From Acceleration to Assurance:

Governing AI-Assisted Tools in SOX and ICFR Audit Analytics

Karishma Velisetty, Independent Researcher

AGENDA

What We'll Cover Today

01

The AI-Assisted Audit Landscape

Tools in use today across three categories: code assistants, language AI, and specialist audit AI

02

Use Cases Across SOX Testing

Code review · Contracts · Invoice recon · Narratives · Test sheets · JE testing · UAR/SOD · Flux · Walkthroughs

03

SOX & ICFR Risk by Use Case

Specific risks, failure modes, and what external auditors scrutinize

04

Governance Framework

People, process, and technical controls to make AI-assisted work reliance-grade

05

Implementation Roadmap & Q&A

A phased approach you can start this quarter

AI Tools Already Operating in SOX Audit Shops

Code & Analytics Assistants

- GitHub Copilot
- Cursor
- Amazon Q Developer
- Tabnine

Script generation, query writing, debugging audit analytics code

Language & Document AI

- Claude (API / Claude.ai)
- ChatGPT / GPT-4
- Microsoft Copilot
- Google Gemini

Drafting narratives, reviewing contracts & policies, summarizing evidence

Specialist Audit AI

- AuditBoard/Optro AI
- TeamMate+ AI
- Workiva AI
- Caseware IDEA AI

Control testing, risk scoring, reconciliation, built-in SOX workflow integration

Where AI Is Being Applied in SOX & ICFR Testing Today

UC1
Code Review
Reviewing analytics scripts

UC2
Contract Review
Extracting SOX-relevant clauses

UC3
Invoice Reconciliation
3-way match & exception detection

UC4
Control Narratives
Drafting & updating narratives

UC5
Test Sheets
Test plans & workpaper generation

UC6
Journal Entry Testing
Full-population JE anomaly detection

UC7
UAR & SOD Analysis
Access review & conflict detection

UC8
Flux Analysis
Financial statement variance review

UC9
Walkthrough Docs
Interview capture & process docs

UC 01

AI-ASSISTED CODE REVIEW

Reviewing Audit Analytics Scripts with AI

How It Works in Practice

- Auditor writes or receives a Python / SQL script for JE testing, expense sampling, or SOD analysis.
- AI reviews the code for logic errors, edge cases (nulls, negatives, duplicates), and deviations from the intended test objective.
- AI flags issues and suggests corrections. Auditor evaluates and documents accepted / rejected suggestions.
- Reviewed code is committed via change management before use in SOX evidence.

SOX / ICFR Risks

- AI misses domain-specific financial logic errors
- No audit trail for AI review actions taken
- Over-reliance replaces expert judgment
- Code accepted without human validation

- ✓ Human reviewer must formally sign off on all AI review findings before code is approved.
- ✓ Log AI review session (tool, date, findings) as part of the change management ticket.
- ✓ Second-level review by CPA / audit lead required for all scripts used in SOX reliance analytics.

UC 02

AI-ASSISTED CONTRACT REVIEW

Reviewing Contracts & Agreements for SOX-Relevant Clauses

How It Works in Practice

- Auditor uploads revenue contracts, vendor agreements, or lease arrangements to an AI assistant.
- AI extracts key terms: performance obligations, termination clauses, variable consideration, effective dates, approval signatories.
- AI flags clauses relevant to ASC 606, ASC 842, or related-party thresholds for auditor review.
- Auditor validates AI extractions against the source document before documenting in workpapers.

SOX / ICFR Risks

- AI misreads clause context or legal nuance
- Confidential contract data sent to uncontrolled AI API
- Incomplete extraction leads to missed controls
- AI hallucination on complex ASC 842 terms

- ✓ Only use enterprise AI with a DPA (data processing agreement) for contract uploads. Consumer tools are prohibited.
- ✓ Classify contracts before selecting AI tool. MNPI-containing contracts must stay within approved on-premise or private deployments.
- ✓ Auditor must validate every AI-extracted clause against the source contract. Document as a workpaper procedure step.

UC 03

AI-ASSISTED INVOICE RECONCILIATION

Invoice Matching, Exception Detection & Reconciliation

How It Works in Practice

- AI ingests AP ledger, purchase orders, and scanned invoices to perform 3-way matching at scale.
- Exception rules (amount tolerances, duplicate vendor IDs, non-PO purchases above threshold) are parameterized and applied automatically.
- AI flags unmatched items, pricing variances, and split transactions potentially designed to circumvent approval thresholds.
- Results are piped into audit sampling selection or directly to the CAPA log for control deficiency tracking.

SOX / ICFR Risks

- Incorrect exception thresholds set or inherited from AI
- Undetected population gaps — incomplete coverage
- AI model drift increases false negative rate over time
- No documentation of matching logic — no external reliance

- ✓ Document matching logic, exception rules, and tolerance thresholds formally — this is audit evidence, not a config file.
- ✓ Confirm population completeness: reconcile record count from source ERP to AI tool input before relying on results.
- ✓ Re-validate exception rule effectiveness quarterly — AI rules degrade as transaction patterns evolve.

UC 04

AI-ASSISTED CONTROL NARRATIVES

Drafting & Updating SOX Control Narratives

How It Works in Practice

- Auditor provides AI with process documentation, prior-year narratives, and interview notes from process owner walkthroughs.
- AI drafts updated control narrative covering: process objective, control owner, frequency, evidence produced, and COSO linkage.
- AI cross-checks narrative consistency against the corresponding risk-control matrix (RCM) and identifies gaps.
- Auditor reviews draft, edits for accuracy, and routes through standard narrative approval workflow.

SOX / ICFR Risks

- AI fabricates plausible-but-inaccurate control steps
- Prior-year errors propagated unchanged into new narrative
- RCM-narrative misalignment not detected by AI
- Confidential process data sent to non-approved AI tool

- ✓ Process owner must formally validate the AI-drafted narrative — their signature confirms accuracy, not just formatting.
- ✓ AI-drafted narratives must be reviewed against the live RCM for completeness before approval. Cross-reference is a required step.
- ✓ Treat narrative drafting prompts and source documents as Confidential — use enterprise-only AI tools.

UC 05

AI-ASSISTED TEST SHEETS

Generating SOX Test Plans & Audit Workpapers

How It Works in Practice

- Auditor provides AI with the control description, risk assertion (completeness, accuracy, existence), and prior test template.
- AI generates test steps, sample size recommendation, evidence requirements, and pass/fail criteria aligned to the control objective.
- AI proposes attribute testing criteria and suggests exception thresholds based on materiality guidance provided.
- Auditor reviews and approves test sheet; it enters the standard workpaper review and sign-off workflow.

SOX / ICFR Risks

- AI test steps don't cover all relevant COSO assertions
- Sample size guidance not calibrated to entity risk appetite
- Test sheet accepted without senior reviewer sign-off
- Evidence requirements too generic for external auditor reliance

- ✓ Require senior auditor or manager sign-off on all AI-generated test sheets before testing begins. No self-approval.
- ✓ Verify COSO assertion coverage explicitly: completeness, accuracy, existence, valuation, rights & obligations.
- ✓ Calibrate AI-suggested sample sizes against your firm's sampling methodology. Document the basis for any override.

Four Emerging AI Applications in SOX Testing

UC6

Journal Entry Testing

Full-population JE anomaly & fraud detection

AI analyses 100% of journal entries for fraud risk indicators: round numbers, after-hours postings, unusual account combinations, and entries that circumvent approval thresholds.

UC7

UAR & SOD Analysis

User access review & conflict detection

AI cross-references user provisioning data against the SOD rule matrix to flag conflicts, identify orphaned accounts, and prioritize access for recertification.

UC8

Flux Analysis

Financial statement variance & analytical review

AI performs period-over-period and ratio analysis on GL data, drafts management explanations for unusual movements, and highlights items requiring substantive investigation.

UC9

Walkthrough Documentation

Interview transcription & process memo generation

AI transcribes walkthrough interviews, extracts control activities and evidence types, and generates structured walkthrough memos mapped to the SOX process documentation template.

UC 06

JOURNAL ENTRY TESTING & FRAUD ANALYTICS

AI-Assisted Full-Population JE Analysis

How It Works in Practice

- AI ingests the full JE population from the ERP (SAP, Oracle, Workday) and applies rules-based plus ML anomaly detection.
- Flags entries by risk indicator: round-dollar amounts, non-standard accounts, after-hours timestamps, entries near period-end cutoff.
- AI scores and ranks exceptions by risk level, generating a prioritized sample for auditor follow-up with preparer/approver details.
- Auditor investigates flagged items, documents conclusions, and links results to the JE testing control in the SOX workpaper.

SOX / ICFR Risks

- AI model tuned on wrong entity produces invalid risk scores
- Population completeness not verified before run
- High false positive rate dilutes auditor focus
- AI flags not individually documented in workpapers

- ✓ Validate population completeness: reconcile JE count from ERP source to AI tool input before relying on any outputs.
- ✓ Document all risk scoring criteria and thresholds formally — parameters are part of the control design and must be version-controlled.
- ✓ Auditor must document an investigation conclusion for every flagged item — not a disposition of 'reviewed by AI'.

UC 07

USER ACCESS REVIEW & SOD ANALYSIS

AI-Assisted UAR & Segregation of Duties Conflict Detection

How It Works in Practice

- AI ingests user role assignments from ERP, IAM, and privileged access systems, and maps them against the approved SOD rule matrix.
- Conflicts are identified at role-level and transaction-level; AI cross-references compensating control documentation to determine net risk.
- Orphaned accounts, dormant users, and over-provisioned super-user access are flagged for recertification or immediate remediation.
- Results are structured as a UAR evidence package with role-owner sign-off tasks routed for completion.

SOX / ICFR Risks

- Stale SOD rule matrix leads to false conflict clearances
- AI misclassifies compensating controls as mitigating
- Incomplete role provisioning data creates gaps in UAR coverage
- AI-generated UAR accepted without role-owner sign-off

- ✓ SOD rule matrix must be reviewed and approved by management at least annually before use in AI-assisted UAR.
- ✓ Role-owner formal sign-off is required for every access certification decision — AI can flag, but only a human can certify.
- ✓ Compensating control assessments generated by AI require independent validation by the controls owner before reducing the SOD risk rating.

UC 08

FINANCIAL STATEMENT FLUX ANALYSIS

AI-Assisted Period-over-Period Variance & Analytical Review

How It Works in Practice

- AI ingests TB/GL data across reporting periods and computes absolute and percentage variances at account and sub-account level.
- Variances exceeding materiality or percentage thresholds are flagged; AI generates draft explanations based on pattern matching with prior-period notes.
- AI performs ratio analysis (gross margin, DSO, inventory turns) and highlights movements inconsistent with operating trends.
- Auditor reviews flagged items, edits AI-drafted explanations, and documents final management representation in the SOX flux memo.

SOX / ICFR Risks

- AI-generated flux explanations accepted without CFO validation
- Threshold miscalibration masks material variances
- AI draft management rep letter treated as a final document
- AI pattern-matching on prior-period data propagates errors

- ✓ Reviewer must formally review and approve all AI-generated flux explanations before entering SOX documentation.
- ✓ Variance thresholds must be calibrated against current-year materiality and approved by the audit lead — not auto-inherited from prior year.
- ✓ AI-drafted management representation letters are working documents only. Final signed representations require the full management approval cycle.

UC 09

WALKTHROUGH DOCUMENTATION

AI-Assisted Interview Capture & Process Documentation

How It Works in Practice

- Walkthrough interviews are transcribed in real-time using AI (e.g., Teams Copilot, Otter.ai, or a dedicated audit tool).
- AI identifies control activity mentions, evidence types, frequency, and control owner references, and maps them to the process flow.
- AI generates a structured walkthrough memo: process objective, control activities, COSO linkage, evidence observed, and gaps noted.
- Auditor reviews the memo for accuracy, adds context from direct observation, and routes for process-owner confirmation.

SOX / ICFR Risks

- Transcription errors create inaccurate control descriptions
- Recording consent from process owners not obtained
- AI memo describes stated intent, not actual operating procedure
- Audio / transcript data stored in non-compliant AI system

- ✓ Obtain written recording consent from all interview participants before using AI transcription tools. Document consent in the workpaper.
- ✓ Use only approved, compliant AI tools for processing audio. Interview audio containing process details is Confidential.
- ✓ Process owner must formally validate the AI-generated walkthrough memo before it is used as SOX evidence. Verbal confirmation is insufficient.

Making AI-Assisted Work Reliance-Grade

People & Process

- Appoint AI Governance Owner
- Publish Acceptable Use Policy
- Mandatory AI literacy training
- Human sign-off at every gate
- Prompt & session documentation

Technical Controls

- Approved enterprise tool list only
- Data classification enforcement
- Change mgmt tagging for AI code
- Output validation test suites
- Version control for all AI scripts

Monitoring & Assurance

- Monthly AI usage log review
- Quarterly KRI dashboard update
- Control testing for AI procedures
- External auditor walkthrough
- Annual policy review & attestation

All three pillars must be operational before any AI-assisted work product can be used as SOX reliance evidence

From Acceleration to Assurance: The Bottom Line

1

AI-assisted tools now span the full SOX testing lifecycle — governance must match.

Nine distinct use cases are in active use. Each carries measurable, manageable risk. The governance gap is real and widening.

2

Confidentiality is the highest-stakes risk across use cases.

Consumer AI tools must be prohibited from receiving Confidential or MNPI data. Enterprise AI with a DPA is the minimum bar.

3

AI-drafted content requires formal human validation — not just review.

Narratives, test sheets, flux memos, and walkthrough memos require process-owner or senior-auditor sign-off confirming accuracy.

4

Thank you

Key References

- COSO Internal Control — Integrated Framework (2013 Edition)
- IIA Global Technology Audit Guide: Artificial Intelligence
- PCAOB AS 2201 — Audit of ICFR / AS 2301 — Audit Evidence
- Fed SR 11-7 Guidance on Model Risk Management
- AICPA SOC 2 Trust Services Criteria (TSP Section 100)