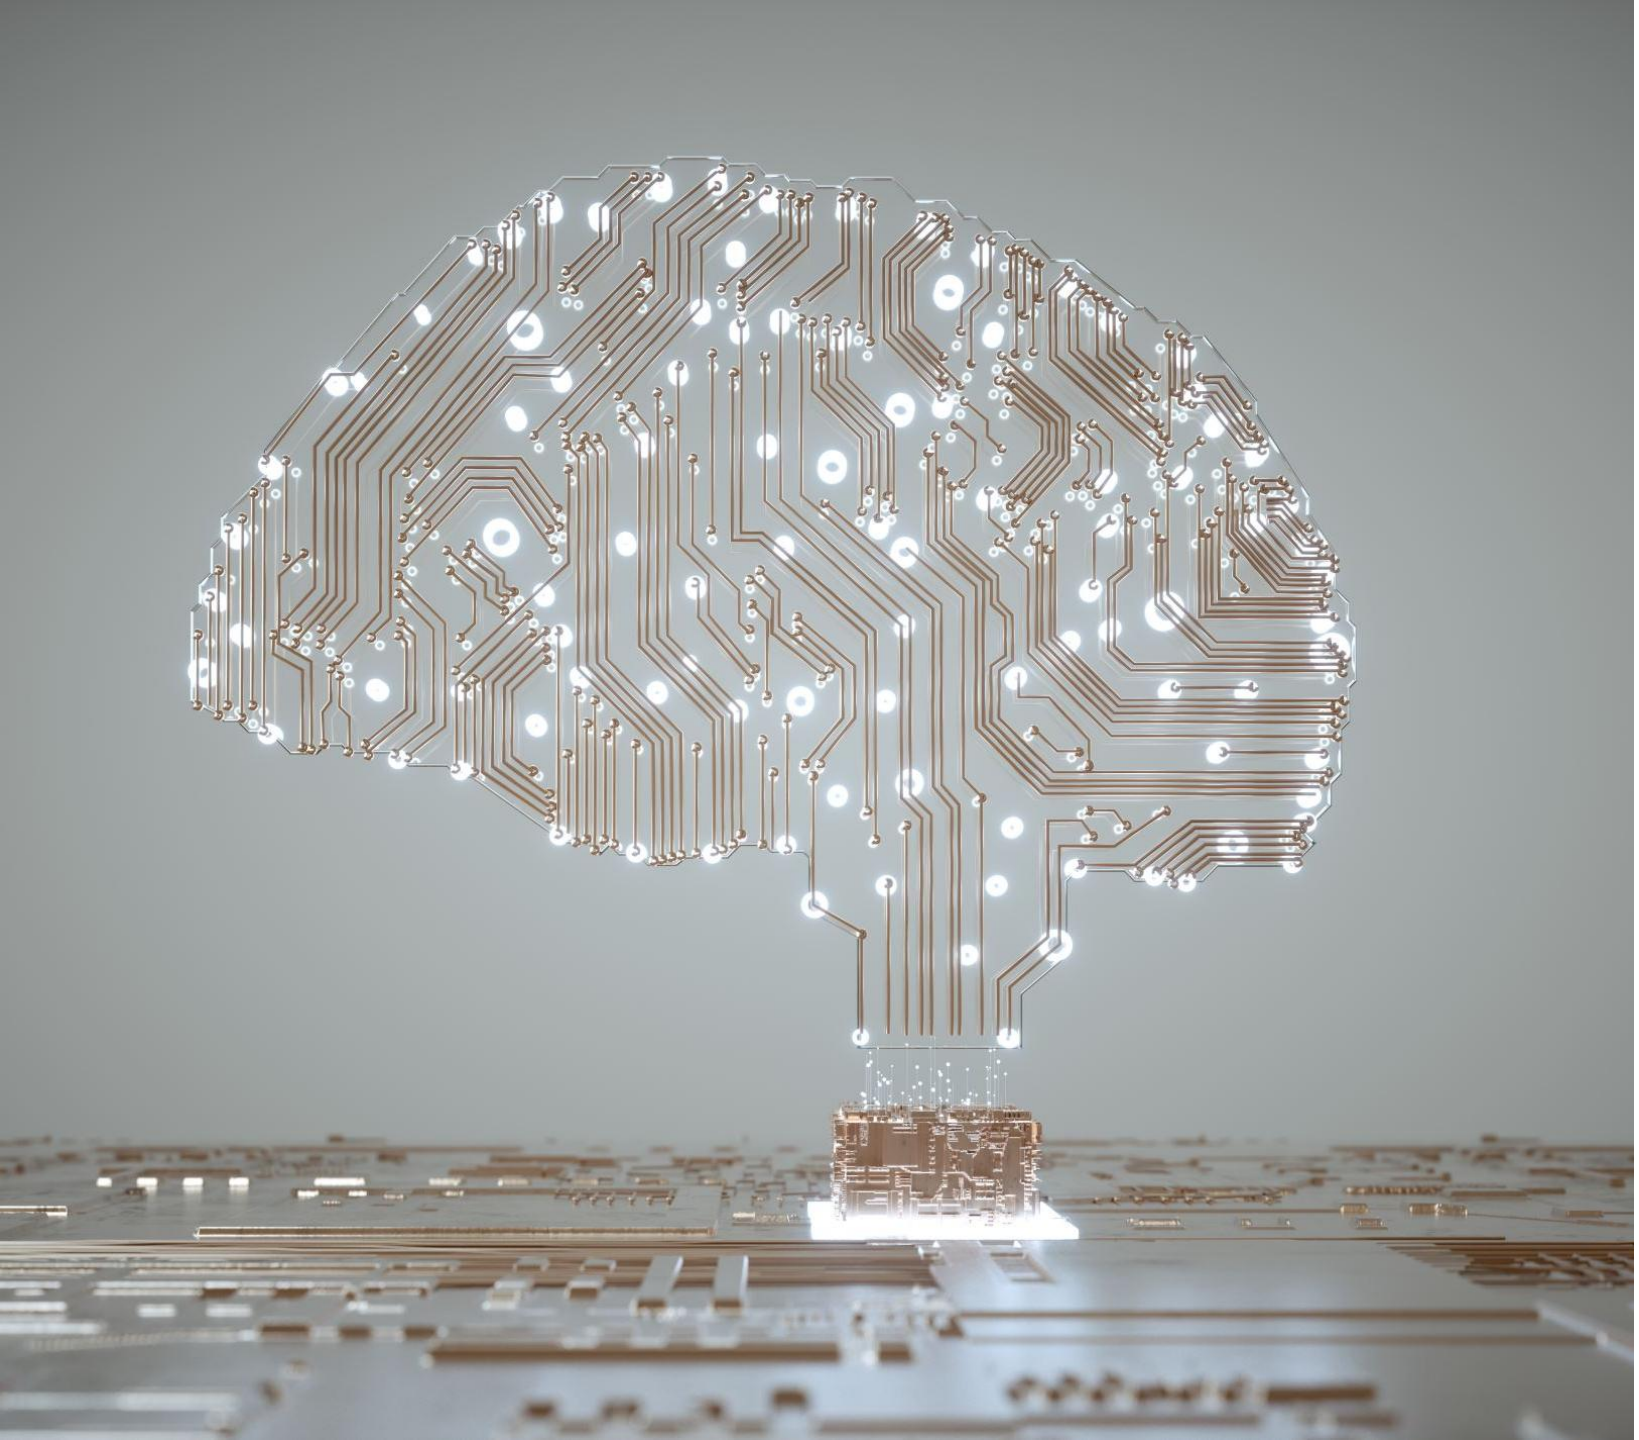

ARTIFICIAL INTELLIGENCE

Practical Guidance to Mitigate Risks

Presenters

Laurel Sykes, CRCM

EVP, Chief Risk Officer

lsykes@arb.bank

Angela Rudd

SVP, Chief Risk Officer

angelarudd@tcbk.com

Jason Levingston

SVP, Chief Information Officer

jasonlevingston@tcbk.com

Paul Abramson

EVP, Chief Technology Officer

pabramson@arb.bank

Agenda

- Definitions
- Use Cases, General Risks, Controls
- Controls for Specific Use Cases
- Governance
- Final Tips & Takeaways

DEFINITIONS

Generative AI

Specializes in crafting new content, including text, images, and videos, by recognizing and replicating learned patterns using advanced models trained on massive datasets.

Large Language Models

LLMs, such as ChatGPT and Gemini, are AI systems capable of understanding and generating human-like language for various applications based on extensive text.

Machine Learning

Empowers AI systems to learn from vast datasets, refining outputs through experience and without the need for explicit programming.

Robotic Process Automation (RPA)

Automates repetitive, rule-based tasks involved in content production – such as data entry, formatting, and distribution – freeing up human capacity for more strategic and creative work.

Speech Recognition and Synthesis

Facilitates the creation and analysis of audio content, making it possible to generate podcasts, voiceovers, and interactive audio experiences.

General Use Cases

Assess performance: “train” models to improve future decisions/content based on real-world outcomes and interactions.

Personalization: by analyzing user data – such as purchase history, browsing behavior, and engagement patterns – AI can customize content for individual recipients or segments.

Analyze massive datasets: Analyze available industry content, trends, seasonal relevance, and key risks or opportunities. Learn from interactions or prior decisions – such as likes, shares, and conversions – to determine what actions to suggest in the future.

Content creation: while AI can autonomously create draft content, humans should review the content for:

- Factual accuracy and compliance
- Sensitivity to cultural or contextual nuances
- Opportunities to add creativity or strategic insight

Strategic Risk

“AI will not replace people, but people who understand how to use AI will replace those that don’t.”

General Risks

- Overreliance creates risk of unidentified errors, miscalculations, or operational disruptions
- Ethical concerns related to copywrite infringement and plagiarism, model bias, job displacement
- Confidential data accessible to the models can increase infosec risks

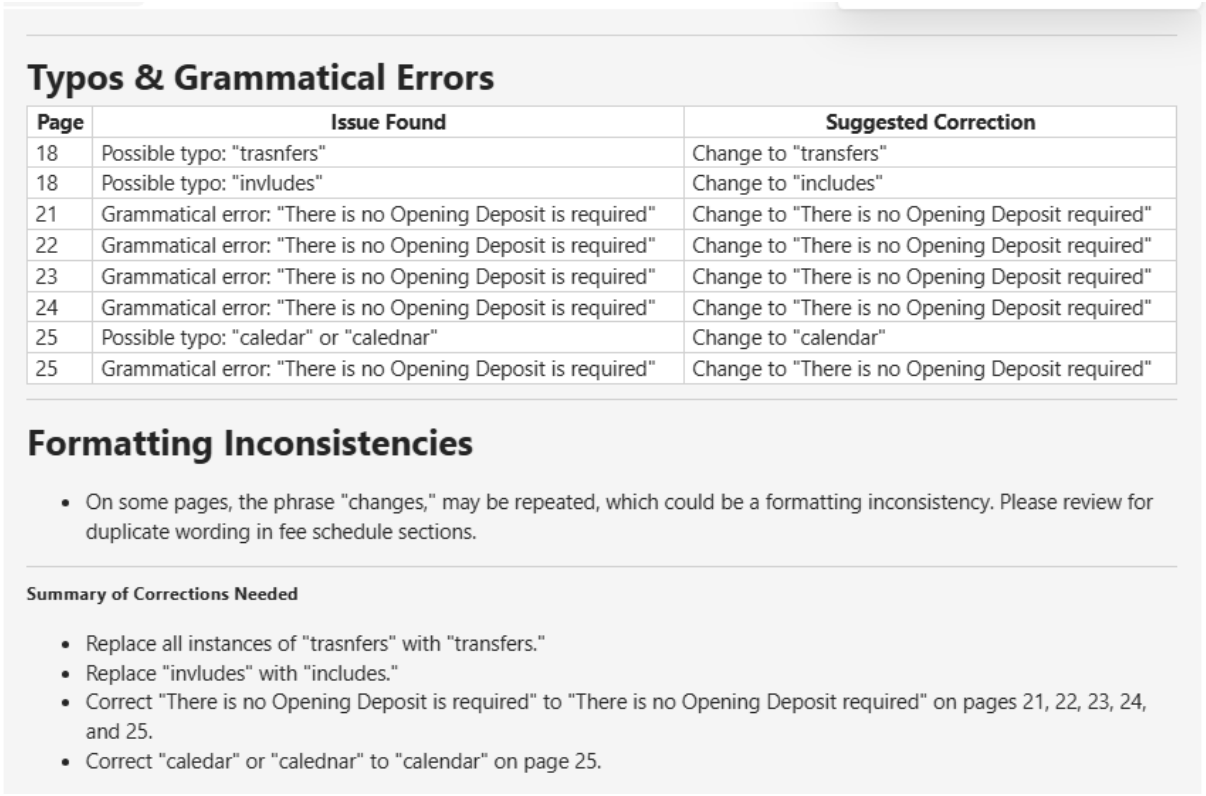

Typos & Grammatical Errors

Page	Issue Found	Suggested Correction
18	Possible typo: "trasnfers"	Change to "transfers"
18	Possible typo: "invludes"	Change to "includes"
21	Grammatical error: "There is no Opening Deposit is required"	Change to "There is no Opening Deposit required"
22	Grammatical error: "There is no Opening Deposit is required"	Change to "There is no Opening Deposit required"
23	Grammatical error: "There is no Opening Deposit is required"	Change to "There is no Opening Deposit required"
24	Grammatical error: "There is no Opening Deposit is required"	Change to "There is no Opening Deposit required"
25	Possible typo: "caledar" or "calednar"	Change to "calendar"
25	Grammatical error: "There is no Opening Deposit is required"	Change to "There is no Opening Deposit required"

Formatting Inconsistencies

- On some pages, the phrase "changes," may be repeated, which could be a formatting inconsistency. Please review for duplicate wording in fee schedule sections.

Summary of Corrections Needed

- Replace all instances of "trasnfers" with "transfers."
- Replace "invludes" with "includes."
- Correct "There is no Opening Deposit is required" to "There is no Opening Deposit required" on pages 21, 22, 23, 24, and 25.
- Correct "caledar" or "calednar" to "calendar" on page 25.

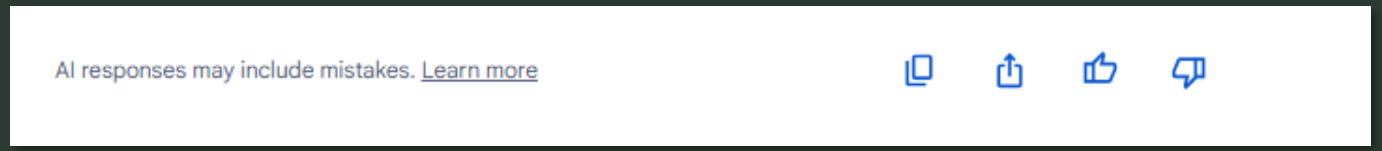

Mitigating AI Risks

Model Risk Management, Change Management, Vendor Management

Schedule regular validation activities, such as back-testing. Continuously retrain for correct results using false positives/negatives. Integrate with existing change management and vendor management programs.

Maintain a Balance Between AI-driven Automation and Human Oversight

Transparent decision-making and clear documentation enhance trust in AI systems and enable accountability for outcomes. Institute quality control and callback processes.

Data Privacy and Access Management

Strong data privacy measures and access management prevent unauthorized use and protect sensitive information in AI systems. Secure the systems and platforms leveraging the AI models.

Ongoing Employee Training and Acceptable Use Policies

Continuous training ensures employees understand ethical AI practices, current risks, and compliance requirements. Limit use without dual controls to only those who are fully trained.

Human in the Middle

- Consider the risks of hallucinations or inaccurate interpretations
- Review and re-prompt, citing specific sources
- Quality control by qualified individuals
- AI an assistant, not a replacement, for an employee

“The executive order directs federal agencies to review and revise examination criteria for community banks participating in Small Business Administration (SBA) programs, with particular attention to the lending practices affecting small businesses across diverse industries. Its intent is to ensure that account closures are not motivated by "discriminatory or overly broad risk assessments" within sectors such as retail, hospitality, manufacturing, and technology – industries historically vulnerable to abrupt debanking actions. The order further mandates "transparency and consistency in how banks assess risk and make decisions about account closures or denials for SBA borrowers," emphasizing that such actions must be based on "objective, non-discriminatory standards." By singling out these key industries, the executive order aims to protect legitimate small business borrowers from unfair exclusion and to promote fair lending practices in the broader marketplace.”

Specific Use Cases

Risks and Controls

Compliance/Legal Use Case

Use Case	How it Works	Risks	Controls
Compliance/Legal	• AI systems equipped with Natural Language Processing (NLP) can parse complex legal language, identify key clauses, and cite reference documents. • AI generates a draft policy for review and approval by management. • Scan vast databases of legal documents, statutes, and case law, identifying relevant information and precedents • Interpret laws and regulations to ensure compliance. By summarizing legislation and identifying obligations, they assist businesses in navigating complex legal frameworks without overlooking critical details. • Contract review tools can detect inconsistencies, flag unusual terms, and suggest revisions.	• Dependence on AI tools can lead to diminished critical thinking. If the AI system fails or produces errors, the legal consequences could be far-reaching. • Risk of employees not learning the material and/or overlooking key issues. • Hallucination, where the tool provides information outside of your prompt • Database can include incomplete or obsolete references. • Risk of employees revealing sensitive data such as attorney client protected information through the AI prompt. • Incorrect summaries missing key obligations or details that exist in other governing documents. • Lack of transparency in how conclusions are drawn can undermine the credibility of AI-assisted legal interpretations	• Human in the middle • Limit use to individuals who are formally trained in the regulations they are working with • Employee training to craft effective and appropriate prompts • Use it to enhance interpretation and not replace it • Understand the risks associated with revealing sensitive information through the prompt • Include legal citations to call the information back to the source of the information.

Marketing Use Cases

- Copilot
- Grammarly Business
- Writesonic

How it Works	Risks	Controls
Articles and Blog Posts: • AI creates outlines, headlines, introductions, and other content based on provided topics and tone from historical posts. • It can expand bullet points into paragraphs, summarize research, and even suggest images or infographics based on what is available online. • The AI checks for logical flow, readability, and keyword integration to support SEO. Advertisements: • AI crafts concise, attention-grabbing headlines and calls-to-action, using A/B testing data and audience segmentation insights. • It adapts ad copy for different platforms (Google Ads, Facebook, LinkedIn) by adjusting tone, length, and format. Social Media Posts: • AI generates multiple versions of posts tailored to each network's character limits and audience preferences. • It suggests optimal posting times, hashtags, and visual elements to maximize engagement. Emails: • AI writes subject lines, body content, and personalized greetings using recipient data and behavior history. • It can segment email lists and tailor messages for different customer journeys (welcome, re-engagement, upsell, etc.). • Emails are tested for readability, spam triggers, and conversion likelihood. • A/B testing to "learn" from the effectiveness of marketing campaigns to ensure a better response on future campaigns.	• Plagiarizing • Hallucination, where the tool provides information outside of your prompt • Content that is not authentic	• Acceptable use policies • Redraft for tone • Human in the middle • A/B campaigns to measure effectiveness against previous campaigns that were human generated

Credit Scoring and Underwriting

- Using Copilot to summarize lengthy documents and/or create executive summaries
- Use of FICO could be replaced by VantageScore or alternative scoring methods
- Vendors are developing “assistants” to perform spreads and identify trends using multiple years of financial data

How it Works	Risks	Controls

• Credit scoring models use your credit report information—specifically payment history, amounts owed, length of credit history, types of credit, and credit inquiries—to calculate a three-digit score that predicts likelihood of repaying debts.
• In underwriting, AI can be used to summarize lengthy documents by using large language models to analyze content, identify key themes and important information in a user-defined format, such as bullet points or summaries. AI leverages publicly available information as well as information we provide, and/or information stored in the user’s environment to assist with tasks such as a summary of loan information that a junior credit analyst would otherwise review.
• Alternative credit scoring methods may also be used to analyze non-traditional data sources such as social media activity, online behavior, and local economic trends.
  ○ Alternative credit scoring methods assess creditworthiness using data beyond traditional credit reports, including utility and rent payments, bank account transaction history, Buy Now, Pay Later (BNPL) and payment app data, and employment and income sources.
  ○ Some methods also use non-financial data like digital footprints, device behavior, or social media activity to build a more inclusive and accurate financial profile for individuals lacking traditional credit history.
	• Model risk, including lack of transparency in how decisions are made
	• Bias in training data can lead to unfair credit decisions or discriminatory practices
	• Risk of employees not learning the material and/or overlooking key issues.
	• AI-driven decisions in areas such as loan approvals and credit scoring can raise ethical issues, especially if decisions lack transparency or accountability
		• Independent validation of model
		• Validation activities, such as backtesting of model inputs and outputs, including false positives/negatives
		• Human in the middle without full reliance on AI
		• Limit use to individuals who are formally trained in the regulations they are working with i.e. only a senior analyst may review AI generated content

Copilot Use Case

Data stored and generated in Microsoft 365

- Efficiency in daily tasks
- Meeting minutes
- Summaries of email threads
- Drafts of documents
- Image and graphic creation

Use Case	How it Works	Risks	Controls
Efficiency in daily tasks such as responding to emails, reading committee materials, taking minutes	Copilot builds on top of the existing Microsoft 365 environment and when enabled can access all data stored and generated on the platforms including Teams, OneDrive, SharePoint, Exchange Online, and Office. The use cases for Copilot are vast and include some of the previously mentioned activities. Capability to generate minutes in Teams meetings is also something that the Bank has approved.	• Improper data segmentation • Access control violations	• Microsoft Enterprise Data Protection to force data segmentation and prohibit the use of the public Copilot interface within the Bank environment • Acceptable use policy • Authorized access for full Copilot licensing • All other access control policies in place to protect Microsoft 365 provide protection for Copilot interfaces (e.g. Conditional Access, and Tenant Restriction)

Copilot Configuration Checklist

Baseline: Leverage or add Microsoft 365 controls including:

- ☐ Conditional Access Policies (e.g. IP restriction, MFA, session risk, endpoint status, geolocation, etc.)
- ☐ Tenant Restrictions (Prevent access to 3rd party Microsoft tenants)
- ☐ Allow only environments with a business need to collaborate
- ☐ Consider blocking personal Microsoft accounts
- ☐ Other security policies related to granular file sharing and access control by application
- ☐ Microsoft Digital Rights Management tools
- ☐ InTune policies

Intermediate: Additional Copilot controls:

- ☐ Microsoft Commercial Data Protection/Enterprise Data Protection (Requires sign-in to Bing chat and public Copilot site. Ensures privacy between tenants and enforces access control)
- ☐ Verify Enterprise Data Protection is enabled by default
- ☐ Configure DNS forwarder to require sign-in and block interaction with public resources

Advanced: Additional General AI controls:

- ☐ Use enterprise browsers for more granular control over website functionality
- ☐ Avoid overly restrictive or “whack a mole” filtering policies

AI Implementation Steps

Define and Approve Use Cases

Begin by evaluating business requirements and identifying key use cases where AI can add value.

Prepare Infrastructure and Train Users

Provide your users with vetted and controlled AI solutions or they will adopt “shadow IT” solutions and introduce security privacy risks.

Pilot and Gather Feedback

Launch a pilot with selected employees, monitor usage, and collect feedback to measure effectiveness.

Analyze and Scale Organization-wide

Review pilot results, refine processes, deploy more globally using an authorized user approval process. Incorporate into vendor management process.

Implement Policies and Acceptable Use Standards

Integrate acceptable use standards into Employee Handbook and create a guiding principles for those with access.

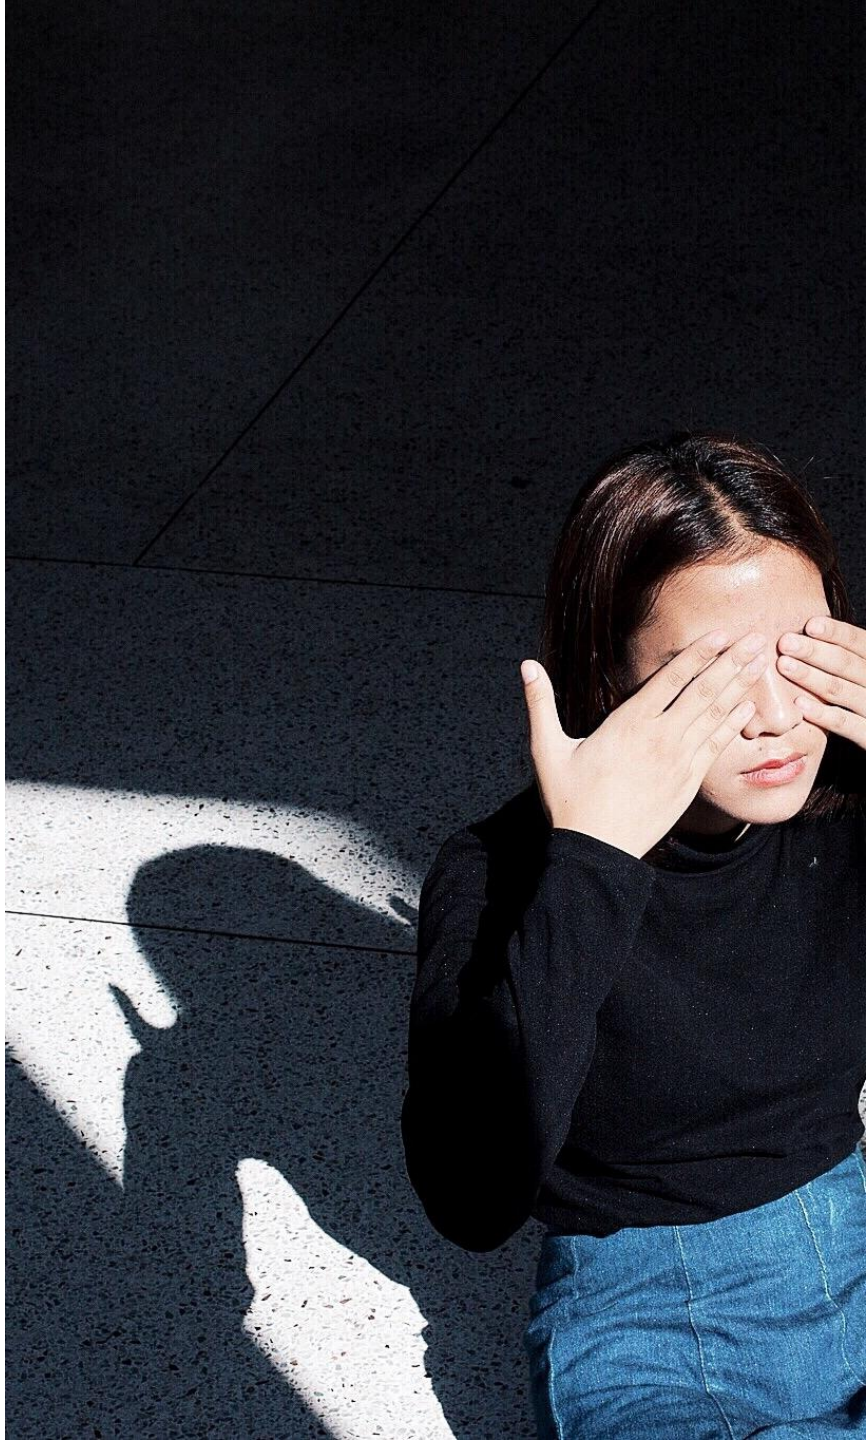

AI Governance

- Guiding Principles
- Strategic Alignment
- Innovate / Manage Risk

Responsible AI Principles

Data Privacy

AI applications must respect user privacy. Data must not be used outside of agreed upon terms and must be compliant with privacy norms and regulations.

Accountability

AI policies must outline the individuals or groups accountable for the planning and deployment of any AI system and must document how it will be governed.

Explainability & Transparency

AI applications will be transparent about how data is used and will provide users and key stakeholders insights into how outcomes are produced.

Fairness & Bias Detection

AI applications must include checks and balances to ensure results are unbiased and there is fair and equitable representation across users.

Security & Safety

AI applications must be resilient to attacks and other risks that could provide physical or digital detriment to individuals or groups.

Validity & Reliability

AI applications must produce results that are accurate and consistent to mitigate AI risk and foster trust in the application.

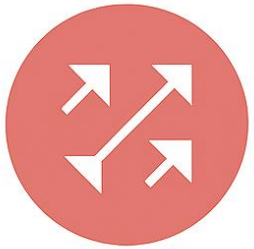

CHANGE MANAGEMENT

VENDOR MANAGEMENT

CONTRACT MANAGEMENT

MODEL RISK MANAGEMENT

Integrate with Existing Governance Frameworks

Change Management

AI Steering Committee

- Strategic oversight (e.g., Executive Sponsor)
- Representation from Risk functions (e.g., IT/IS, Compliance, Risk, Legal, Operations, TPRM, etc.)

Center of Excellence

- Innovation hub
- Best practices and training

Working / Advisory Groups

- Pilots
- Risk and ethics reviews

Board and Management Reporting

- Oversight & accountability

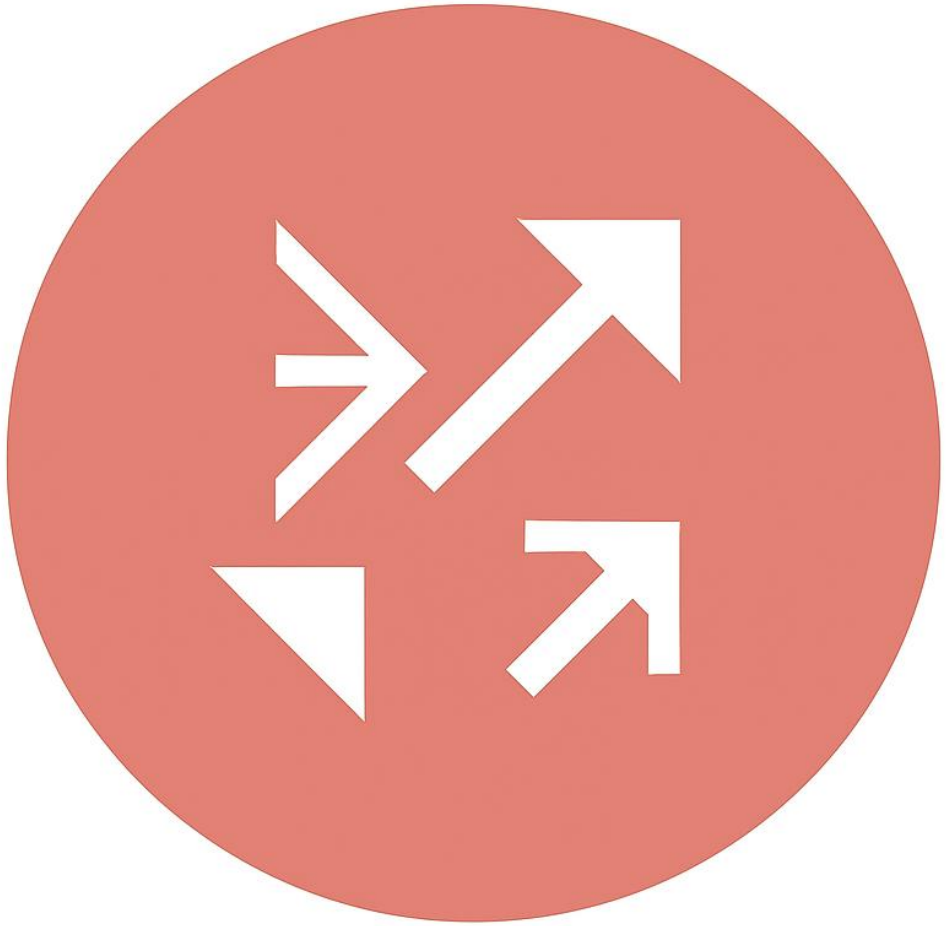

Vendor Management

Due Diligence

- Governance maturity, transparency, regulatory alignment.
- Model training, bias mitigation, auditability

Ongoing Monitoring

- Establish continuous oversight mechanisms

Collaboration between Risk, Legal, Compliance and IT/IS

Use “Pilot” to test before scaling

- Establish acceptable guardrails and be prepared to fail fast

Vendor oversight is key!

Contract Management

Contractual Safeguards Are Non-Negotiable

Embed AI-specific clauses in vendor contracts:

- Intellectual property ownership
- Liability for model errors
- Data usage rights
- Right to audit
- Responsibility to notify as use of AI changes
- Termination triggers for non-compliance

Model Risk Management

Risk Management Practices must evolve

- Standard practices include validation, governance, regular review, and documentation.
- “Black box” - *Explainability and fairness*

Governance must span the lifecycle

- Include ethical and regulatory checkpoints
- Continuous validation

Model inventory and classification

- Centralized
- Complexity , impact, and regulatory exposure

Cross-functional collaboration is critical

- Shared accountability and clear escalation paths

Next Steps/ Takeaways

1. Try it out!
2. Determine initial use cases to align with business strategy
3. Conduct a risk assessment
4. Institute cross-functional team, such as an AI Governance Committee
5. Implement policies, procedures, and guiding principles
6. Incorporate acceptable use standards into employee handbooks
7. Inventory vendors to assess risks
8. Integrate into change management processes
9. Document tools/model inventory that includes AI

Compliance Officer Prompts

1. I think [insert idea]. Pretend to completely disagree. Give me the strongest arguments against it so I can see the blind spots.
2. Explain [insert the legal issue] to me as if I'm 5 years old. Then explain it again for a smart banking executive with no time to read.
3. Here's my draft [paste text] to explain the risks associated with this compliance issue [insert citation]. Rewrite it so that anyone with zero content understands it instantly.
4. I'm stuck choosing between three alternatives [option a], [option b] and [option c] as codified here [insert citation]. Lay out the risks and benefits side-by-side so the tradeoffs are obvious.