AI Governance:
Building the Right Oversight Structures

Matt Kelly
Radical Compliance
mkelly@RadicalCompliance.com

5 June 2025

RADICAL COMPLIANCE
SHARP THINKING ABOUT COMPLIANCE, AUDIT, AND RISK

Who is this guy?

- Writing, blogging independently at www.RadicalCompliance.Com
- Previously editor at Compliance Week, 2006-2016
- Doing various research, writing projects for private clients

Order of Events …

- What do we know about the risks of AI, compliance or otherwise?
- What should an AI risk committee do?
- What are the frameworks or other guidance that can be a roadmap?
- What are you going to monitor and report?
- **Q&A**

Part I:

What do we know about the risks of AI?

There are lots of them!

There are lots of them!

Regulatory

- Privacy laws
- Cybersecurity
- Anti-discrimination
- Consumer
- Explainability

Operational

- Fraud
- IP infringement
- Model drift
- Errors
- Poor testing, QA

Reputational

- Alienated customers
- Brand embarrassment
- Regulatory scrutiny
- Employee distrust

Let’s freak out more specifically, shall we?

- 49 percent of employees admit to using AI in ways that contravene employer’s policies.
- 60 percent say they have witnesses other employees using AI tools in wrong ways.
- 66 percent say they rely on AI output without evaluating the information it provides.

— ‘*Trust, Attitudes, and Use of AI*,’ KPMG, May 2025

Now consider <u>how</u> employees use AI

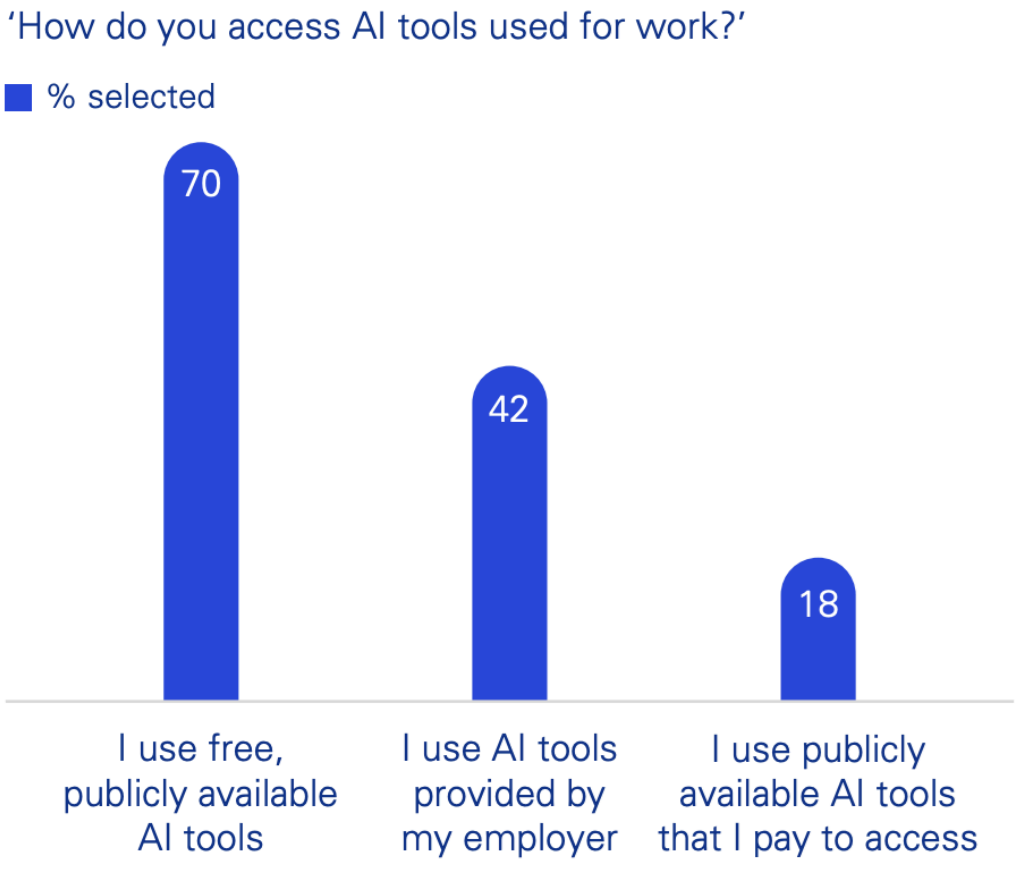

— '*Trust, Attitudes, and Use of AI*,' KPMG, May 2025

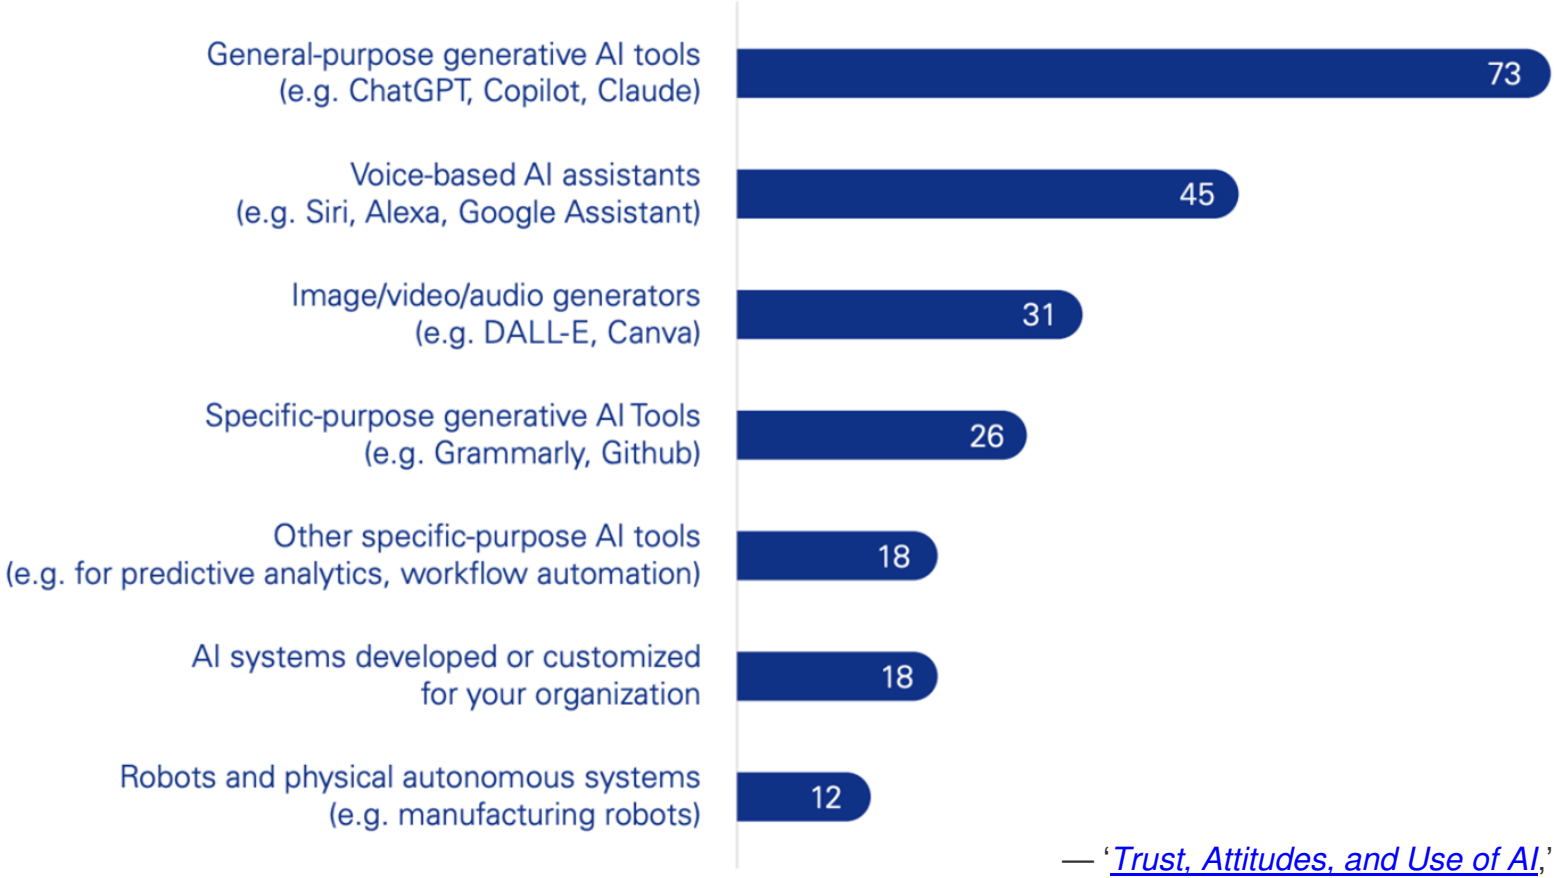

— '*Trust, Attitudes, and Use of AI*,' KPMG, May 2025

So are employees using AI…

So are employees using AI…

- **Embedded within a specific workflow process?**
  - **Salesforce**
  - **Oracle**
  - **Workday**

So are employees using AI…

- **Embedded within a specific workflow process?**
  - **Salesforce**
  - **Oracle**
  - **Workday**
- **Or more like desktop software?**
  - **MS Word**
  - **Excel**
  - **PowerPoint**

Part II:

From AI risks to AI risk committee

The AI risk committee

- What does it do?
- Who should serve?
- Who should lead?
- How can things go wrong?

Start with what it does…

The Big Vision Stuff

- Ethical & responsible use of AI
- Alignment with company values
- Accords with laws & regulations
- Supports business objectives
- Embraces innovation

The Small Tactical Stuff

- Assess risks of AI
- Assure necessary policies, training, controls all in place
- Test AI performance
- Enforce accountability

Have a charter

- Scope of duties
- Responsibilities
- Membership of committee
- Authority
- Meeting minutes, etc.

Think about the big questions…

- How can AI help the company achieve its business objectives?
- How might employees use AI to achieve their daily tasks?
- What big risks arise from Questions 1 and 2?
- What controls can address Question 3?

The biggest challenge…

… everyone not in the room!

Part III:

Frameworks, risk matrices, and more

Risk. Category. Control. Repeat

Risk Category	Description	Likelihood	Impact	Example Controls
Bias & Discrimination	AI leads to unfair treatment in hiring, lending, policing, etc.	High	High	Bias audits, diverse training data, explainability tools
Privacy Violations	Use of personal data without proper consent or safeguards	High	High	Data minimization, purpose limitation, consent management
Lack of Explainability	Inability to justify AI decisions, especially in high-risk contexts	Med.-High	High	Model interpretability tools, human oversight

Challenge: Tracking all risks, remediation

- Inventory of all current, planned, or recently used AI systems
- Description of how each AI operates, including data sources and restrictions on data use
- Description of how you track changes to AI use over time
- Description of how you monitor AI use and performance
- Records of testing done and any ‘model drift’ identified in testing

— *New York Dept. of Financial Services guidance*

Frameworks do abound

- NIST AI Risk Management Framework
- ISO 42001: AI Management Systems
- OECD AI Principles
- IEEE Ethically Aligned Design
- EU Ethics Guidelines for Trustworthy AI
- Microsoft Frontier Governance Framework
- Google Secure AI Framework

MIT Risk Repository

- 1,600+ risks
- 7 ‘domains’
- 23 sub-domains

Great for…

- Risk assessments
- New risks

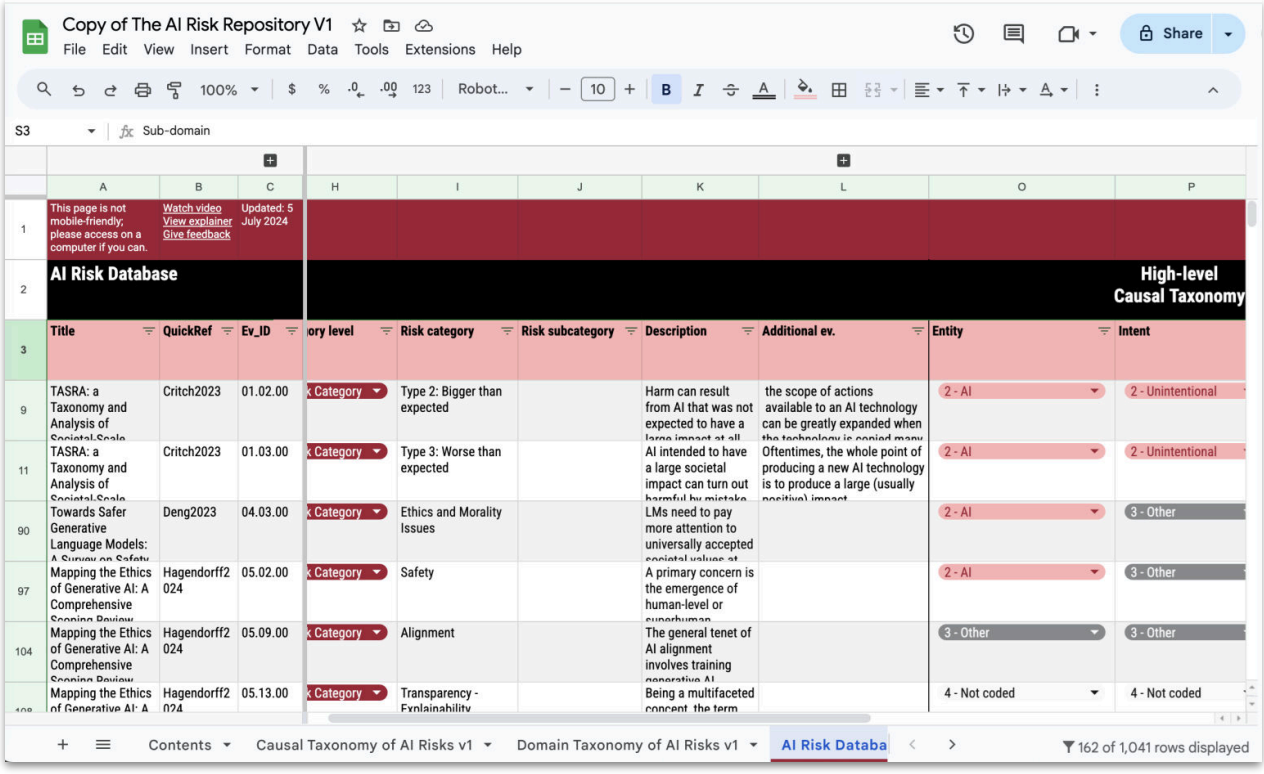

RADICAL COMPLIANCE
SHARP THINKING ABOUT COMPLIANCE, AUDIT, AND RISK

Part VI:

What to monitor and report?

Report the risks and remediation…

- Incidents of bias or discrimination
- Gaps in explainability of AI decisions
- Violations of privacy or AI regulations (GDPR, HIPAA, EU AI Act)
- System failures (AI gone off-line) or model drift
- Third-party risks
- Gaps in governance (inability to test AI inputs or outputs, or to track employee AI usage)

…but talk about the ethical dilemmas

- Where ‘the human point’ should be in an AI-driven process
- Disruptions to stakeholders
- Bottlenecks to innovation that slow down AI usage: good or bad?
- The next leaps in AI that might complicate life even more

…but talk about the ethical dilemmas

- Where ‘the human point’ should be in an AI-driven process
- Disruptions to stakeholders
- Bottlenecks to innovation that slow down AI usage: good or bad?
- The next leaps in AI that might complicate life even more

Tie all of it back to company’s ethical values and business objectives: ‘Who do we want to be as an organization?’

Thank you

Matt Kelly, editor & CEO
www.RadicalCompliance.com
mkelly@RadicalCompliance.com