

YOUR SPEAKER

## Romanos Kapsalis

AI Engineer code4thought®

 [romanos@code4thought.eu](mailto:romanos@code4thought.eu)

SESSION Thematic Noon: Automated auditing & inspection of AI systems under current regulatory frameworks.

HOST IIA Greece – Institute of Internal Auditors

WHEN 15 – 06 – 2026



A PLAIN-LANGUAGE BRIEFING FOR INTERNAL AUDITORS

# Auditing **the AI** you didn't build.

*What an AI audit is, why your office needs one, and how iQ4AI makes it possible.*

## TODAY'S ROADMAP

# Four things you'll **walk away with.**

No AI jargon — just what your office needs to know, and what it can do.

/ 01

## The problem

AI is already making decisions in your jurisdiction. Where — and why it matters.

/ 02

## What “auditing AI” means

The same discipline you already use ; applied to algorithms. 3 questions every audit answers.

/ 03

## iQ4AI

How non-engineers run a rigorous AI audit. No code, no maths, no ML team.

/ 04

## Live demo

An internal promotions model, audited end-to-end on stage. Fair? Accurate? Both?

*By the end, you'll have concrete tools —* **not just slides.**

THE PROBLEM YOU'VE ALREADY INHERITED

# Decisions in your jurisdiction **are being made by algorithms.**

Often quietly. Often by third-party vendors. And — until very recently — **nobody was auditing them the way we audit every other public-sector decision.**

/ WELFARE

## Benefits & fraud

Eligibility, risk-scoring of recipients, automatic flagging.

/ JUSTICE

## Risk assessment

Bail, sentencing recommendations, pre-trial detention.

/ HEALTH

## Triage & access

Care prioritisation, waiting-list ordering, screening.

/ EDUCATION

## Allocation

School placement, exam grading, scholarship awards.

*If a human caseworker did this, you'd audit the file.*

**The algorithm deserves the same scrutiny.**

## FIRST — WHAT IS AN “AI MODEL”?

# It's a **rule of thumb** learned from old cases.

A casework supervisor reviews ten thousand past files and notices patterns: **which factors tended to predict which outcomes**. They turn that hunch into a checklist new staff can apply.

An AI model is the same idea — automated. It studies past examples, finds patterns, and reproduces them on new cases at scale.

---

**NOT MAGIC.** NOT UNDERSTANDING. NOT JUDGEMENT. JUST STATISTICS ON PAST DATA.

## / STEP 01 — THE OLD FILES

## Historical cases

Past applicants, past inspections, past outcomes — whatever the office has on file.



## / STEP 02 — THE PATTERN

## Pattern-finding

The model fits a rule: “cases that look like  $X$  tended to result in  $Y$ .”



## / STEP 03 — THE DECISION

## Applied to new cases

Same rule, applied — automatically — to every new file that arrives.

## WHY A LEARNED RULE OF THUMB CAN FAIL

# Three ways an AI system **goes wrong** — even when the maths is “correct”.

/ 01 — Bad mirror

## It learned from biased history.

The past it studied was already unequal. The model faithfully reproduces that inequality — at scale, and with the air of objectivity.

**For example** — *past inspectors flagged certain neighbourhoods more aggressively, the model learns to do the same.*

/ 02 — Sneaky proxies

## It found a stand-in for a forbidden trait.

Even if you remove protected attributes like ethnicity, the model can rediscover them through a postcode, a school, a surname — and decide based on that.

**For example** — *a “networking score” that quietly correlates with gender, doing the work the original variable is no longer allowed to do.*

/ 03 — Drift

## The world changed, the model didn't.

It was trained on the population of five years ago. Today's applicants don't look like that population — but the model is still applying yesterday's rule.

**For example** — *a fraud detector tuned to pre-pandemic spending suddenly misclassifying ordinary post-pandemic behaviour.*

IT HAS ALREADY HAPPENED – IN PUBLIC-SECTOR CONTEXTS

# Three cases where an **unaudited algorithm** caused public harm.

/ NETHERLANDS · 2013–2019

## Childcare-benefits scandal

Tax authority · risk-scoring of fraud applicants

A risk-scoring algorithm flagged families — disproportionately those with **dual nationality or non-Dutch surnames** — as likely fraudsters. Benefits were clawed back. Families were ruined. The cabinet eventually resigned over it.

FAMILIES WRONGLY ACCUSED

≈ **26,000**

/ UNITED KINGDOM · 2020

## A-level grading algorithm

Department for Education · exam grade prediction

With exams cancelled during COVID, an algorithm assigned final grades. It **systematically downgraded students from state schools** while leaving private-school grades largely intact. The decision was reversed within days under public pressure.

GRADES DOWNGRADED

≈ **40 %**

/ UNITED STATES · 2016 ONWARDS

## Criminal-justice risk scoring

State courts · pre-trial risk assessment

An investigation found a widely-used risk-scoring tool was **nearly twice as likely** to incorrectly flag Black defendants as future re-offenders, and to incorrectly label white defendants as low risk. The tool was already in use across multiple states.

BIAS RATIO (FALSE FLAGS)


≈ **2× higher**

*/ In every case, an audit at deployment would have caught it. None was performed.*

/ 02

# What “auditing AI” actually means.

The same discipline you already use — applied to algorithms.



SO — WHAT DOES “AUDITING AI” ACTUALLY MEAN?

# The **same audit** you already do. For an algorithm.

You already know how to audit a process: **look at the inputs, examine the procedure, test the outputs, and check it against the rules.**

An AI audit is exactly that — applied to a decision-making system that happens to be a piece of software instead of a casework team.

**Same evidence-based discipline.** *Different subject under audit.*

/ TRADITIONAL AUDIT ASKS	/ AN AI AUDIT ASKS
Are the inputs trustworthy?	Is the training <b>data</b> representative & current?
Was the procedure followed?	Does the model <b>perform</b> as claimed?
Are people treated consistently?	Are outcomes <b>fair</b> across groups?
Can the decision be justified?	Can the model's reasoning be <b>explained</b> ?
Is there a paper trail?	Is there a versioned, <b>traceable</b> record?
<b>FAMILIAR DISCIPLINE</b> → APPLIED TO A NEW SUBJECT	

## THE THREE QUESTIONS EVERY AI AUDIT MUST ANSWER

# Whatever the system, the audit boils down to **three questions.**

/ Question 01

## Does it work?

PERFORMANCE

Forget the marketing. On real cases with known outcomes — how often is it right, how often is it wrong, and what kinds of mistakes does it make?

---

also called · accuracy · reliability · robustness

/ Question 02

## Does it treat people equally?

FAIRNESS

Across the people who pass through it — women and men, citizens and residents, regions, ethnic groups — are outcomes similar, or is one group consistently worse off?

---

also called · equity · disparate impact · non-discrimination

/ Question 03

## Can it be explained?

EXPLAINABILITY

When the system says “no” to a citizen, can we tell them *why* — in terms a person, an auditor, and a court would find acceptable?

---

also called · transparency · traceability · defensibility

*/ Every other technical term you'll hear is a tool to answer one of these three.*

## A PRACTICAL COMPLICATION FOR PUBLIC-SECTOR AUDITORS

Most of the AI in your jurisdiction is a **black box**.

/ TYPE A

## Transparent system

You — or your IT team — built it, or you have full access. You can open the model, see the data, examine the rule it learned.

/ Typically: in-house tools, open-source models, research deployments

**Audit posture:** *the full picture is available — perform a complete audit.*

/ TYPE B

## Black-box system

A vendor supplies the system. You feed in cases, you receive decisions. The model itself, the training data, the inner workings — sealed.

/ Typically: procured platforms, SaaS scoring tools, API services

**Still auditable.** What you see going in and coming out is enough to do real, rigorous evaluation.

VS

*The common excuse — “we can’t audit it, the vendor won’t share the model” — is no longer true. You don’t need the model. You need the predictions and the demographics.*

WHAT YOU BRING TO AN AI AUDIT

# Three ingredients — in plain words.

Not all are required. **The art of the audit is choosing the right approach for what you actually have.**

/ 01 — Always needed

## The cases the system has seen

A spreadsheet of the applicants, files, or events the AI was asked to decide on — **without the names**, but with the relevant attributes.

REQUIRED

/ 02 — Always needed

## What the system decided

The output for each of those cases — the verdict, the score, the risk band. **This is what the audit examines.**

REQUIRED

/ 03 — For fairness audits

## Who is in the data

The demographic attributes — age band, gender, region — needed to test **whether outcomes vary by group.**

FOR FAIRNESS

*/ If you have the first two, an audit is already possible.*

# This is moving from **good practice** to **obligation**.

A few years ago, auditing an algorithm was a discretionary internal-controls choice. **It is becoming a legal one** — and the public sector is in the first wave of scope.

We'll keep this brief; the point is simply that the question has shifted from *“should we?”* to *“when, by whom, and to what standard?”*

## EU AI Act

/ APPLIES IN YOUR JURISDICTION

Classifies most public-sector AI as **high-risk**. Mandates documentation, human oversight, post-market monitoring — and assessments akin to an audit.

## OECD AI Principles

/ REFERENCE FRAMEWORK

Endorsed by member states. Establishes **transparency, accountability, robustness** as principles governments are expected to apply to their own AI.

## National guidance

/ VARIES BY COUNTRY

Algorithm registers, impact-assessment requirements, supreme-audit-office reviews — emerging in most EU member states and the UK.

FROM OBLIGATION TO CAPABILITY

# Required. So **how** does your office do it?

Most public-sector offices don't have ML engineers on staff and can't afford external consultants per system. **Three practical paths forward.**

/ Path 01

## Build it yourself.

Hire data scientists, build the tooling, train the team. **Years of investment** and ongoing maintenance — and a tough budget line to defend to the board.

Slow · expensive · hard to staff

/ Path 02

## Hire it out, per audit.

External consultants per engagement. **Per-system cost**, slow turnaround, and the institutional knowledge walks out the door at the end of every project.

Costly · slow · no retention

/ Path 03

## A platform built for auditors.

A tool that **speaks your language**, runs the analysis under the hood, and produces evidence your office can sign off on. **No ML team required.**


Fast · evidence-based · scales

*/ The rest of this session walks through that third path.*

/ 03

# iQ4AI — the platform.

How non-engineers run a rigorous AI audit. No code, no maths.



WHERE WE COME IN

# iQ4AI is the platform that lets **non-engineers** run a rigorous AI audit.

Built for the people who sign off on public-sector decisions — internal auditors, compliance officers, oversight bodies. **You don't need to write code, retrain the model, or understand the maths.** You bring what you have; the platform does the rest.

/ 01

## Built for auditors

No machine-learning background required. The interface speaks the language of governance — protected attributes, impact ratios, traceable findings.

/ 02

## Works on what you have

Three audit modes that adapt to your level of access — from full transparency down to vendor predictions only.

/ 03

## Evidence-based output

Versioned reports, audit trail, exportable findings — designed to defend a decision in front of a regulator, a citizen, or a court.

WHERE MODELS, DATASETS, AND AUDITS ALL LIVE

# A **project** is the question — and everything orbiting it.

A project pins down **what an AI system is supposed to be deciding** — the task, the inputs, the outputs, the protected groups. Then it gathers every model, every dataset, and every audit run against that one decision.

## / INSIDE ONE PROJECT

<b>Problem</b>	What is this system deciding? — the question in <b>one sentence</b> .
<b>Setup</b>	Input shape, output shape, <b>protected attributes</b> , fairness thresholds.
<b>Models</b>	The artifact(s) being audited — <b>one or several versions</b> over time.
<b>Datasets</b>	The cases the audit runs against — and, where available, what actually happened.
<b>Audits</b>	Every evaluation and every report — <b>versioned and attributable</b> .

● PROJECT
PRJ-2026-014

## Benefits eligibility — tier 1

/ HIGH-RISK · WELFARE ALLOCATION

---

PROBLEM	Classify applicants as eligible / not eligible
TASK	Binary classification
PROTECTED	Region · Age band · Gender
THRESHOLD	4/5ths rule (0.80)
OWNER	Internal Audit · Cyprus

3  
DATASETS

2  
MODELS

5  
AUDITS

v4  
CURRENT

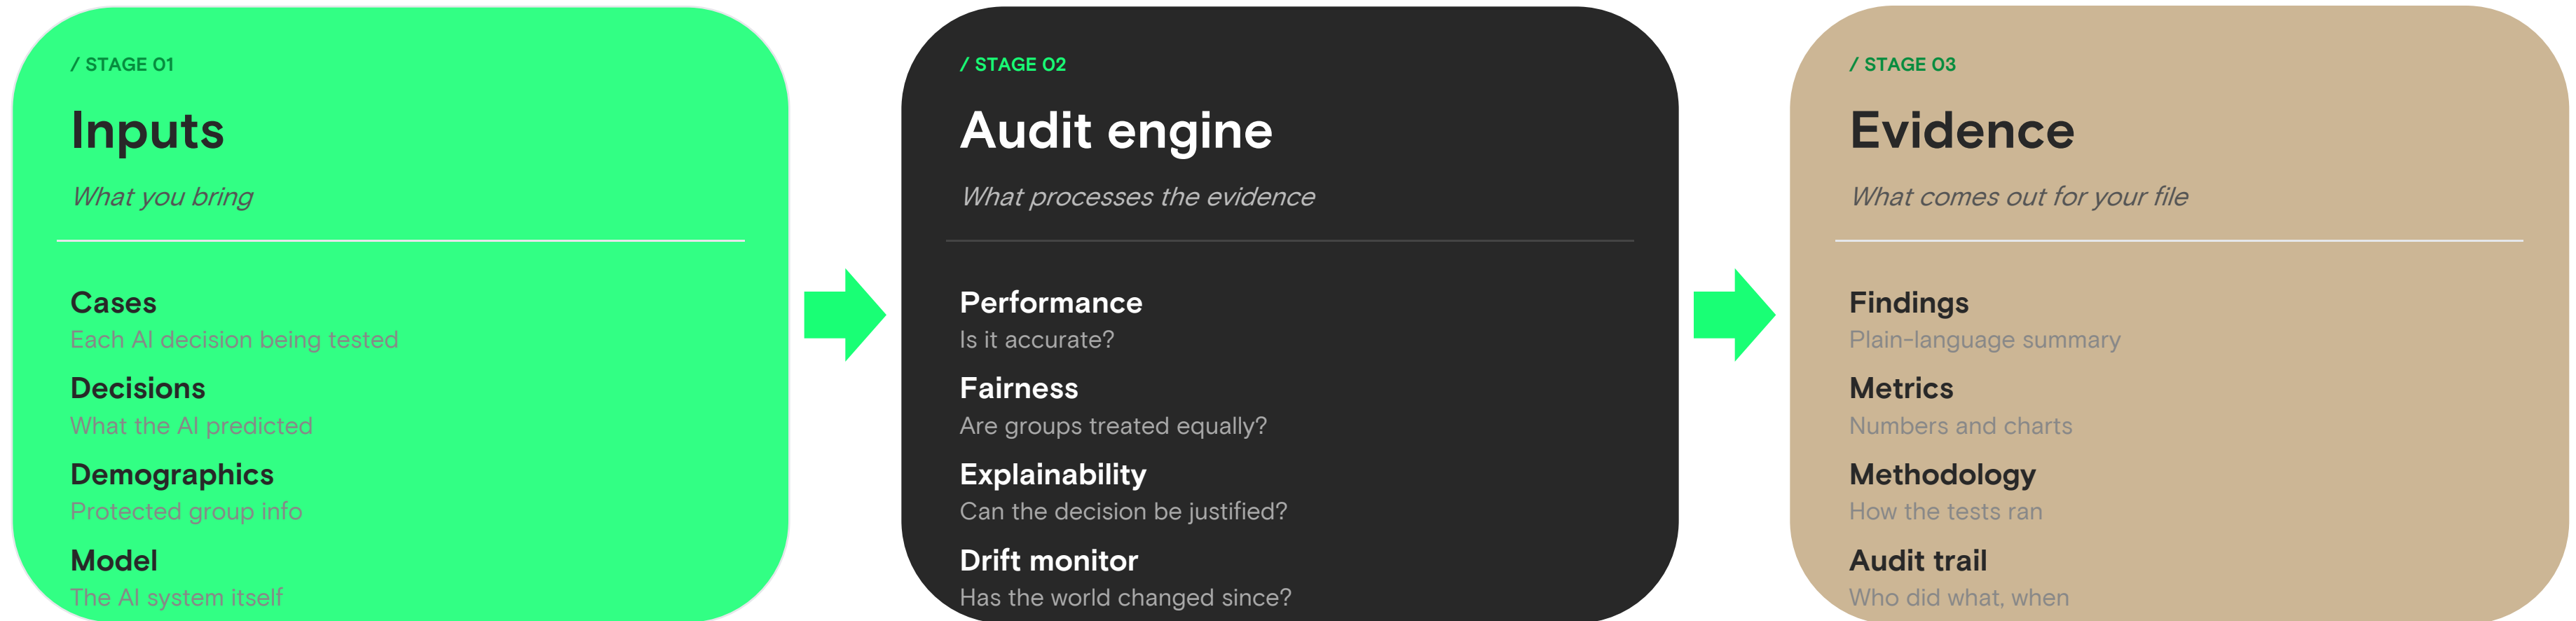
ALL ARTEFACTS · VERSIONED
LAST AUDIT · 2 WEEKS AGO

*/ One project per decision you're auditing. Everything else hangs off it.*

/ UNDER THE HOOD — END-TO-END FLOW

# Three stages: files in, tests run, evidence out.

How any project flows through the platform — with one governance layer running under all three stages.



/ CROSS-CUTTING — UNDER EVERY STAGE

### Built around the stack — governance that every layer relies on.

**Versioning**

Every audit re-runnable on the exact same inputs.

**Access and roles**

Configurator, reviewer, sign-off — separated.

**Audit log**

Every action attributable, time-stamped.

**Integrations**

CSV in. PDF and Excel out. APIs for handoff.

*Same pipeline for every audit. Auditor-friendly inputs in, defensible evidence out.*

COMPONENT · MODELS & PREDICTIONS

# Bring the **model** — or just the **predictions**.

/ 01 · UPLOAD THE MODEL

## You hand us the trained model.

iQ4AI loads it, runs inference on the testing set, and generates predictions for the audit automatically.

<b>ONNX</b> ./ promotion_model_v3.onnx	4.2 MB
FRAMEWORK	XGBoost · ONNX 1.16
INPUTS	11 features
OUTPUT	probability
STATUS	● loaded



iQ4AI  
AUDIT  
ENGINE

/ 02 · PREDICTIONS ONLY

## Or just send the output.

When the model is a black box — a vendor tool, an API, a closed system — predictions alone are enough.

PREDICTIONS.CSV · VENDOR OUTPUT		
CANDIDATE	LABEL	PROB
c_0001	REC	0.871
c_0002	REC	0.624
c_0003	PASS	0.318
c_0004	REC	0.752

/ KEY PRINCIPLE

You don't always need the model. If a vendor provides **only predictions**, iQ4AI can still audit fairness and output quality.

/ TYPE 01

### Labels

class · 0 / 1 · A / B / C

Classification audits. The minimum viable signal.

/ TYPE 02

### Probabilities

confidence · 0.00 – 1.00

Calibration & threshold analysis unlock with confidence scores.

/ TYPE 03

### Continuous

numeric score · regression

Regression & scoring models. Full distributional analysis.

## COMPONENT · DATA INTAKE

# What goes in, and what the platform checks.

Every dataset is logged, validated, and ready for review — before any model touches it.

/ 01

## Source recorded

Filename, upload date, and uploader are captured automatically. Full audit trail from day one.

/ 02

## Data quality checked

Missing values, duplicates, and unexpected formats are flagged before processing begins.

/ 03

## Ready for review

A plain-language summary of every column is generated — no technical knowledge needed to inspect.

## EXAMPLE · WHAT YOU SEE AFTER UPLOAD

### promotion\_history.csv

3,847 employee records · 12 fields · Uploaded 12 / 03 / 2026 by J. Papadopoulos

## WHAT'S IN THE FILE

## WHAT IT MEANS

<b>Employee ID</b>	A unique number for each person — never the actual name.
<b>Gender</b>	Recorded as Male / Female / Prefer not to say.
<b>Years at company</b>	How long the employee has worked here.
<b>Department</b>	Which business unit the employee belongs to.
<b>Performance rating</b>	Annual rating from manager review (1 to 5).

● Validated · Ready for audit review

## DATA QUALITY

- Complete · all 3,847 records present
- Complete · 0 missing values
- Complete · range 0–34 years
- Complete · 8 distinct departments
- 1.2% missing · flagged for follow-up

## ● Traceability

Every dataset has a logged source, date, and owner.

## ● Transparency

Quality issues surface before, not after, analysis.

## ● Reproducibility

Same input always yields the same audit record.

# The Audit Triad Matrix

Three complementary audit pathways — each defined by a distinct purpose, catalyst, and output.

01

## Assessment

### PURPOSE

End-to-end model quality assessment.

### CATALYST

Requires ground truth.

### KEY OUTPUT

Performance, fairness, explainability.

02

## NYC Bias Law Audit

### PURPOSE

Compliance with laws and regulations such as NYC Local Law 144.

### CATALYST

Scoring / selection rates and impact ratios.

### KEY OUTPUT

Official PDF report · impact ratios across gender / ethnicity.

03

## Model Output Analysis

### PURPOSE

Broad output-level statistical analysis.

### CATALYST

Model outputs (no ground truth required).

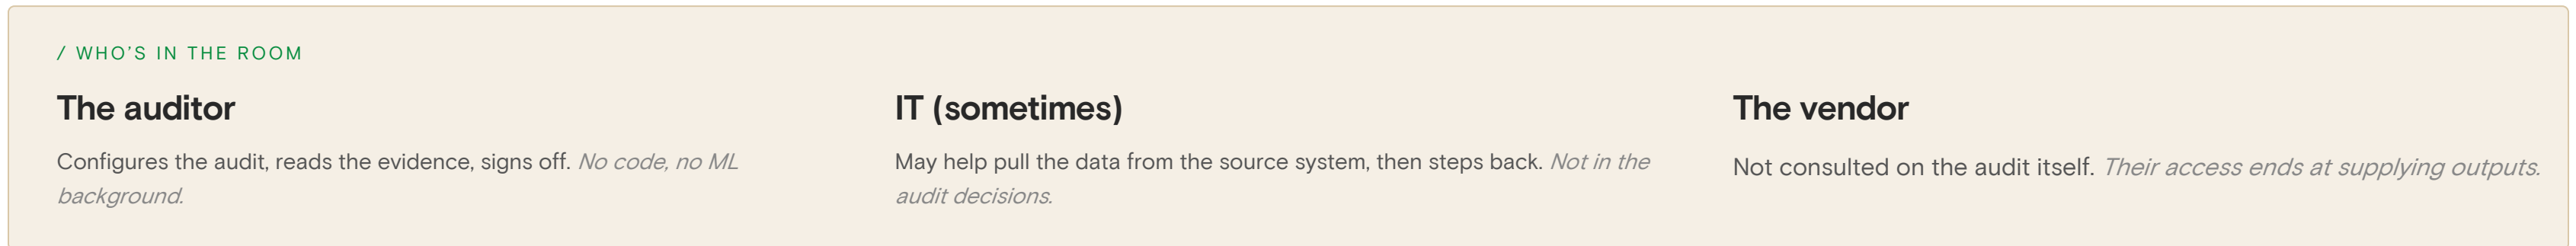
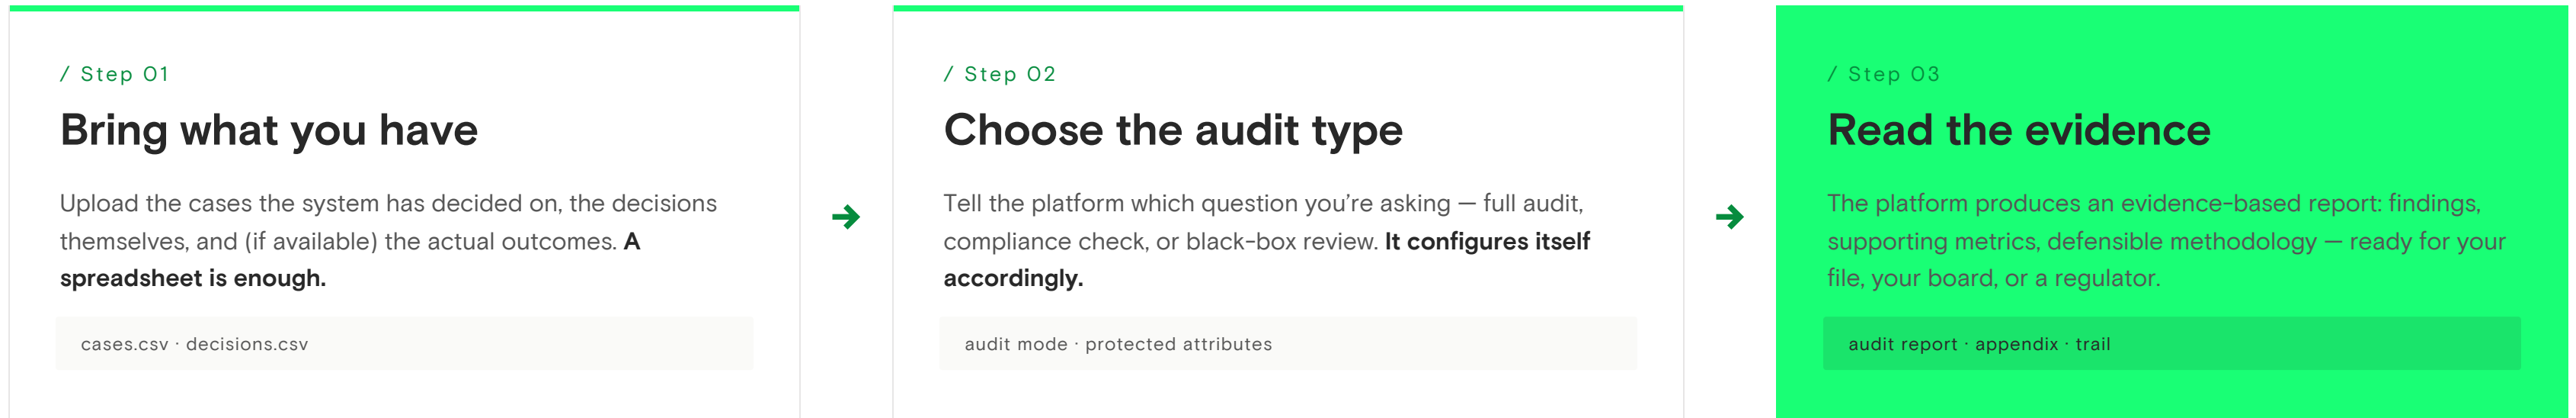
### KEY OUTPUT

14+ modular checks — drift, disparity, proportions.

Note · the EU AI Act audit architecture is reserved for future compliance routing.

HOW iQ4AI WORKS — FROM THE AUDITOR’S SEAT

# Three steps. No code.



*/ THE PROMISE* You bring the **oversight question**. iQ4AI brings the **evidence**. The audit opinion is yours.

## WHAT AN IQ4AI AUDIT GIVES YOU

# An **audit file** — not a black-box answer.

The output is built to live in your file, your committee pack, and your reply to the regulator — not to disappear into a dashboard nobody can quote from.

## / 01 Findings, in writing

Each conclusion stated plainly, with the evidence and the threshold it's measured against.

## / 02 Supporting metrics

The numbers behind the finding — charts, tables, breakdowns by group, exportable.

## / 03 Methodology & parameters

The exact configuration used — thresholds, protected attributes, dataset versions. Re-runnable.

## / 04 Audit trail

Who configured it, when, with what inputs. Every action logged and attributable.

● AUDIT REPORT
№ PS-2026-014

/ Verdict

### Adverse impact detected on region & age

SUBJECT	Benefits-eligibility scoring tool v4
AUDIT TYPE	Compliance check
SELECTION RATE	71.2% / 52.8%
IMPACT RATIO	0.74 · below 0.80 threshold
PERFORMED BY	Internal Audit · S. Marinov
METHODOLOGY	4/5ths rule · norming applied

Reproducible · versioned
Sign-off pending

Filed

A PRACTICAL STARTING POINT FOR YOUR OFFICE

# A four-step **AI audit programme** any office can stand up.

## 01

/ INVENTORY

### List the algorithms in your jurisdiction.

Most offices don't have this list. Start with procured systems — the contracts already say what they do. **You can't audit what you can't see.**

## 02

/ TRIAGE

### Classify each one by risk.

Decisions affecting benefits, justice, health, employment, education are **high-risk under the EU AI Act** — they go first. Lower-stakes tooling can wait.

## 03

/ AUDIT

### Run an evidence-based audit.

Pick the audit type that matches your access. Document findings in your file the same way you would any other audit. **This is where iQ4AI fits.**

## 04

/ MONITOR

### Re-audit on a cycle.


Models drift, populations change, vendors update without telling you. **An annual re-audit is the floor** ; high-risk systems may need quarterly review.

———— / Start with one system. Build the muscle. Scale from there. ————

/ 04

**Live demo** — see it in action.

An internal promotions model, audited end-to-end on stage.



# An internal promotions model. Fair and accurate?

We have full access to the model, the data, and the ground truth — so we run an Assessment.

/ THE SCENARIO

## Internal promotions model

BINARY CLASSIFICATION

/ FULL ACCESS — WHY WE CAN ASSESS

- ✓ **The model**  
We can run it on any input.
- ✓ **The training data**  
We see what it learned from.
- ✓ **The ground truth**  
We know what really happened.

/ 13 FEATURES — WHAT THE MODEL USES

## Inputs to the promotion decision

PROTECTED ATTRIBUTES — WHAT WE TEST FAIRNESS AGAINST

gender age age\_group

BACKGROUND & ROLE

years\_experience education\_level department

PERFORMANCE SIGNALS

performance\_rating certifications projects\_completed training\_hours  
networking\_score leadership\_potential hours\_worked\_weekly

THE QUESTION · Is the model both accurate enough to defend — and fair across gender and age?

— Q&A SESSION

# Questions from the room.

Six questions from the live webinar — with the answers we gave.

# Πώς ελέγχεις ένα proprietary "black box" μοντέλο;

---

## OUR ANSWER

Μέσω black-box auditing: στέλνουμε ελεγχόμενα δεδομένα στο μοντέλο και αναλύουμε τι επιστρέφει. Δεν απαιτείται πρόσβαση στον πηγαίο κώδικα ούτε στα training data\* του παρόχου — αρκεί η πρόσβαση στο API του μοντέλου. Σχεδιάζουμε στοχευμένα test sets, αξιολογούμε fairness, robustness, performance και explainability, και τεκμηριώνουμε τη συμπεριφορά του μοντέλου με μετρήσιμα στοιχεία. Με αυτόν τον τρόπο επαληθεύουμε αν τηρούνται οι δεσμεύσεις του παρόχου — χωρίς να χρειαστεί να ανοίξουμε το κουτί.

\* Σημείωση: τα training data είναι απαραίτητα μόνο στην περίπτωση του explainability analysis

# Πόσο "non-technical" μπορεί να είναι ο ελεγκτής;

---

## OUR ANSWER

Η πλατφόρμα απευθύνεται και σε ελεγκτές, όχι μόνο σε data scientists. Δεν χρειάζεται απαραίτητα γνώση Python, στατιστικής ή Machine Learning (ML). Χρειάζεται όμως κατανόηση του τι ελέγχεις: business context, ρυθμιστικό πλαίσιο, διαχείριση ρίσκου. Η ελεγκτική εμπειρία είναι το ζητούμενο — εμείς αναλαμβάνουμε το τεχνικό κομμάτι μέσα από έτοιμες ροές εργασίας, προκαθορισμένες ρυθμίσεις και αυτοματοποιημένες αναφορές. Επιπλέον, η πλατφόρμα προσφέρει insights και recommendations σε απλή, κατανοητή γλώσσα, ώστε ο ελεγκτής να ερμηνεύει τα ευρήματα χωρίς τεχνικό λεξιλόγιο.

# Καλύπτει η μεθοδολογία σας και generative AI / LLMs;

---

## OUR ANSWER

Ναι, αλλά με διαφορετικό σετ ελέγχων. Τα κλασικά μοντέλα πρόβλεψης ελέγχονται για μεροληψία, ακρίβεια και υποβάθμιση της απόδοσης με την πάροδο του χρόνου. Τα generative μοντέλα (LLMs) απαιτούν επιπλέον ελέγχους: παραγωγή ψευδών πληροφοριών (hallucinations), κακόβουλη χειραγώγηση μέσω εντολών (prompt injection), διαρροή ευαίσθητων δεδομένων, ακατάλληλο περιεχόμενο και συνέπεια στις απαντήσεις. Αυτή τη στιγμή η πλατφόρμα καλύπτει πλήρως τα μοντέλα πρόβλεψης — το εξειδικευμένο module για generative AI βρίσκεται ήδη σε ανάπτυξη και αναμένεται σε επόμενη έκδοση.

# Η πλατφόρμα προτείνει διορθώσεις ή μόνο εντοπίζει ευρήματα;

---

## OUR ANSWER

Και τα δύο. Η πλατφόρμα δεν επισημαίνει απλώς τα ευρήματα, σε ορισμένα σημεία προτείνει ήδη συγκεκριμένες διορθώσεις, με μηνύματα σε απλή και κατανοητή γλώσσα, ώστε ο χρήστης να καταλαβαίνει τι πρέπει να αλλάξει και γιατί. Παράλληλα, βρίσκεται σε φάση ανάπτυξης η δυνατότητα δυναμικών προτάσεων διόρθωσης μέσω τεχνητής νοημοσύνης, που θα προσαρμόζονται ανά εύρημα. Και φυσικά, πέρα από το iQ4AI, η ομάδα μας είναι πάντα διαθέσιμη για να σας καθοδηγήσει σε κάθε βήμα.

# Πώς διασφαλίζουμε ότι το audit αξιοποιεί όλα τα data που εισάγουμε;

## OUR ANSWER

Με πλήρη ιχνηλασιμότητα και καταγραφή κάθε βήματος. Κάθε σύνολο δεδομένων που εισάγεται καταγράφεται αυτόματα — ποιος το ανέβασε, πότε, και σε ποιον έλεγχο αντιστοιχεί. Μέσα στην πλατφόρμα βλέπεις πόσες γραμμές και στήλες περιέχει, αν κάποιες εγγραφές αφαιρέθηκαν και για ποιον λόγο. Οι έλεγχοι αξιοποιούν πάντα το σύνολο των διαθέσιμων δεδομένων. Στις περιπτώσεις που η νομοθεσία επιτρέπει εξαίρεση κάποιων κατηγοριών — όπως π.χ. ο κανόνας του 2%\* στον Bias Audit Law (Local Law 144) της Νέας Υόρκης — αυτό γίνεται αποκλειστικά με απόφαση του χρήστη και με πλήρη, ξεκάθαρη ενημέρωση εντός της πλατφόρμας.

\* Η εξαίρεση του 2% επιτρέπει στους ελεγκτές να παραλείπουν δημογραφικές κατηγορίες που αποτελούν λιγότερο από το 2% των δεδομένων κατά τον υπολογισμό της αναλογίας επιπτώσεων.

# Συμβουλές για internal auditor με ελάχιστη γνώση AI;

---

## OUR ANSWER

Τρεις πρακτικές συμβουλές:

- (1) Ξεκίνα από τον επιχειρηματικό κίνδυνο, όχι από τα μαθηματικά — ποιες αποφάσεις παίρνει το μοντέλο και ποιον επηρεάζουν.
- (2) Μάθε τις βασικές έννοιες, όχι τους αλγορίθμους — τι σημαίνει μεροληψία, αξιοπιστία, δικαιοσύνη, ερμηνευσιμότητα.
- (3) Ξεκίνα με μια δοκιμή σε ένα μοντέλο χαμηλού ρίσκου. Ο πρώτος έλεγχος είναι πάντα ευκαιρία μάθησης — εμείς προσφέρουμε καθοδηγούμενες εκτελέσεις με πραγματικά σενάρια. Και ναι, είναι απολύτως εφικτό χωρίς τεχνικό υπόβαθρο.

— SCAN TO EXPLORE

Discover **iQ4AI** — our AI audit platform



— LET'S STAY IN TOUCH

Thank **you.**

code4thought®

[code4thought.eu](https://code4thought.eu)

[contact@code4thought.eu](mailto:contact@code4thought.eu)

Athens · Patras