveImpact produces that measurement in a sealed, client-side environment: analysis of the source data runs in a container inside the client's estate, and only pre-agreed, reviewed analysis results are processed on EU-resident infrastructure for reporting. Raw and record-level data do not leave the client environment. For the vendor assessment of a regulated institution, that is the difference that matters.

A plain email is enough for a technical question: **valentin.mayr@waveimpact.de**. No form, no newsletter, no registration — and no commitment.

Dr Valentin José Mayr is founder and managing director of waveImpact GmbH in Bremen, Germany. waveImpact tests AI systems for fairness, transparency and regulatory fit.