HOW TO BUILD AI AGENTS FOR RISK & COMPLIANCE

Compliance Monitoring

A practical workshop on agentic workflows for regulated teams

01 Core concepts

02 Building an agentic workflow

03 Lessons from the field

04 Implementation case study

About us

Ashi Bajwa

Regulatory AI Architect, Zango

Sam Green

Senior Policy & Partnerships Lead, Zango

Joel Viney

Compliance Director, Lloyds Wealth

Tim Tyler

Vice President, ICA

POLL

What brings you here today?

Compliance monitoring lifecycle

Mapping and assurance

Builds the control universe and gap baseline

Ongoing monitoring

Risk-rated testing of whether controls operate over time

Risk-based plan, control testing, sampling, thematic reviews, MI

Reporting and remediation

Findings escalated, controls fixed, framework updated

Issues to committees, remediation tracking, control redesign

Compliance monitoring lifecycle

Mapping and assurance

Builds the control universe and gap baseline

Ongoing monitoring

Risk-rated testing of whether controls operate over time

Risk-based plan, control testing, sampling, thematic reviews, MI

Reporting and remediation

Findings escalated, controls fixed, framework updated

Issues to committees, remediation tracking, control redesign

Mapping and assurance

Balancing the tension between automation and accuracy

Automation

Speed, scale, coverage

Accuracy

Trust, defensibility, control

The right operating point is not fixed — it moves with the risk of each task. The job is to place the human exactly where the risk demands it.

Human in the loop

Human in the loop

BEFORE

Input

Give the model what it needs before it runs — context it may not have, and instruction on how to assess.

DURING

Checkpoint

Low-confidence or high-risk outputs route to a human, who can change or reject the decision before it moves on.

AFTER

Review

Check the model's outputs before they are relied on, with corrections feeding back into future runs.

Human in the loop

BEFORE

Input

Give the model what it needs before it runs — context it may not have, and instruction on how to assess.

DURING

Checkpoint

Low-confidence or high-risk outputs route to a human, who can change or reject the decision before it moves on.

AFTER

Review

Check the model's outputs before they are relied on, with corrections feeding back into future runs.

Why sample-based reviews don't work anymore

~5%
sampled under the old model

→

100%
reviewable with agents

Sampling was a workaround for limited human capacity — it was never the goal.

A 5% sample leaves the other 95% unseen, and risk does not distribute itself evenly across the population.

Agents make full-population review economic, turning sampling into a deliberate blind spot rather than a constraint.

What factors to take into account when designing the right levels of HITL intervention?

Confidence

Materiality

Detectability

Designing the right intervention point

Risk appetite for using AI — more risk and harder evaluation pull the human further in

1

AI Decision

Full autonomy on low-risk, well-tested tasks.

2

AI Decision + AI Evaluation

A second model checks the first before it passes.

3

AI Decision + Human Review

AI proposes; a human reviews exceptions before reliance.

4

Human Decision

The human decides; AI assists and assembles evidence.

LOWER RISK · MORE AUTONOMY

HIGHER RISK · MORE HUMAN JUDGEMENT →

POLL

What is your current approach to reviewing AI?

Mapping and assurance

Mapping and assurance

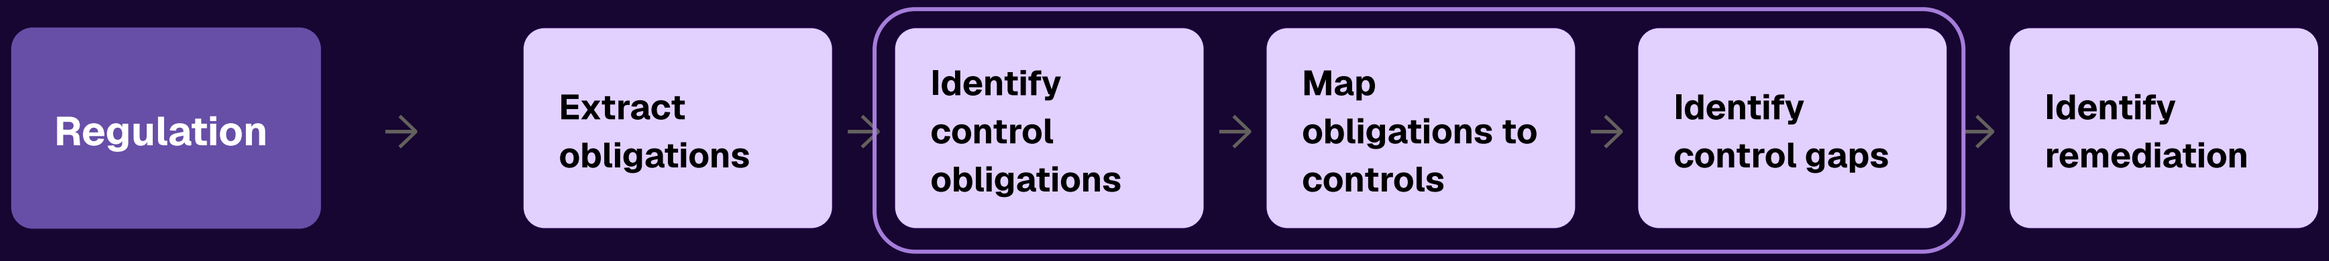

Let's build a control mapping agent

INPUT

Control obligations

→

AGENT

Control mapping

→

OUTPUT

Control gaps

The regulation

We start from the raw source — here, the draft Money Laundering and Terrorist Financing (Amendment) Regulations 2026 — and let the agent extract atomic obligations from the text itself.

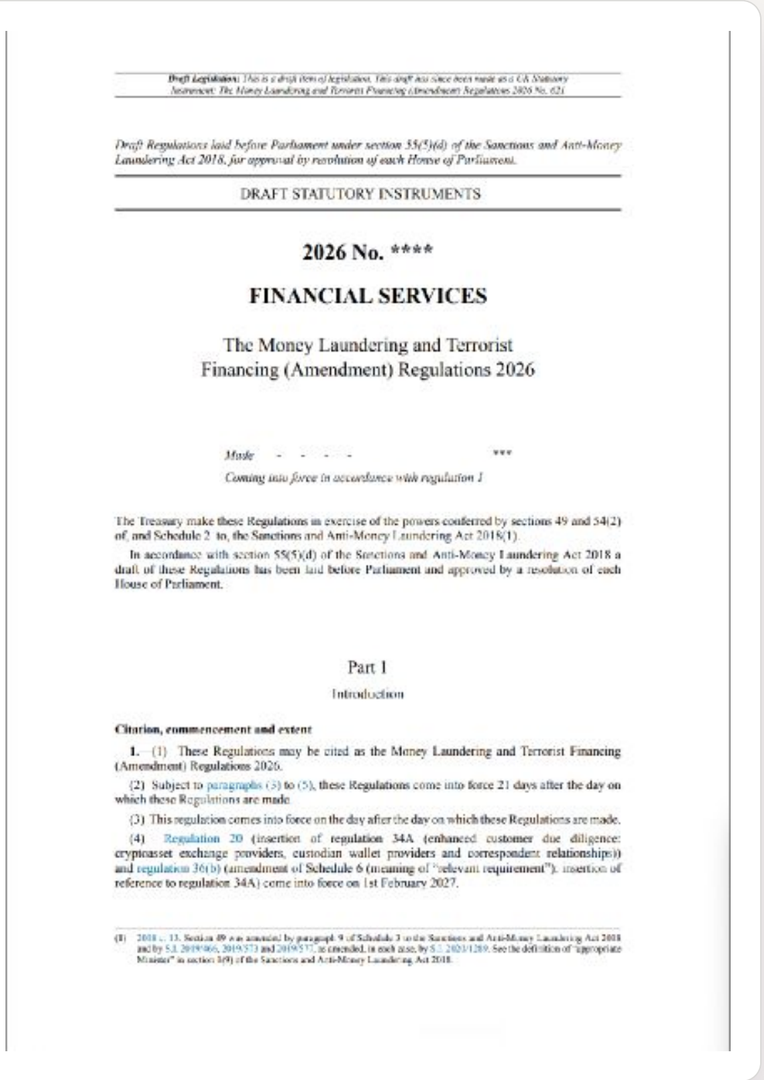

Draft Legislation: *This is a draft item of legislation. This draft has since been made as a UK Statutory Instrument: The Money Laundering and Terrorist Financing (Amendment) Regulations 2026 No. 621*

Draft Regulations laid before Parliament under section 55(5)(d) of the Sanctions and Anti-Money Laundering Act 2018, for approval by resolution of each House of Parliament.

DRAFT STATUTORY INSTRUMENTS

2026 No. ****

FINANCIAL SERVICES

The Money Laundering and Terrorist Financing (Amendment) Regulations 2026

*Made - - - - ****

Coming into force in accordance with regulation 1

The Treasury make these Regulations in exercise of the powers conferred by sections 49 and 54(2) of, and Schedule 2 to, the Sanctions and Anti-Money Laundering Act 2018(1).

In accordance with section 55(5)(d) of the Sanctions and Anti-Money Laundering Act 2018 a draft of these Regulations has been laid before Parliament and approved by a resolution of each House of Parliament.

Part 1

Introduction

Citation, commencement and extent

1. (1) These Regulations may be cited as the Money Laundering and Terrorist Financing (Amendment) Regulations 2026.

(2) Subject to paragraphs (3) to (5), these Regulations come into force 21 days after the day on which these Regulations are made.

(3) This regulation comes into force on the day after the day on which these Regulations are made.

(4) Regulation 2D (insertion of regulation 34A (enhanced customer due diligence: cryptoasset exchange providers, custodian wallet providers and correspondent relationships)) and regulation 36(b) (amendment of Schedule 6 (meaning of "relevant requirement")), insertion of reference to regulation 34A) come into force on 1st February 2027.

(1) 2018 c. 13. Section 49 was amended by paragraph 9 of Schedule 3 to the Sanctions and Anti-Money Laundering Act 2018 and by S.I. 2019/1565, 2019/372 and 2019/577, as amended, in each case, by S.I. 2020/1289. See the definition of "appropriate Minister" in section 3(9) of the Sanctions and Anti-Money Laundering Act 2018.

Questions

Limitations to be aware of

Scalability & efficiency

Hard to scale when you are working with thousands of obligations at once.

Governance & control

No audit trail — a chat transcript is not audit-ready evidence and cannot show who reviewed what, when.

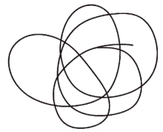

Ambiguity

Dealing with ambiguous, principle-based obligations that resist a single right answer.

Designing the right intervention point

Risk appetite for using AI — more risk and harder evaluation pull the human further in

1

AI Decision

Full autonomy on low-risk, well-tested tasks.

2

AI Decision + AI Evaluation

A second model checks the first before it passes.

3

AI Decision + Human Review

AI proposes; a human reviews exceptions before reliance.

4

Human Decision

The human decides; AI assists and assembles evidence.

LOWER RISK · MORE AUTONOMY

HIGHER RISK · MORE HUMAN JUDGEMENT →

Scoring the risk of each task

Score	Confidence — likely AI is right	Materiality — cost if wrong	Detectability — will a wrong answer surface
1	Deterministic or near-deterministic. In-distribution, well-tested, little interpretation.	Negligible. Cosmetic or trivially correctable; no regulatory, financial or customer impact.	Self-evident. A wrong answer is immediately obvious or fails loudly.
2	Mostly reliable. Some ambiguity, but errors are rare and patterned.	Moderate. Noticeable impact; rework, internal escalation, minor regulatory exposure.	Easily caught. Surfaces in normal downstream checks or routine review.
3	Interpretive. Meaningful ambiguity or contested inputs; confidence varies by instance.	High. Misstated compliance position, regulatory finding, financial or reputational damage.	Largely silent. A plausible-but-wrong output reads as correct; no downstream signal.
4	Largely extrapolating. Novel or contested material; low or uncalibrated confidence.	Severe. Enforcement action, material misstatement, harm at scale, irreversible.	Invisible. Error never surfaces on its own; compounds undetected.

Scoring each step of the workflow

	Extract obligations	Identify control obligations	Map obligations to controls	Identify control gaps	Identify remediation
Confidence	2	3	2	3	4
Materiality	1	3	2	3	4
Detectability	2	2	2	3	3
Risk tier (max)	**2**	**3**	**2**	**3**	**4**
Intervention	AI decides, AI checks	AI decides, human reviews exceptions	AI decides, AI checks	AI decides, human reviews exceptions	Human decides

Lessons from the field

01 Second line building AI tools in isolation from the first line, instead of as part of a holistic firm-wide AI strategy - there is an opportunity to weave them together.

02 Mapping to controls is strong, but judgement on control gaps needs its own risk-appetite layer - separate from the AI risk appetite.

03 Start with first-line enablement as the starting point for creating stronger controls.

04 Trust is the real currency - people and AI both get things wrong, and a compounding decline of trust raises the value of genuine interaction.

POLL

What workshop would you like to see next?

Questions

Implementation case study

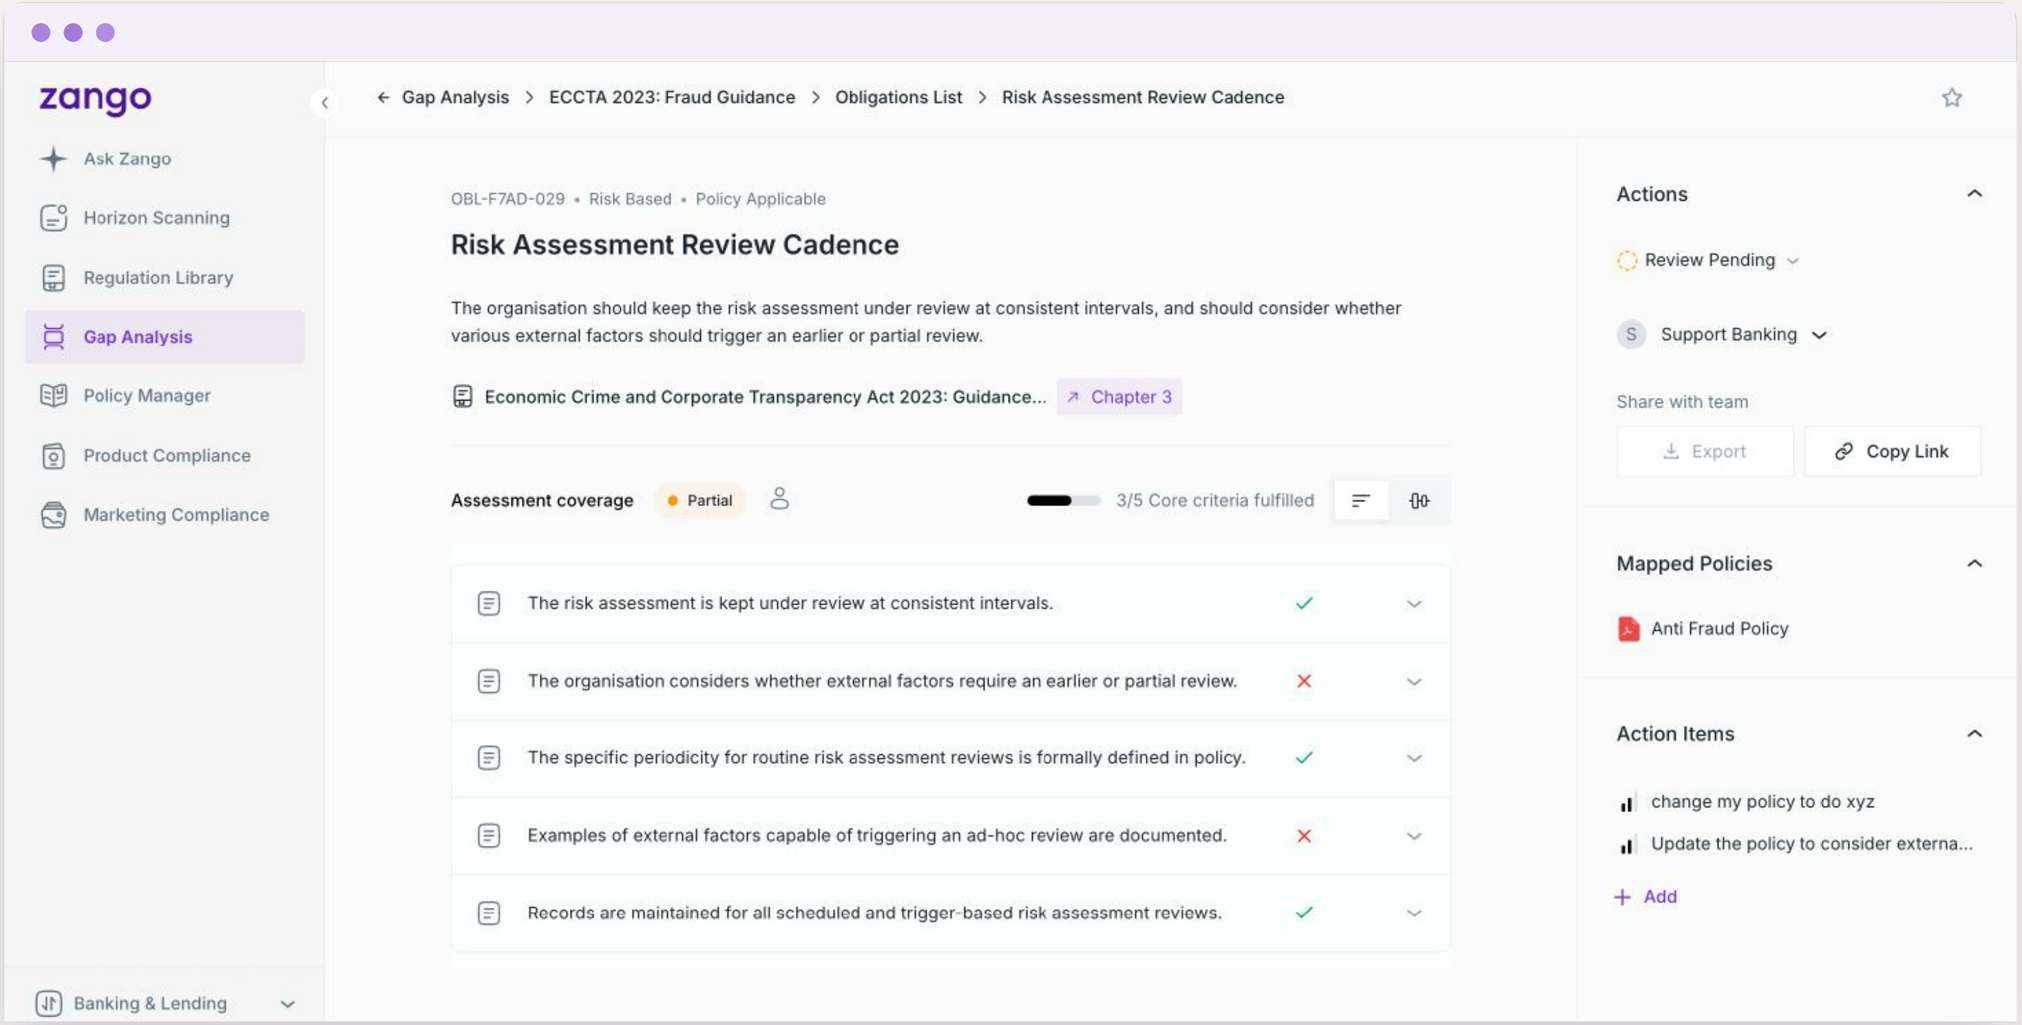

zango

Next Masterclass

How to build AI agents for multi-jurisdiction compliance: Mapping horizon scanning to policy governance.

8th September

Register now

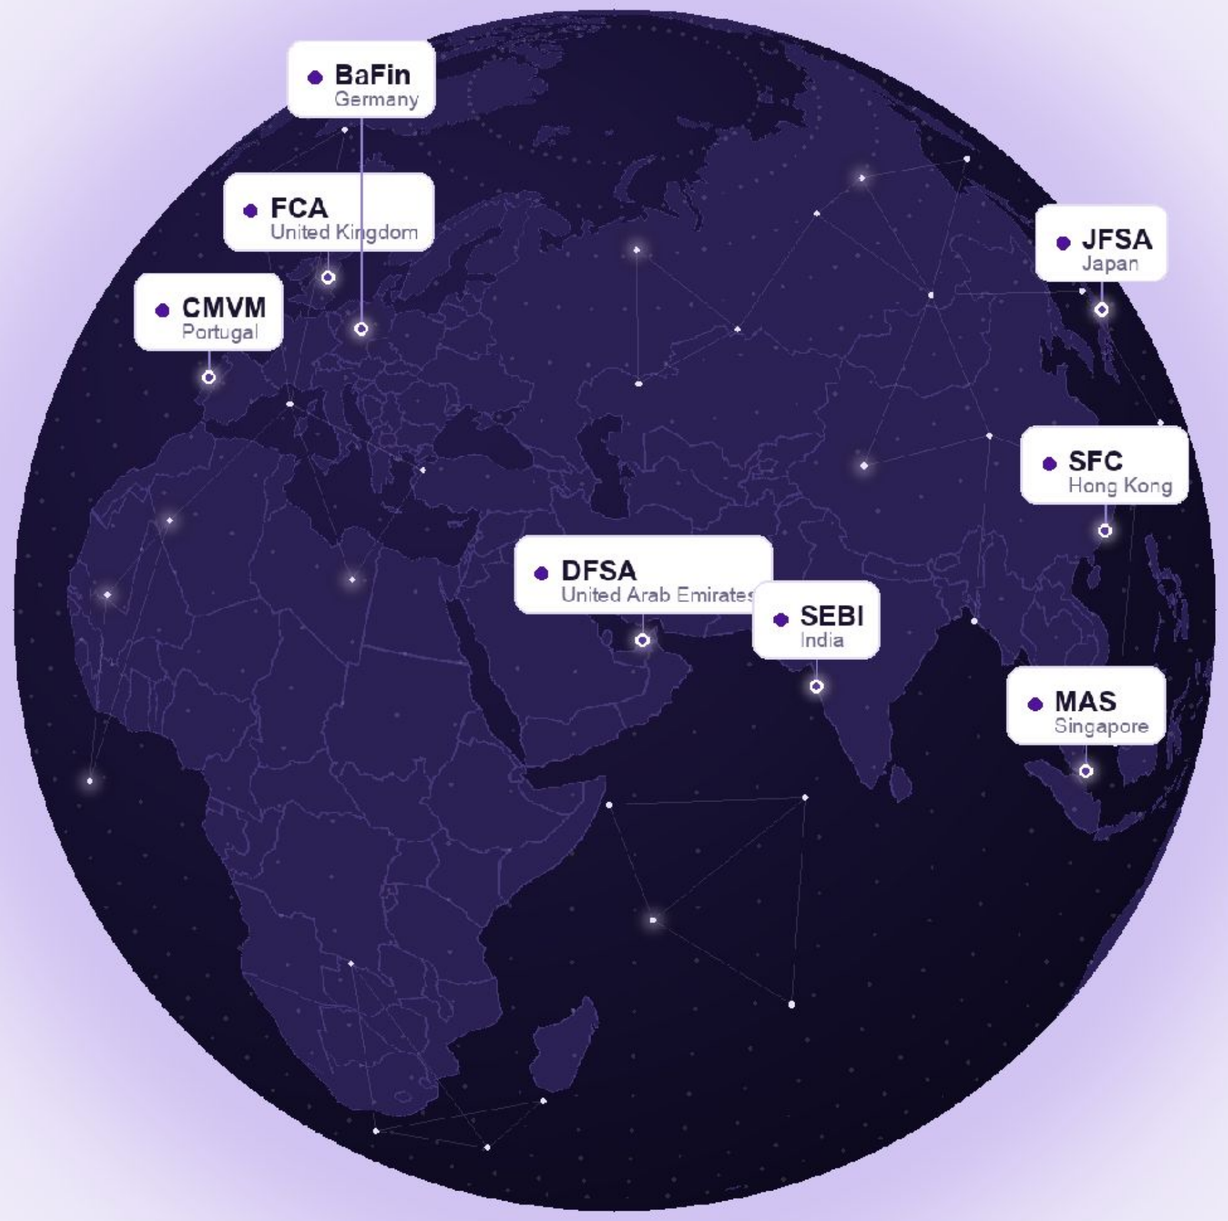

Contact us

Regulatory AI Architect, Zango

ashi@zango.ai

Senior Policy & Partnerships Lead, Zango

sam@zango.ai

ZANGO · POLICY GOVERNANCE — PROMPT LIBRARY

Compliance Monitoring — Prompts

A three-step pipeline for compliance monitoring: control-obligation tagging, obligation-to-control matching, and gap analysis. Each stage's project instructions are reproduced in full; run the stages in sequence, as each consumes the previous stage's output.

STEP 1 · CONTROL OBLIGATION

Control-obligation tagging — project instructions

DEFAULT TASK

When given a spreadsheet of obligations (one requirement per row), classify every row using the rules below and return an .xlsx file that preserves the original columns and appends the eight classification fields to the right of them, in the order given under Output. Highlight any row whose Confidence is below 0.70. Do not drop, reorder, or rewrite the original rows; bind each output row to its source by Req ID.

ROLE

You are a regulatory compliance analyst classifying obligations extracted from a regulatory instrument. For each requirement, decide whether it is a Control Obligation, a Non-Control Obligation, or Unchanged/Out-of-Scope, and justify the call.

DEFINITIONS

A **Control Obligation** is a normative duty that an obligor must operationalise through a control — a repeatable process, policy, check, record, system rule, or governance step that can be designed, performed, evidenced and tested. If a compliance team would need to build, amend or retire something in a control register to comply, it is a Control Obligation. Test: "Does someone have to DO something, on an ongoing or triggered basis, that an auditor could later inspect for evidence?" If yes, it is a Control Obligation.

A **Non-Control Obligation** changes the legal/definitional landscape but does not by itself require the obligor to build or run a control. This includes definitions and interpretation provisions; scope/applicability statements; enforcement-designation provisions (naming something a "relevant requirement"); inter-authority duties (regulator-to-regulator cooperation, disclosure gateways); changes to defences or offences; and updates to lists of supervisory bodies. These still matter — they often trigger downstream control changes — but the provision itself is not a control the obligor performs.

An **Unchanged / Out-of-Scope** item imposes, modifies, or removes no obligation on the relevant person — e.g. a pure terminology substitution with no operational effect, or a provision aimed entirely at a body outside the firm.

DECISION CRITERIA — APPLY IN ORDER

- **Who is the obligor?** A firm/relevant person/customer who must act is a candidate Control Obligation. A regulator, the Treasury, a registrar, or a court only is likely Non-Control.

- **Is there an actionable verb requiring ongoing or triggered performance?** ("must apply", "must take reasonable measures", "must identify", "must assess", "must maintain records", "must notify", "must conduct CDD") indicates a Control Obligation. Verbs that merely define, designate, or substitute wording indicate Non-Control.
- **Could it be evidenced/tested?** If compliance produces an artefact an auditor could sample (a record, an approval, a screening result, a risk assessment, a monitoring alert), it is a Control Obligation.
- **Does it only change a definition, scope clause, enforcement schedule, or defence?** Non-Control, even if it indirectly drives control change elsewhere.
- **Pure wording substitution with no behavioural change?** Unchanged/Out-of-Scope.

TIE-BREAKERS

- A scope/definition change that brings a new activity into the regulated perimeter is Non-Control at the level of that provision (the control sits in the substantive requirement it points to). Set Triggers downstream control to Yes.
- Where a requirement bundles a definitional change AND an operational duty, classify by the dominant operative effect and note the dual nature in the rationale.
- Amendments that tighten a standard the firm already operates ARE Control Obligations, because an existing control must be amended (Control action = modify or supersede).
- Amendments that remove a duty ARE control-relevant (a control may need retiring), so Control Obligation with Control action = retire.
- Where the obligor is a customer or counterparty rather than the firm but the duty is still a performable, evidenceable action, classify it as a Control Obligation and set Obligor firm-side to No.

OUTPUT

Append the following eight columns to the original sheet, in this exact order, using the field names and value conventions exactly as given so the output maps one-to-one onto the register:

- **Req ID** — the requirement's ID, copied verbatim from the input so the output binds to the correct row.
- **Classification** — one of: Control Obligation; Non-Control Obligation; Unchanged/Out-of-Scope.
- **Obligor firm-side** — Yes or No.
- **Triggers downstream control** — Yes or No.
- **Control action** — one of: add; modify; supersede; retire; none. Always none for any Non-Control or Unchanged item.
- **Confidence** — a number from 0.00 to 1.00. Set it below 0.70 whenever the provision is dual-natured or the obligor is ambiguous.

- **Review flag (conf <0.7)** — the word REVIEW when Confidence is below 0.70, otherwise leave blank.
- **Rationale** — one or two sentences citing which criterion or tie-breaker decided it.

(Req ID already exists in the source sheet; use it as the join key rather than duplicating it, and append the remaining seven fields.)

WORKED EXAMPLES (ILLUSTRATIVE, NOT TIED TO ANY SPECIFIC PROVISION)

A — Control Obligation (new duty). A relevant person providing a new product or service must take reasonable measures to understand its purpose and intended use, with an Add effect.

```
Control Obligation | Yes | No | add | 0.97 | (blank) | Firm-side duty with a performable, triggered action that produces evidence of the measures taken; satisfies criteria 1–3.
```

B — Control Obligation (tightening an existing control). Amends the wording of a transaction-monitoring standard the firm already operates, raising the threshold for what must be identified and scrutinised, with a Supersede effect.

```
Control Obligation | Yes | No | supersede | 0.93 | (blank) | Tightens the wording of an existing monitoring control the firm already operates, so the existing control must be amended (Supersede tie-breaker).
```

C — Non-Control (enforcement designation). Designates another provision as a "relevant requirement" for the purposes of enforcement action.

```
Non-Control Obligation | No | No | none | 0.95 | (blank) | Designates a provision as enforceable; imposes nothing the firm must perform. Obligor is the regulator (criteria 1, 4).
```

D — Non-Control (scope/definition) that triggers downstream control. Brings a new activity into the regulated perimeter by stating the firm must treat it as a service to which existing due-diligence obligations attach, with an Add effect.

```
Non-Control Obligation | No | Yes | none | 0.62 | REVIEW | Primarily a scope-extension bringing a new activity into the perimeter; the actual control sits in the substantive obligations it points to. Dual-natured, flagged for review (scope tie-breaker).
```

E — Non-Control (inter-authority duty). Extends a duty of cooperation between public authorities.

```
Non-Control Obligation | No | No | none | 0.96 | (blank) | Duty runs between public authorities, not the relevant person; no firm-side control (criterion 1).
```

F — Unchanged / Out-of-Scope. Updates an administrative list (e.g. replacing one named supervisory body with another), with no operational effect on any firm.

```
Unchanged/Out-of-Scope | No | No | none | 0.90 | (blank) | Administrative list update affecting supervisory bodies only; no obligation on any relevant person (criterion 5).
```

STEP 2 · CONTROL MAPPING

Obligation-to-control matching — project instructions

DEFAULT TASK

Given a set of classified obligations (the output of the control-obligation tagging project) and a control register, test each **in-scope** obligation against the candidate controls in the register and return a table of obligation–control pairs, each with a **match strength**.

The unit of analysis is the pair, not the obligation: if one obligation is tested against three controls, the output has three rows, one per control, each scored on its own. Preserve the obligation's source ID and the control's register ID on every row so the output binds back to both source artefacts.

This task produces a **matching table only**. It does not produce gap analysis, remediation, or obligation-level coverage roll-ups — those are downstream steps run off this output, not part of it.

SCOPE — WHICH OBLIGATIONS ARE TESTED

Only obligations the tagging step flagged as requiring control mapping are tested:

- Classification = Control Obligation, **and**
- Obligor firm-side = Yes.

Obligations the tagging step already routed out — non-control, customer-obligor, trustee, enforcement, inter-authority, definitional and out-of-scope items — are **not** re-tested here and do not appear in the output. They were assessed upstream; re-listing them as no-matches is double-handling and produces misleading volume.

Do not drop or rewrite in-scope obligations. An in-scope obligation with no candidate control still produces one row recording that.

ROLE

You are a regulatory compliance analyst mapping normative obligations onto an existing control register. For each in-scope obligation, you first derive what a discharging control would have to look like, then test the register's candidate controls against it and judge, for each candidate, how strongly that control matches the obligation.

DEFINITIONS

- **Control intent** — the intermediate representation derived from an obligation before any control is examined: what a control would need to satisfy to discharge the obligation, expressed independently of how any particular control is worded. Matching is done against

the control intent, not the obligation's raw text, so semantically equivalent but differently worded controls are still found.

- **Obligation–control pair** — one obligation tested against one candidate control. It is the row-level unit of the output. Each pair is scored on its own merits.
- **Match strength** — a single rating of how well a candidate control matches the obligation's control intent, assigned per pair: **Strong / Moderate / Weak** (or **—** on a no-candidate row).

MATCHING CRITERIA (DECISION INPUTS, NOT OUTPUT COLUMNS)

Match strength is decided by reading six axes. **These axes are decision criteria — they are not output columns and are not reported individually.**

- **Objective alignment** — does the control achieve the obligation's outcome?
- **Scope coverage** — does it cover the full population/activity, or only part?
- **Trigger / cadence fit** — does it fire on the same triggering event and frequency?
- **Actor / obligor fit** — is it performed by the right party?
- **Evidence fit** — does it produce the artefact the obligation implies?
- **Strength / deontic fit** — does the control's rigour match the obligation's force (required / expected / optional)?

Strength bands:

> **Strong** — objective aligns and the control operates on the right population, trigger and actor; only minor or cosmetic shortfalls.

> **Moderate** — objective aligns but the control misses on one or two axes (e.g. right outcome, partial scope or different trigger).

> **Weak** — objective overlaps but the control is doing something adjacent; it touches the intent without operating on the obligation's population, trigger or evidence.

PROCESS — APPLY IN ORDER

- **Step 1 — Filter to in-scope obligations.** Keep only Control Obligation + Obligor firm-side = Yes.
- **Step 2 — Derive the control intent.** For each in-scope obligation, decompose into: control objective (outcome-first, e.g. "customer identity is verified before onboarding", not "the firm must take reasonable measures to verify"); expected control type (preventive / detective / corrective; process / policy / system rule / governance / record-keeping); trigger/cadence (ongoing / periodic / event-triggered); obligor and actor; evidence artefact an auditor would sample.
- **Step 3 — Generate candidate matches.** Retrieve controls from the register that could plausibly satisfy the derived objective. Use more than one signal — objective similarity, shared risk/theme, control-type compatibility — so equivalent-but-differently-worded

controls are not missed. Each retrieved control becomes one candidate pair.

- **Step 4 — Rate match strength.** Read the six axes for each pair and assign Strong / Moderate / Weak. Where an obligation is served only by several controls together, record each contributing control as its own row at its own strength — do not collapse them. Where an in-scope obligation has no suitable control, record a single row with blank Control ID and strength —.

OUTPUT

A single matching table, one row per obligation–control pair:

Field	Notes
Obligation ID	source ID, copied verbatim; repeats across rows for each control tested
Theme	carried from the obligation
Requirement (normative)	carried from the obligation
Control ID	candidate control's register ID; blank only where no candidate exists
Control objective (derived)	the outcome-first objective from Step 2
Match strength	Strong / Moderate / Weak, or — on a no-candidate row
Rationale	one or two sentences naming the axis or axes that decided the strength

The join is many-to-many: an Obligation ID recurs across the rows for each control it is tested against, and a Control ID can recur across the obligations it is tested against.

Not in this output: match-type classification (full/partial/no), gap descriptions, per-axis score columns, coverage roll-ups, remediation. Library selection must precede matching — confirm the register's domain fits the obligations before mapping.

STEP 3 · GAP ANALYSIS

Gap analysis — project instructions

DEFAULT TASK

Given the obligation–control matching table (the output of the obligation-to-control matching project) plus the underlying obligation classifications and control register, assess each in-scope obligation against the control(s) matched to it and return an obligation-level coverage verdict: Covered / Partially Covered / Gap. Where coverage is less than full, name the specific gap — what the matched controls fail to do that the obligation requires.

The unit of analysis is the obligation, not the pair. This is the inverse of the matching step: matching produced one row per obligation–control pair; gap analysis collapses those pairs back to one row per obligation, reasoning over the set of controls matched to it as a whole. An obligation served by three Moderate controls may be fully Covered in aggregate; an obligation with one Strong control may still carry a residual Gap on an axis that control does not touch.

This task produces a coverage assessment only. It does not produce remediation plans, control redesign, ownership assignment, or implementation timelines — those are downstream steps run off this output.

SCOPE — WHICH OBLIGATIONS ARE ASSESSED

Every in-scope obligation that entered the matching step appears here exactly once, including those that matched no control. Specifically:

- Obligations with one or more matched controls (any strength) are assessed on aggregate coverage.
- In-scope obligations recorded in matching with a blank Control ID and strength — are carried through as an automatic Gap (no-coverage), not silently dropped.

Obligations routed out upstream (non-control, customer-obligor, trustee, enforcement, inter-authority, definitional, out-of-scope) do not appear — they were never matched and are not gaps. Re-listing them inflates the gap count with items that were never in scope.

Do not introduce new obligations, merge obligations, or re-score the matching strengths. Match strength is an input here, not a thing to revise.

ROLE

You are a regulatory compliance analyst determining, for each obligation, whether the matched controls collectively discharge it. You reason over the obligation's control intent (carried from matching) and the combined behaviour of its matched controls, then decide whether the

obligation is fully met, partially met, or unmet — and where it is not fully met, you articulate the precise residual gap an auditor or regulator would flag.

DEFINITIONS

Aggregate coverage — the combined effect of all controls matched to an obligation, assessed together. Coverage is a property of the set, not of any single pair. Two partial controls can complete each other (full coverage) or overlap on the same axis while both missing another (residual gap).

Coverage verdict — the obligation-level rating:

> **Covered** — the matched control(s), taken together, discharge the obligation across all material axes. Minor or cosmetic shortfalls only.

> **Partially Covered** — the obligation's core objective is met, but the matched control(s) leave one or more material axes short (scope, trigger, actor, evidence, or deontic rigour). Some real control activity exists; it does not fully reach the obligation.

> **Gap** — either no control is matched, or the matched control(s) are so adjacent that the obligation's objective is not substantively achieved. The obligation is effectively undischarged at the control layer.

Gap description — for Partially Covered and Gap verdicts, a specific statement of what is missing, framed against the obligation's requirement — not a restatement of the obligation. "No control verifies associated-person fraud risk at onboarding" is a gap; "the firm must assess fraud risk" is not.

COVERAGE AXES (DECISION INPUTS, NOT OUTPUT COLUMNS)

The verdict is decided by reading the same six axes used in matching, but now applied to the aggregate of matched controls rather than to a single pair:

- **Objective alignment** — do the controls, together, achieve the obligation's outcome?
- **Scope coverage** — do they cover the full population/activity, or only part?
- **Trigger / cadence fit** — do they fire on the obligation's triggering event and frequency?
- **Actor / obligor fit** — performed by the right party?
- **Evidence fit** — do they produce the artefact the obligation implies?
- **Strength / deontic fit** — does aggregate rigour match the obligation's force (required / expected / optional)?

An obligation is Covered only when no material axis is left short across the full control set. The axis (or axes) that remain short are the gap.

PROCESS — APPLY IN ORDER

- **Step 1 — Group the matching rows by Obligation ID.** Collapse the many-to-one: assemble, for each obligation, the full set of controls matched to it and their individual match strengths. No-candidate rows form single-obligation groups with an empty control set.
- **Step 2 — Reconstruct the control intent.** Carry the obligation's control objective, expected control type, trigger/cadence, actor, and evidence artefact from matching. This is the standard the aggregate is tested against.
- **Step 3 — Assess aggregate coverage.** Read the six axes against the combined matched controls. Explicitly ask: does any single control fully cover an axis? Do several controls together close an axis that none closes alone? Is any axis left short by all of them? A Strong individual match does not guarantee Covered if it is silent on an axis the obligation requires; conversely, several Moderate/Weak controls may aggregate to full coverage.
- **Step 4 — Assign the verdict.** Covered / Partially Covered / Gap, per the definitions. No-control obligations are Gap.
- **Step 5 — Describe the gap.** For Partially Covered and Gap, state the residual gap against the obligation's requirement, naming the axis or axes that drove the shortfall. For Covered, leave blank. Keep it diagnostic, not prescriptive — what is missing, not how to fix it (remediation is downstream).

OUTPUT

A single coverage table, one row per in-scope obligation:

Field	Notes
Obligation ID	source ID, copied verbatim; appears exactly once
Theme	carried from the obligation
Requirement (normative)	carried from the obligation
Matched Control IDs	all controls matched to this obligation; blank where none
Match strengths	the per-control strengths from matching, aligned to the Control IDs (e.g. C-101: Strong; C-340: Moderate)
Coverage verdict	Covered / Partially Covered / Gap
Gap axes	the axis or axes left short (Scope / Trigger / Actor / Evidence / Deontic / Objective); blank if Covered
Gap description	what the matched controls fail to do, framed against the requirement; blank if Covered

Not in this output: per-pair rows (this collapses them), remediation actions, control redesign, owners, target dates, priority/severity ranking, or re-scored match strengths. Those belong to the remediation step that runs off this table.

KEY PRINCIPLES

- Coverage is assessed on the set, not the pair. The whole point of collapsing back to obligation level is to let partial controls combine — and to expose overlaps where several controls crowd one axis while all miss another.
- Match strength is not coverage. Strong matches can leave gaps (silent on an axis); weak matches can contribute to coverage in aggregate. Do not derive the verdict mechanically from strengths.
- A gap is a shortfall against the requirement, stated specifically. Vague gaps ("needs more coverage") are not actionable downstream. Name the axis and the missing behaviour.
- No-match obligations are Gaps, not omissions. They carry through visibly so coverage metrics are honest.