

General Risks Associated with Using AI

1. Data Privacy and Security

Collection and storage of sensitive customer data used by AI models can increase the risk of experiencing cyberattacks and data breaches.

2. Bias in AI Algorithms

Bias in training data can lead to unfair credit decisions or discriminatory practices, potentially harming the bank's reputation and leading to regulatory criticism.

3. Over-reliance on AI

Excessive dependence on AI models may result in operational disruptions if these systems fail.

4. Ethical Concerns

AI-driven decisions in areas such as loan approvals and credit scoring can raise ethical issues, especially if decisions lack transparency or accountability and can enable willing or unwilling acts of plagiarism.

5. Regulatory Challenges

The rapid advancement of AI has outpaced regulatory frameworks. Making it more difficult to anticipate and conform with potential changes in regulations.

6. Model Risks

AI models are susceptible to errors, miscalculations, or malicious manipulation, which can lead to financial losses or incorrect decisions. This is an emerging technology and many of the risks are currently unknown/undiscovered.

General Controls

Banks should implement the following controls:

POTENTIAL USE CASES AND RISKS

- Establish controls to ensure compliance with data privacy laws.
- Regularly audit AI models to identify and mitigate biases. Do not inherently trust the accuracy of the output.
- Maintain a balance between AI-driven automation and human oversight.
- Invest in employee training to adapt to AI-enhanced workflows.
- Collaborate with regulators to anticipate and address evolving compliance requirements.
- Secure the systems and platforms leveraging the AI models.
- Identify which AI models providers use and maintain a list of approved models with applicable controls for each depending on use case.

Definitions

AI-powered content creation draws upon a suite of advanced technologies:

- Natural Language Processing (NLP): Equips machines to comprehend, interpret, and generate human language, enabling more nuanced and effective communication.
- Machine Learning (ML): Empowers AI systems to learn from vast datasets, refining their outputs through experience and without the need for explicit programming.
- Large Language Models (LLMs): Sophisticated AI models—such as ChatGPT and Gemini—trained on extensive text to produce coherent, context-aware content.
- Generative AI: Specializes in crafting new content, including text, images, and videos, by recognizing and replicating learned patterns.
- 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.
- Computer Vision: Allows AI to interpret and process visual information, supporting tasks like image recognition, tagging, and custom graphic generation within content workflows.
- Speech Recognition and Synthesis: Facilitates the creation and analysis of audio content, making it possible to generate podcasts, voiceovers, and interactive audio experiences.

Together, these technologies form the backbone of contemporary AI-driven content creation, each contributing to more efficient, creative, and tailored business process strategies.

POTENTIAL USE CASES AND RISKS

General Uses

- Assess performance: AI systems continuously learn 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 customizes content for individual recipients or segments. This increases relevance and the likelihood of conversion. Automation tools can dynamically adjust text, images, product recommendations, and prices in real-time to match audience preferences.
- Analyzing massive datasets: Analyze available industry content. Assess 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

Use Case	How it Works	Risks	Controls
Process Automation	Robotic Process Automation (RPA) can automate repetitive tasks such as document verification, client reviews, and account onboarding	• Risk of employees not learning the material and/or overlooking key issues.	Callback / Review / Dual approval
AML	• AI algorithms analyze transactional data in real-time to identify anomalies and detect potential ML activity. • Machine learning models help banks detect less obvious ML activity patterns that human analysts could miss.	• Data Privacy and Security especially related to the customer information required to operate the model • Model Risks • Chatbots/LLMs not designed for forensic analysis • Overreliance on efficiency gains	• Ensuring robust security controls such as data segmentation, vendor review, access control • Compliance with data protection regulations • Validation activities, such as backtesting of model inputs and outputs, including

POTENTIAL USE CASES AND RISKS

Use Case	How it Works	Risks	Controls
			false positives/negatives • Adequate contingency plans in the event it is unavailable
Fraud Detection and Prevention	• AI algorithms analyze transactional data in real-time to identify anomalies and detect potential fraud. • Machine learning models help banks prevent unauthorized transactions by learning what is out-of-pattern for a client.	• Model risk, including lack of transparency in how decisions are made • Bias in training data can lead to unfair decisions or discriminatory practices	• Independent validation of model • Validation activities, such as backtesting of model inputs and outputs, including false positives/negatives
Credit Scoring and Loan Underwriting	• 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.	• 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

POTENTIAL USE CASES AND RISKS

Use Case	How it Works	Risks	Controls
	○ 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.		
Marketing	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.	• 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

POTENTIAL USE CASES AND RISKS

Use Case	How it Works	Risks	Controls
	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.		
Drafting documents	Tools such as co-pilot have the capability to search through other content on the Bank’s file storage that share the same content as the prompt.	• Plagiarizing • Hallucination, where the tool provides information outside of your prompt • Improper data segmentation, privacy, access control violations	• Employee training to craft effective and appropriate prompts • Human in the middle • Acceptable use policies • Access control • Validation of data segmentation • Approved vendor list
Risk Management	AI models can assess market conditions, generate summaries of risk reports, and assist with recommending controls.	• Without proper human review, errors, inconsistencies, or fabricated data (hallucinations) can go unnoticed	• Employee training to craft effective and appropriate prompts • Human in the middle

POTENTIAL USE CASES AND RISKS

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.	• 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.

POTENTIAL USE CASES AND RISKS

Use Case	How it Works	Risks	Controls
		Lack of transparency in how conclusions are drawn can undermine the credibility of AI-assisted legal interpretations	
Customer Service or Internal Support	AI-driven chatbots and virtual assistants provide 24/7 support, or assist with common customer queries by ingesting FAQs, policies documents, disclosures, etc.	- Untrained models can provide incorrect answers. - Risk of employees not learning the material and/or overlooking key issues. - Public access to AI interfaces with access to privileged data or providing advice/service that carries liability opens up the possibility for abuse. This is an emerging risk and attack vectors are novel and being actively developed and used. - Canned responses may not feel authentic from a customer service standpoint.	- Human in the middle to customize canned responses - Limited scope, data access, and capability - Penetration testing - Mystery shopping
Predictive Analytics for Customer Retention	AI analyzes customer behavior to predict potential churn. Banks use these insights to implement targeted campaigns and offer incentives to retain valuable clients.	- Model risk, including lack of transparency in how decisions are made - Potential noncompliance with customer privacy and cybersecurity laws	- Validation of results and other output - Measurement of effectiveness

POTENTIAL USE CASES AND RISKS

Use Case	How it Works	Risks	Controls
		• Bias in training data can lead to unfair decisions or discriminatory practices • When fed incomplete, outdated, or biased data, AI can produce flawed outputs	
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 of the 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 • Inaccurate results	• 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)
Financial analysis	AI can enable accelerated financial reconciliation by matching transactions, flagging discrepancies, and preparing reports; perform variance analysis that highlights anomalies, surfaced potential drivers, and generated draft explanations;	• Over-reliance on AI can create a "black box" scenario where finance teams lose critical analytical	• Human in the middle without full reliance on AI

POTENTIAL USE CASES AND RISKS

Use Case	How it Works	Risks	Controls
	and use data preparation to clean, reshape, and structure financial data exports into analysis-ready reports.	- skills and become complacent. - Without proper human review, errors, inconsistencies, or fabricated data (hallucinations) can go unnoticed, compromising the integrity of financial statements. - Inaccurate financial information from AI can lead management to make misguided choices, resulting in financial losses and missed market opportunities. - Risk of employees not learning the material and/or overlooking key issues. - Incorrect outputs.	- 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 - Callback / Review / Dual approval

USE IN FINANCIAL REPORTING

Inaccuracies in AI-generated financials pose significant risks across multiple business functions, including legal, regulatory, and reputational areas. These errors can arise from flawed training data, inherent model limitations, or intentional misuse, leading to flawed decision-making and substantial financial and legal consequences.

Strategic and operational risks

- Flawed business strategies:** Financial reports and forecasts drive crucial business decisions like resource allocation, budgeting, and expansion strategies. Inaccurate financial information from AI can lead management to make misguided choices, resulting in financial losses and missed market opportunities. For example, overestimating revenue could lead to overspending and cash flow issues.

POTENTIAL USE CASES AND RISKS

- **Poor data quality:** AI models are only as reliable as the data they are trained on. When fed incomplete, outdated, or biased data, AI can produce flawed outputs and misrepresent a company's financial health. Relying on outdated historical data without context, for instance, can cause AI to miss new economic trends.
- **Lack of human oversight:** Over-reliance on AI can create a "black box" scenario where finance teams lose critical analytical skills and become complacent. Without proper human review, errors, inconsistencies, or fabricated data (hallucinations) can go unnoticed, compromising the integrity of financial statements.
- **Inconsistent outputs:** Even when given the same data, some AI models, especially generative ones, can produce inconsistent or non-repeatable results due to algorithmic randomness. This makes it difficult to maintain reliability for financial reporting, which requires consistency for credibility and compliance.

Legal and compliance risks

- **Regulatory penalties:** Inaccurate financial reporting can lead to non-compliance with standards and regulations set by governing bodies like the SEC. Authorities are increasing their scrutiny of AI use in finance and have already issued fines for false AI claims, or "AI washing". In extreme cases, companies could face large penalties and sanctions.
- **Litigation and liability:** Companies face potential lawsuits from investors, customers, and other stakeholders who suffer financial losses due to reliance on flawed AI-generated financials. Misleading investment advice from an AI, or failure to disclose the risks of AI products, can lead to claims of fraud or negligence.
- **Fraud and manipulation:** Malicious actors can use generative AI to create believable fake financial documents, social media profiles, or clone voices to conduct investment scams and other financial fraud. Corporate insiders could also use AI to create fraudulent documentation for embezzlement or accounting fraud.
- **Bias and discrimination:** If an AI is trained on historical data reflecting past bias, it can perpetuate or amplify that discrimination in areas like credit scoring or loan approvals. This can lead to unfair or illegal outcomes and increase liability under anti-discrimination laws.

Reputational and market risks

- **Loss of stakeholder trust:** The confidence of investors, creditors, and customers is foundational to a company's success. If it is discovered that a company's financials were based on inaccurate AI outputs, trust can be severely eroded, leading to a damaged brand image, loss of customers, and reduced market value.

POTENTIAL USE CASES AND RISKS

- **Investor distrust:** False or misleading AI-generated financial projections can lead to poor investment decisions by investors. An AI's "hallucinations"—invented information presented as fact—could cause financial institutions to present a stronger financial position than reality, damaging investor faith.
- **Market instability:** Widespread adoption of similar AI financial models could lead to herd behavior during a market crisis. If many firms use AI models with the same inherent biases or flaws, they might make similar decisions simultaneously, amplifying market volatility and potentially threatening financial system stability.
- **Negative publicity:** When AI systems fail, the fallout can quickly become public, with social media amplifying the negative impact. Handling these crises diverts significant resources and can have immediate and severe reputational consequences.