Scientific Research Publishing

Journal of Computer and Communications, 2026, 14(11), 204-239
https://www.scirp.org/journal/jcc
ISSN Online: 2327-5227
ISSN Print: 2327-5219

The Identity-to-Transaction Gap: Rethinking U.S. Bank Customer Identification Rules in the Age of Generative AI and Deepfake Financial Fraud

N. B. M. Abrar Faisal

Supreme Court of Bangladesh, Dhaka, Bangladesh
Email: abrar09faisal@gmail.com

How to cite this paper: Abrar Faisal, N.B.M. (2026) The Identity-to-Transaction Gap: Rethinking U.S. Bank Customer Identification Rules in the Age of Generative AI and Deepfake Financial Fraud. *Journal of Computer and Communications*, **14**, 204-239.
https://doi.org/10.4236/jcc.2026.1411010

Received: September 12, 2026
Accepted: November 17, 2026
Published: November 20, 2026

Copyright © 2026 by author(s) and Scientific Research Publishing Inc.
This work is licensed under the Creative Commons Attribution International License (CC BY 4.0).
http://creativecommons.org/licenses/by/4.0/



Open Access

Abstract

Generative artificial intelligence has changed the cybersecurity conditions under which banks can rely on digital identity. Deepfake video, cloned voice, manipulated documents, credential compromise, and digital injection attacks can allow an unauthorized actor to operate behind the identity of a customer whom a bank has correctly identified. This article examines whether U.S. banking regulation provides a sufficiently coherent standard for that post-onboarding problem. Using doctrinal and policy analysis of customer identification, customer due diligence, suspicious activity reporting, FinCEN materials, and the 2025 NIST Digital Identity Guidelines, it develops the concept of an "Identity-to-Transaction Gap": the separation that can arise between a bank's knowledge of its verified customer and its assurance about the person exercising authority over a later high-risk transaction. The article argues that this is not an absence of regulation but a cybersecurity coordination problem across identity proofing, authentication, fraud monitoring, and transaction decision-making. It proposes Transaction-Level Identity Assurance (TLIA), a technology-neutral and risk-based framework in which materially elevated indicators of impersonation or account takeover trigger proportionate additional assurance before execution. TLIA is designed to complement rather than replace existing customer identification, due diligence, authentication, monitoring, and reporting controls. For cybersecurity research and policy, the contribution is an interdisciplinary model that connects legal duties with modern digital-identity threat models while preserving proportionality, privacy, accessibility, and operational flexibility. The analysis incorporates three publicly documented case and threat materials, distinguishes genuine-user payment deception from account takeover, and identifies measurable outcomes

for future empirical evaluation of the proposed framework. No primary empirical dataset or demonstrated effectiveness is claimed.

Keywords

Generative AI, Deepfake Fraud, Banking Regulation, Transaction-Level Identity Assurance, Account Takeover

1. Introduction

1.1. Background: Generative AI and the Changing Nature of Financial Fraud

Banks have always faced the problem of establishing who their customers are. What has changed is the ease with which a person's identity can now be imitated. Generative artificial intelligence (GenAI) can produce realistic images, video, audio, and documents at relatively low cost and with limited technical skill. In financial services, this creates a problem that goes beyond conventional identity theft. A fraudster may not need to steal a physical identity document or impersonate a customer in person. AI-generated or manipulated material can be used to imitate the evidence on which digital identity systems themselves rely.

This is no longer a hypothetical concern. In November 2024, the Financial Crimes Enforcement Network (FinCEN) reported an increase in suspicious activity reports describing the suspected use of deepfake media in fraud schemes targeting financial institutions and their customers. According to FinCEN, criminals have used GenAI-generated or manipulated identity documents to circumvent identity verification and authentication controls. FinCEN has also identified cases in which fraudulent identities believed to have been created with GenAI were used to open accounts and subsequently receive or launder proceeds of other fraud schemes [1].

The scale of the broader identity problem is significant. In its 2024 analysis of identity-related suspicious activity, FinCEN examined Bank Secrecy Act reports filed in 2021 and identified approximately 1.6 million reports associated with identity-related exploitation, representing about 42 percent of the reports filed during that year and approximately $212 billion in suspicious activity. Importantly, FinCEN did not treat identity abuse merely as an account-opening problem. Its framework considered exploitation across three stages of the customer relationship: account creation, account access, and transaction processing [2] [3].

This three-stage understanding is particularly important in the age of GenAI. The traditional question "Who is this customer?" remains necessary, but it is no longer sufficient in every case. A bank may correctly identify a legitimate customer when an account is opened, yet a different person may later obtain access to that account through stolen credentials, social engineering, synthetic media, voice

cloning, deepfake video, or other forms of AI-assisted impersonation. The identity recorded by the bank may therefore remain accurate while the person actually controlling a particular interaction or transaction is not the verified customer. This distinction creates an important challenge for banking law and financial regulation.

1.2. The Problem of AI-Enabled Identity Exploitation

The regulatory architecture governing customer identification was largely developed around the assumption that identity could be established through reliable documentary or non-documentary evidence and that appropriate authentication could subsequently protect access to the account. GenAI complicates this assumption because it can attack the evidentiary layer itself.

A fraudulent document may appear genuine. A facial image may belong to a real customer but be artificially generated or manipulated. A voice used to authorize an instruction may sound like the account holder without being produced by that person. More sophisticated attacks may inject forged digital media directly into remote verification systems rather than merely presenting a fake image to a human reviewer.

The technical response to this problem is already developing. NIST Special Publication 800-63-4, finalized in 2025, revised the federal Digital Identity Guidelines to address changes in the digital identity environment. Among other developments, the revision expanded fraud-related requirements and introduced controls addressing injection attacks and forged media, including deepfakes. NIST SP 800-63A-4 specifically recognizes that GenAI can be used to create or modify images and videos to defeat document validation, biometric systems, and visual comparisons [4]-[6].

These developments reveal an emerging tension. Technical standards are increasingly designed around the possibility that digital evidence, biometric information, and apparently live interactions may be manipulated. Banking regulation, however, still contains important obligations that were formulated before these technologies became widely accessible.

The problem is therefore not simply whether banks possess technology capable of detecting deepfakes. The more difficult legal question is what level of identity assurance the law should require from a financial institution after a customer has already been identified and an account has been established.

1.3. Research Problem: The Identity-to-Transaction Gap

This article describes this problem as the "Identity-to-Transaction Gap". The term is used here as an analytical concept to distinguish between two related but different propositions: first, that a financial institution has established the identity of its customer; and second, that the institution has sufficient assurance that the person initiating or controlling a later transaction is in fact that customer or another authorized person.

These propositions can increasingly diverge. For example, a bank may have properly identified Customer A when the account was opened. The records held by the bank may be entirely accurate. If Fraudster B later gains control of the account and uses AI-generated voice, video, documents, or compromised authentication information to initiate a high-risk transaction, the original customer identification has not necessarily failed. The bank still knows who its customer is. The immediate problem is different: does the bank know who is acting through the customer's account at the point of transaction?

The distinction becomes clearer when customer identity is examined across the banking relationship. At account creation, the institution establishes the customer's identity through prescribed verification procedures. During subsequent account access, authentication controls assess whether the claimant possesses valid credentials. At the transaction stage, however, additional uncertainty may arise concerning whether the person exercising financial authority is genuinely authorized. Although these stages are connected, the controls applicable to each do not necessarily provide equivalent assurance against AI-enabled impersonation [7] [8].

The legal question is therefore not whether CIP is obsolete. Nor does this article argue that banks currently have no duties concerning authentication, fraud detection, or suspicious transactions after account opening. Rather, the concern is whether the existing combination of customer identification, authentication, transaction monitoring, and AML obligations provides a sufficiently coherent legal standard for determining identity assurance at the point of a high-risk transaction when GenAI makes convincing impersonation increasingly scalable.

The resulting gap is therefore not a failure to identify the customer, but a potential failure to connect that established identity with sufficient assurance of authority at a later, higher-risk transaction. Section 4 develops this distinction and Section 7 examines its regulatory implications.

1.4. Research Questions

Against this background, the article addresses one principal research question: Does the existing U.S. banking regulatory framework provide an adequate standard of identity assurance between customer onboarding and transaction execution in cases involving AI-enabled impersonation? Four subsidiary questions follow. First, how do existing U.S. customer identification, due diligence, authentication, AML, and suspicious activity reporting requirements address identity risks arising after account opening? Second, does the emergence of deepfakes, synthetic media, and AI-enabled impersonation expose a regulatory gap between the identification of the customer and verification or authentication of the person controlling a subsequent transaction? Third, what can banking regulators learn from the risk-based approach reflected in NIST SP 800-63-4 and SP 800-63A-4, particularly their treatment of forged media and injection attacks? Finally, should U.S. financial regulation develop a more explicit transaction-level identity assurance require-

ment for defined categories of high-risk transactions, rather than relying primarily on identity proofing at onboarding combined with existing authentication and transaction-monitoring controls? These questions are deliberately framed around high-risk transactions rather than all banking activity. Requiring renewed identity proofing for every ordinary transaction would be costly, intrusive, and operationally unrealistic. The relevant regulatory question is whether stronger identity assurance should be triggered where the nature, value, channel, behavioral pattern, or other risk indicators associated with a transaction create a materially increased risk of impersonation.

1.5. Objectives and Contribution of the Study

The first objective of this article is to clarify the legal distinction between customer identification and transaction-level identity assurance. These concepts are closely connected but should not be treated as interchangeable. Establishing who owns an account does not necessarily establish who is controlling it at every later point in time.

The second objective is to examine whether GenAI has widened the practical significance of this distinction. FinCEN's treatment of identity exploitation across account creation, account access, and transaction processing provides an important regulatory foundation for this inquiry [2] [3].

Third, the article examines the relationship between banking regulation and emerging technical standards. NIST's 2025 Digital Identity Guidelines are especially relevant because they expressly recognize threats from forged media and injection attacks and introduce technical controls intended to mitigate them. The article considers whether similar risk principles can inform banking regulation without converting technical guidance designed for digital identity systems into an inflexible legal mandate for financial institutions.

The principal conceptual contribution of the study is the development of the Identity-to-Transaction Gap as a framework for analysing identity risk across the lifecycle of a banking relationship. Building on this framework, the article develops a proposed Transaction-Level Identity Assurance (TLIA) model. TLIA does not require a bank to re-identify every customer before every transaction. Instead, it proposes a risk-based approach under which enhanced identity assurance would be triggered by specified indicators associated with higher-risk transactions. Such indicators could include unusual transaction value, anomalous account behaviour, changes in devices or credentials, suspicious remote instructions, indicators of synthetic media, or other circumstances suggesting an increased risk that the person controlling the transaction is not the verified customer.

The proposed approach therefore seeks to move the regulatory discussion from a binary question—whether a customer has or has not been identified—to a more functional question: when should a financial institution be required to obtain additional assurance that a high-risk transaction is actually being controlled or authorized by the person whose verified identity is attached to the account?

2. Research Design, Methodology and Evidentiary Limits

This study adopts a doctrinal and policy-oriented legal research methodology. It examines the existing U.S. legal and regulatory framework governing customer identification, identity verification, anti-money laundering compliance, authentication, and suspicious transaction reporting. The analysis focuses particularly on the Bank Secrecy Act framework, the Customer Identification Program rules, relevant FinCEN materials, and applicable federal regulatory guidance.

The doctrinal analysis is supplemented by an examination of technical standards where those standards help explain how the underlying risk has changed. Particular attention is given to NIST SP 800-63-4 and SP 800-63A-4. NIST's framework is not treated as a substitute for banking law. Rather, it is used as a technical benchmark against which the assumptions embedded in existing legal requirements can be evaluated.

The study also employs a functional gap analysis. Instead of examining customer identification, account authentication, and transaction monitoring as entirely separate compliance functions, it maps regulatory protections across three stages identified in FinCEN's identity-related analysis: account creation, account access, and transaction processing. This approach makes it possible to identify areas in which an institution may satisfy an obligation at one stage while remaining exposed to identity manipulation at a later stage [2] [3].

The research is primarily limited to the United States banking regulatory framework. Comparative references to other jurisdictions may be used where they provide a useful regulatory contrast, but the purpose is not to conduct a comprehensive comparative study. Similarly, the article does not attempt to provide a technical evaluation of particular deepfake-detection products or biometric systems. Technology is considered only to the extent necessary to understand the legal and regulatory problem.

The methodological inquiry is confined to whether existing legal duties and technical standards can be coordinated at the point of heightened impersonation risk. It does not infer a new binding obligation from NIST guidance or from the incidents considered below.

Finally, the article does not assume that every transaction requires continuous biometric surveillance or repeated identity proofing. Such an approach could impose disproportionate compliance costs, create privacy concerns, increase false rejections, and unnecessarily interfere with ordinary banking activity. The inquiry is narrower: whether the emergence of scalable AI-enabled impersonation justifies a risk-based legal framework that connects verified customer identity more directly with the authorization of high-risk transactions.

In this sense, the central argument of the article is relatively simple. The first generation of digital banking regulation concentrated heavily on establishing who the customer is. The emerging challenge is ensuring, in appropriate circumstances, that the person acting through that customer's account remains sufficiently connected to that verified identity. Closing that distance—between iden-

tity and transaction—is the regulatory problem this article seeks to examine.

3. AI-Enabled Threat Typology and Documentary Case Studies

Identity fraud existed long before generative AI. However, GenAI has made impersonation more accessible, scalable, and convincing. Fraudsters can now create or manipulate identity documents, images, video, voice, and personalized communications, increasing the difficulty of distinguishing a genuine customer from an impersonator.

3.1. AI-Generated Documents and Synthetic Identities

GenAI can be used to create or modify identity documents and combine fabricated information with stolen personal data. FinCEN has warned that fraudulent identity documents and GenAI-generated media may be used to circumvent identity verification and authentication controls [1].

Synthetic identity fraud presents a related risk by combining real and fabricated information to construct an apparently credible identity. These attacks primarily challenge the reliability of identity proofing at account creation.

3.2. Deepfake Video, Voice Cloning, and Social Engineering

Deepfake technology can imitate a customer's face or voice, while GenAI can produce highly personalized communications for social-engineering attacks. These capabilities may undermine remote verification processes that rely on facial images, video, telephone communication, or other evidence of customer presence.

The emerging problem is therefore not simply whether a document, face, or voice appears genuine, but whether the institution can reasonably trust that the evidence originates from the genuine customer.

3.3. Account Takeover after Legitimate Onboarding

Account takeover illustrates a particularly significant post-onboarding vulnerability. A legitimate customer may successfully complete identity verification and establish an account, after which an unauthorized actor obtains access through compromised credentials, phishing, or social engineering. In such circumstances, the institution's original customer identification records may remain accurate even though control of the account has shifted to a different person. The resulting vulnerability therefore concerns the reliability of subsequent authentication and transactional authority rather than the accuracy of the initial identification process.

3.4. AI-Assisted Fraud at the Transaction Stage

AI-assisted impersonation may become particularly consequential when a compromised account is used to initiate a high-risk transaction. Even where the customer's original identity was correctly verified, an unauthorized actor may exploit synthetic voice, video, or manipulated communications to satisfy additional con-

firmation procedures. This creates uncertainty about whether the apparent authorization genuinely originates from the customer or another authorized person. The appropriate regulatory response is not universal re-identification but proportionate additional assurance where credible indicators of impersonation are present.

3.5. Evidence and Case-Study Methodology

To address the evidentiary limitations of purely doctrinal analysis, this revision adds a structured qualitative examination of publicly documented incidents and official threat advisories. The material was purposively selected for its relevance to the distinction between 1) customer identification, 2) authentication of access, and 3) authorization of a consequential payment. The evidence base consists of the Hong Kong SAR Government's June 2024 account of a deepfake-enabled corporate payment fraud, the FBI's December 2024 warning concerning generative-AI financial fraud, and FinCEN's November 2024 alert concerning deepfake media targeting financial institutions [9]-[11]. These are documentary case materials, not interviews, banking transaction logs, controlled experiments, or a representative sample of incidents. The analysis consequently tests conceptual plausibility and identifies control-design implications; it does not establish prevalence, causal effectiveness, or the success rate of TLIA.

3.6. Case Study I: Hong Kong SAR Deepfake Corporate-Payment Fraud

The Hong Kong SAR Government reported that police received a complaint in late January 2024 after an employee was induced to authorize transfers totalling approximately HK$200 million to five local bank accounts. According to the official account, the deception began with a phishing email purporting to come from the chief financial officer of the employee's UK head office. The employee then attended what appeared to be a video conference with senior colleagues. Investigators believed the pre-recorded meeting had been fabricated using publicly available video and voice material; the subsequent transfer instructions continued by instant messaging [9]. The reported event is a concrete example of synthetic-media-based apparent authority at the payment-instruction stage. It is not evidence that the receiving banks' CIP procedures failed, that a U.S. bank violated an applicable rule, or that a specific bank authentication product was defeated. The corporate employee, rather than an account-takeover intruder, reportedly initiated the transfers. Thus, the episode illustrates a boundary of the proposed theory: a genuine authorized user can be socially engineered into a fraudulent payment, and simply re-authenticating that user would not establish the legitimacy of the underlying instruction. For TLIA to add value in analogous corporate-payment settings, escalation must incorporate independent confirmation of the instruction's origin, beneficiary-change controls, and separation of duties, rather than treating biometric similarity as proof of authorization.

3.7. Case Study II: FBI-Documented AI-Enabled Bank Account Access

The FBI's public service announcement of 3 December 2024 identifies several ways generative AI can facilitate financial fraud. Of direct relevance, the Bureau reports that criminals use AI-generated audio to obtain access to bank accounts by impersonating individuals; it also describes fraudulent identity documents, convincing synthetic images, and real-time video impersonation of executives [10]. The notice supplies a documented law-enforcement threat typology rather than a named adjudicated bank incident with disclosed transaction logs. It therefore supports the plausibility of the account-access and identity-assurance threat but does not quantify the probability of successful compromise. In the framework proposed here, a cloned-voice attempt at account recovery should trigger an assessment of whether the channel and authenticator have become unreliable; confirmation by a previously bound, independently authenticated channel is preferable to another voice-only challenge. The case also demonstrates that "deepfake detection" cannot be the sole proposed control: detection may be uncertain, while independent possession-based or cryptographic evidence can reduce reliance on imitable speech.

3.8. Case Study III: FinCEN Suspicious-Activity Indicators

FinCEN's alert FIN-2024-Alert004, dated 13 November 2024, describes an increase in suspicious activity reporting involving suspected deepfake media, especially fraudulent identity documents used to circumvent verification and authentication. Its red flags include use of a third-party webcam plugin, attempts to switch communication channels during remote verification, reluctance to use multifactor authentication, inconsistency between device or geographic data and identity documents, and unusual rapid movement of funds through newly opened accounts [11]. These indicators are not findings that each flagged customer committed fraud: each requires contextual assessment and appropriate investigative judgment. They nonetheless show how the proposed TLIA trigger can be operationalized without imposing re-identification on every transaction. Where suspicious synthetic media coincides with a new payee, recent account recovery, or high-risk payment, the institution may require independent confirmation, delay execution for investigation where lawful, or refer the activity to its established fraud and BSA/AML processes. The FinCEN alert itself does not impose TLIA as a new binding legal duty.

3.9. Cross-Case Synthesis and Limits of Inference

The three sources illuminate different failure modes. The Hong Kong SAR incident concerns false authority presented to a genuine employee; the FBI notice concerns imitated identity used for account access and payment deception; and FinCEN identifies observable indicators across identity verification, authentication and transactional behaviour. A useful distinction follows: transactor authen-

tication establishes whether a claimant controls a recognized credential, whereas payment-intent verification addresses whether the instruction reflects a genuine, uncoerced and correctly understood intention. TLIA should therefore be framed as a decision protocol integrating both identity/authority assurance and fraud-context review, not as a claim that one additional biometric check solves authorized-push-payment scams. The documentary evidence does not demonstrate that TLIA would have prevented any of these incidents. It supports a testable regulatory design hypothesis that should be assessed prospectively with false-positive rates, payment delay, accessibility impacts, intervention rates and fraud losses.

4. Conceptual Framework: Identity and Transactional Authority

The Identity-to-Transaction Gap describes the potential separation between a bank's knowledge of its customer's verified identity and its assurance that the person later controlling a transaction is that customer or another legitimately authorized person. This distinction becomes increasingly important where AI-generated documents, deepfakes, voice cloning, or compromised credentials can be used to imitate a legitimate customer.

4.1. Identity Verification vs. Transaction Authentication

Identity verification and authentication serve different purposes. Identity verification establishes whether an applicant is genuinely connected to a claimed identity, while authentication provides assurance that a later claimant controls an authenticator associated with the account.

Accordingly, successful KYC answers: "Who is the customer?" A later transaction may raise a different question: "Who is acting through the account now?" The distinction matters because a customer's identity may remain completely accurate even after an unauthorized person obtains control of the account.

4.2. Identity Exploitation across the Customer Lifecycle

FinCEN's lifecycle-based analysis distinguishes identity exploitation at three stages: account creation, account access, and transaction processing. At the first stage, criminals may exploit stolen or synthetic identities to establish financial relationships. At the second, unauthorized actors may compromise access to accounts belonging to legitimately identified customers. At the third, compromised access may facilitate unauthorized financial transactions. GenAI can intensify vulnerabilities across these stages by generating convincing documents, biometric representations, and communications.

Taken together, these documentary sources illustrate three analytically distinct risks: fraudulent instructions to an authenticated user, impersonation during account access, and warning indicators observable by financial institutions. They support a proposal for risk-sensitive escalation, not a claim that an additional identity check would prevent every fraud.

Synthetic identities primarily challenge the reliability of customer identification, while deepfakes and AI-generated impersonation can affect both onboarding and later authentication. Modern remote verification may rely on documents, facial images, video, biometrics, or other digital evidence. GenAI creates the possibility that such evidence may appear genuine even when the person presenting it is not the legitimate identity holder. The central problem can therefore be expressed as: A genuine identity can be represented by a false actor. This is particularly important after legitimate onboarding because the bank's customer record may remain correct while the person exercising control has changed.

4.3. Why Successful KYC Does Not Guarantee Transaction Authenticity

Consider a customer who successfully completes identity verification and establishes a legitimate banking relationship. If an unauthorized actor subsequently compromises the account and employs AI-generated impersonation to initiate a transaction, the original customer identification process may remain entirely valid. The failure arises from the separation between the verified customer's legal identity and the authority exercised during the subsequent transaction. This illustrates why successful onboarding verification cannot, by itself, guarantee the authenticity of every later transaction.

Thus, verifying a customer's identity does not necessarily establish that the person initiating a transaction is the genuine account holder or an authorized representative. This distinction does not suggest that banks should conduct a fresh Know Your Customer (KYC) verification for every transaction. Equally, successful authentication cannot, by itself, establish that a transaction is legitimate, as even genuine customers may be misled into authorizing fraudulent payments. Against this background, the article raises a more specific question: Where the risk of impersonation associated with a particular transaction increases significantly, should banks take additional steps to verify that the person initiating or authorizing the transaction is genuinely linked to the previously verified customer?

This question forms the conceptual foundation of the Identity-to-Transaction Gap and the proposed Transaction-Level Identity Assurance framework developed later in this article.

5. U.S. Banking Regulation and Existing Safeguards

The Identity-to-Transaction Gap cannot be evaluated without first considering what U.S. banking law already requires. The United States does not regulate customer identity through a single rule. Instead, identity-related obligations are distributed across the Bank Secrecy Act (BSA), the Customer Identification Program (CIP) rules adopted following the USA PATRIOT Act, Customer Due Diligence (CDD) requirements, suspicious activity reporting obligations, and related regulatory guidance.

Together, these rules create a layered framework. Banks must identify custom-

ers, understand the nature and purpose of customer relationships, develop customer risk profiles, monitor activity, and report certain suspicious transactions. The framework is therefore considerably broader than a one-time KYC exercise at account opening.

At the same time, these obligations perform different functions. Identifying a customer, monitoring an account for suspicious activity, and determining whether the person controlling a particular transaction is genuinely connected to the verified customer are not necessarily the same regulatory task. This distinction is important when considering AI-enabled impersonation.

5.1. Bank Secrecy Act and the AML Framework

The foundation of the U.S. anti-money laundering framework is the Bank Secrecy Act of 1970, as subsequently amended. The BSA authorizes the Department of the Treasury to impose recordkeeping, reporting, and other requirements on financial institutions for the purpose of detecting and preventing money laundering and other financial crime.

Over time, the BSA has developed from a primarily recordkeeping and currency-reporting statute into a broader framework for financial intelligence and AML/CFT compliance. Its requirements include, among other matters, recordkeeping, Currency Transaction Reports, Suspicious Activity Reports, customer identification, and institutional AML compliance programs.

The framework was strengthened through several subsequent statutes. The Money Laundering Control Act of 1986 criminalized money laundering at the federal level, while the Annunzio-Wylie Anti-Money Laundering Act of 1992 strengthened the suspicious activity reporting framework. The USA PATRIOT Act of 2001 further expanded the BSA architecture, particularly in relation to customer identification, due diligence, information sharing, and terrorist financing.

For banks, the significance of the BSA is not limited to identifying customers. An effective BSA/AML program requires institutions to maintain systems capable of recognizing financial activity that may indicate money laundering, fraud, terrorist financing, or other unlawful conduct.

This distinction is relevant to the present study. Customer identity is one component of the BSA framework, but the broader system also focuses on the manner in which accounts and transactions are used. The BSA therefore creates an important regulatory bridge between customer identity and customer activity. However, the existence of that bridge does not necessarily answer a more specific question raised by AI-enabled impersonation: whether the bank has adequate assurance that the person controlling a particular high-risk transaction is the verified customer or another legitimately authorized person.

5.2. USA PATRIOT Act and the Customer Identification Program (CIP)

Customer identification became a more explicit component of U.S. banking regulation following the USA PATRIOT Act of 2001. Section 326 required the devel-

opment of minimum standards for financial institutions concerning customer identification and verification.

For banks, the implementing rule is now found principally in 31 C.F.R. § 1020.220 [7]. The central idea of the CIP framework is important to the argument developed in this article. A bank must establish a written Customer Identification Program appropriate to its size and type of business. The program must include risk-based procedures that enable the bank to form a reasonable belief that it knows the true identity of each customer [7] [8]. The bank must generally obtain specified identifying information before opening an account, including the customer's name, date of birth for an individual, address, and identification number. It must then verify the customer's identity to the extent reasonable and practicable through documentary methods, non-documentary methods, or a combination of both.

The rule does not require a bank to prove every element of identifying information with absolute certainty. The regulatory objective is instead to establish sufficient verification procedures for the institution reasonably to believe that it knows the customer's true identity.

This flexible and risk-based formulation has practical advantages. Banks differ considerably in size, customer base, delivery channels, products, and risk exposure. A prescriptive identity-verification process suitable for one institution or product may therefore be inappropriate for another.

The CIP framework also contains an important temporal dimension. Its central operation is closely associated with the establishment of an account relationship. Existing customers may, in defined circumstances, be excluded from the definition of a new "customer" where the bank already has a reasonable belief that it knows the person's true identity.

FinCEN's longstanding CIP guidance also makes clear that CIP is only one component of a bank's broader BSA/AML obligations. Successful implementation of CIP does not, by itself, satisfy suspicious activity reporting or other regulatory duties.

This qualification is important. It would be inaccurate to argue that U.S. banking regulation stops monitoring risk once a customer has successfully completed CIP. It does not. The narrower issue raised by this article is different. The statutory and regulatory language of CIP is primarily directed toward establishing who the customer is. It does not, by itself, establish a comprehensive transaction-by-transaction rule determining when a bank must reconfirm the connection between that verified identity and the individual exercising control over a later transaction.

This difference becomes increasingly important where GenAI allows an unauthorized person to imitate characteristics of a legitimate customer without changing the customer's underlying identity record.

5.3. Customer Due Diligence (CDD) Requirements

Customer Due Diligence extends the regulatory framework beyond basic identification. FinCEN's CDD framework identifies four core elements of customer due

diligence: identifying and verifying customers; identifying and verifying beneficial owners of covered legal entity customers; understanding the nature and purpose of customer relationships for the purpose of developing customer risk profiles; and conducting ongoing monitoring to identify and report suspicious transactions and, on a risk basis, maintain and update customer information [12]-[14]. CDD is therefore particularly relevant to the Identity-to-Transaction Gap because it demonstrates that U.S. AML regulation is not limited to account opening.

The concept of a customer risk profile requires the financial institution to develop an understanding of the nature and purpose of the relationship. This provides a baseline against which subsequent activity can be considered. Where activity materially departs from what would reasonably be expected, that departure may contribute to the institution's assessment of whether further investigation or suspicious activity reporting is appropriate.

The ongoing monitoring requirement is even more significant. Covered institutions must monitor customer activity to identify and report suspicious transactions and must update customer information when appropriate on the basis of risk.

However, the CDD framework should not be interpreted as requiring banks to continuously re-identify every customer. FinCEN's guidance emphasizes a risk-based approach. Updating customer information is generally connected to information discovered through normal monitoring that is relevant to assessing or reassessing the customer's risk.

The beneficial-ownership component of CDD has also evolved. In February 2026, FinCEN granted exceptive relief from the requirement that covered financial institutions identify and verify the beneficial owners of a legal entity customer every time that customer opens a new account. Under the relief, an institution may generally limit renewed beneficial-owner identification and verification to the customer's first account opening, circumstances in which facts reasonably call previous information into question, and situations in which the institution's risk-based ongoing CDD procedures require renewed identification or verification [15] [16].

This development reinforces an important feature of contemporary U.S. AML regulation: risk-based reassessment rather than universal repetitive verification. For the present study, CDD therefore occupies an intermediate position. It clearly extends regulatory responsibility beyond onboarding and requires continuing attention to customer behavior. Yet its principal focus remains customer risk, suspicious activity, and maintenance of relevant customer information. It does not necessarily answer the distinct question of what level of identity assurance must attach to the person controlling a specific transaction when impersonation risk becomes unusually high.

5.4. FinCEN Guidance on Identity-Related Financial Crime

FinCEN's recent work on identity-related financial crime provides particularly important support for examining identity risk as a lifecycle problem rather than

merely an onboarding problem. In January 2024, FinCEN published a Financial Trend Analysis based on BSA reports filed during calendar year 2021. Approximately 1.6 million reports, or 42 percent of roughly 3.8 million BSA reports filed during the review period, were associated with identity-related suspicious activity, representing approximately $212 billion in suspicious activity. The scale of these figures is significant, but the conceptual structure used by FinCEN is even more relevant to this study.

FinCEN's analysis demonstrates that identity-related exploitation extends beyond the establishment of a banking relationship. By examining vulnerabilities associated with account creation, subsequent account access, and transaction processing, the agency provides a regulatory basis for understanding identity risk throughout the customer lifecycle. This perspective supports the article's central argument that successful onboarding verification does not eliminate the need for appropriate safeguards against later impersonation and unauthorized transactional control.

At the transaction-processing stage, identity-related exploitation may affect the movement of funds or other activity conducted through the financial institution. FinCEN's analysis also distinguishes among forms of identity exploitation such as impersonation, circumvention, and compromise and connects these risks to identity validation, verification, and authentication processes. This regulatory framing is important because it supports a central premise of the present article: identity risk continues after the customer has been successfully identified. FinCEN's subsequent alert concerning fraud schemes involving deepfake media strengthens that conclusion. The agency has warned financial institutions about suspicious activity involving the use of GenAI-generated media and fraudulent identity documents to circumvent identity verification and authentication methods.

The importance of this development should not be overstated. FinCEN's alerts and analyses do not themselves create a general legal obligation requiring banks to perform new biometric identity verification before every transaction. Nor do they establish the Transaction-Level Identity Assurance model proposed later in this article.

They do, however, provide strong regulatory evidence that identity exploitation now occurs across multiple stages of the banking relationship and that GenAI can be used against both identity verification and authentication controls.

The issue is therefore no longer simply whether the financial institution initially identified the correct customer. It is also whether existing controls remain sufficiently reliable when a malicious actor can manipulate the signals through which identity and authority are subsequently assessed.

5.5. Banks' Existing Obligations for Monitoring Suspicious Transactions

Any argument that U.S. banking law contains an Identity-to-Transaction Gap must account for existing suspicious activity monitoring and reporting obliga-

tions. Under 31 C.F.R. § 1020.320, banks are required to report certain suspicious transactions through Suspicious Activity Reports (SARs). In general, a transaction conducted or attempted by, at, or through a bank falls within the reporting rule where it involves or aggregates at least $5000 in funds or other assets and the bank knows, suspects, or has reason to suspect circumstances specified in the regulation [17].

These circumstances include transactions involving funds derived from illegal activity or intended to conceal such funds, transactions designed to evade BSA requirements, and transactions that lack an apparent lawful or business purpose where no reasonable explanation is available after examining the relevant facts.

Compliance with this obligation necessarily requires more than passive record-keeping. Banks must maintain systems and procedures capable of identifying activity that may warrant further investigation and reporting.

In practice, suspicious activity monitoring may draw on a wide range of indicators: transaction size, frequency, destination, deviations from expected customer behavior, account history, geographic exposure, device or channel information, and other relevant risk factors. Identity-related anomalies may also become relevant where they suggest account takeover, impersonation, fraud, or laundering.

These obligations significantly reduce the conceptual space in which an Identity-to-Transaction Gap can be claimed. A bank cannot simply argue that it completed KYC at onboarding and therefore has no continuing responsibility for suspicious use of the account.

But suspicious activity monitoring and transaction-level identity assurance are not necessarily interchangeable. Consider again a legitimate customer whose account has been compromised. A fraudster gains access and attempts an unusual high-value transfer while using AI-generated voice or video to imitate the customer. A transaction-monitoring system may identify the payment as anomalous because of its value, destination, timing, or inconsistency with previous activity. This may trigger investigation, intervention, or eventually a SAR. The identity question, however, is slightly different:

What level of assurance should the institution obtain, before or during execution of a sufficiently high-risk transaction, that the person controlling the transaction is genuinely connected to the verified customer identity?

SAR obligations are primarily designed to ensure that suspicious activity is detected and reported. Transaction-level identity assurance, by contrast, would concern the circumstances in which the institution should strengthen the connection between customer identity and transactional authority before relying on the apparent identity of the person initiating or confirming the transaction.

This distinction also highlights a temporal difference. Suspicious activity reporting can operate as a detection and reporting mechanism surrounding or following suspicious conduct. A transaction-level assurance control may instead operate as a preventive measure at the moment at which a transaction is being au-

thorized or processed.

The distinction should nevertheless be treated cautiously. Existing authentication, fraud prevention, cybersecurity, and risk-management requirements may already address parts of this problem, and the regulatory framework should not be artificially divided merely to create a research gap.

The relevant inquiry is therefore one of regulatory coherence rather than regulatory absence. The United States already has a substantial framework governing customer identification and suspicious financial activity. CIP establishes a risk-based standard for forming a reasonable belief about the customer's true identity. CDD extends obligations into customer-risk assessment and ongoing monitoring. SAR requirements require banks to identify and report qualifying suspicious activity. FinCEN's recent analyses expressly recognize identity exploitation at account creation, account access, and transaction processing [2] [3].

Yet these obligations developed for different regulatory purposes. The question raised by GenAI is whether their combined operation produces a sufficiently clear standard for circumstances in which a bank has correctly identified its customer but faces evidence that the person controlling a high-risk transaction may be an AI-assisted impersonator. The issue can therefore be summarized as follows: CIP asks: Who is the customer? CDD asks: What is the nature and risk of the customer relationship, and has relevant information changed?

Transaction monitoring asks: Is the activity suspicious? The Identity-to-Transaction question asks: When impersonation risk is materially elevated, who is actually exercising authority over this transaction, and what level of assurance should the bank be expected to obtain before relying on that apparent authority? It is the relationship among these questions—not the absence of U.S. AML regulation—that forms the legal foundation for the analysis that follows.

6. Digital Identity Standards and NIST

The cybersecurity significance of this analysis lies in connecting digital-identity threats with transaction-stage banking decisions. Rather than assessing a particular detection product, this section examines how identity proofing, authentication assurance, resistance to forged media, and risk-sensitive escalation can inform institutional responses to suspected impersonation [18] [19].

NIST's Digital Identity Guidelines provide an important technical framework for understanding modern identity risks. The 2025 Revision 4 distinguishes identity proofing (IAL) from authentication (AAL) and applies different levels of assurance according to risk [4]-[6].

The revised guidelines also respond to emerging threats from deepfakes, forged media, and digital injection attacks. NIST SP 800-63A-4 recognizes that remote identity systems must consider whether digital evidence, including images and video, is genuine and originates from a trustworthy source [5].

A key lesson from NIST is that identity assurance should not be treated as a single event. Identity proofing establishes the claimed identity, while authentica-

tion provides later assurance that the claimant controls an authenticator associated with that identity. Higher risks may justify stronger controls.

NIST does not itself impose a banking-specific Transaction-Level Identity Assurance (TLIA) duty. Instead, its risk-based assurance principles provide a technical basis for the distinct regulatory proposal developed here: banks could calibrate additional transaction-stage confirmation to credible indicators of impersonation, without treating technical guidance as binding banking law.

NIST's risk-based approach recognizes that the strength of identity proofing and authentication controls should reflect the consequences and likelihood of identity-related failures. Although the Digital Identity Guidelines do not establish a banking-specific obligation to verify transactional authority, their assurance-based principles provide a useful technical foundation for considering proportionate additional controls where impersonation risk is materially elevated.

7. The Identity-to-Transaction Regulatory Integration Problem

The preceding sections demonstrate that U.S. banks operate within a substantial regulatory framework governing customer identification, due diligence, authentication, and suspicious activity. It would therefore be misleading to describe AI-enabled impersonation as an entirely unregulated problem. Banks are not permitted to identify a customer at onboarding and then disregard the manner in which the account is subsequently used.

The more difficult question is whether these existing obligations collectively establish a sufficiently clear standard for a different problem: what should a bank be expected to do when it knows the identity of its customer but has reason to doubt that the person exercising control over a high-risk transaction is that customer or another legitimately authorized person?

The relevant problem is one of regulatory coordination. CIP explicitly addresses the establishment of customer identity; CDD, authentication, fraud prevention, and monitoring address subsequent risks through different mechanisms. Yet those mechanisms do not always express a common decision rule for escalating assurance when credible, transaction-specific indicators suggest that the apparent transactor lacks authority. GenAI makes this coordination question more pressing because convincing identity signals can be fabricated without altering the bank's accurate customer record.

7.1. The CIP "Reasonable Belief" Standard

The starting point is the Customer Identification Program rule. Under 31 C.F.R. § 1020.220, a bank's CIP must contain risk-based procedures for verifying the identity of each customer to the extent reasonable and practicable. Those procedures must enable the institution to form a reasonable belief that it knows the true identity of each customer [7]. The standard is deliberately risk-based rather than absolute. Banks are required to design verification procedures in light of relevant

factors, including the types of accounts offered, methods of account opening, identifying information available, and characteristics of the institution and its customer base.

The reasonable-belief standard is tailored to the institution and the circumstances of account opening. Its limitation for this inquiry is functional: successful CIP establishes the customer's identity but does not conclusively determine control of a later payment. A subsequent credential compromise therefore implicates access and authorization safeguards rather than retroactively invalidating CIP.

Accordingly, the limits of CIP should be understood as limits of regulatory function rather than proof that customer-identification law is defective. The subsequent inquiry belongs to the interaction of authentication, fraud controls, and risk-based transaction decisions.

7.2. When Does a Bank's Identity Verification Duty End?

Asking when a bank's identity-verification duty "ends" requires caution because U.S. regulation does not support a simple answer such as "when the account is opened." CIP is strongly associated with account opening and establishment of the customer relationship, but the broader BSA/AML framework imposes continuing obligations. Customer Due Diligence requires banks to understand the nature and purpose of customer relationships, develop customer risk profiles, conduct ongoing monitoring to identify and report suspicious transactions, and, on a risk basis, maintain and update customer information [12] [14].

The ongoing nature of CDD is important. A material change discovered through normal monitoring may require customer information to be updated and the customer's risk profile to be reconsidered. Accordingly: CIP verification is not equivalent to the entirety of the bank's continuing customer-related obligations. Yet continuing CDD should not automatically be equated with continuous identity verification either. The FFIEC framework describes updating customer information as risk-based and event-driven. Its purpose includes maintaining an appropriate understanding of the customer relationship and supporting the identification of suspicious activity. It does not impose a general requirement that a customer repeat full identity proofing every time the customer conducts a transaction.

Universal repetition of documentary or biometric proofing for ordinary card payments, withdrawals, transfers, and inquiries would be costly and potentially exclusionary. The proportionate alternative is to reassess assurance when transaction-specific evidence materially changes the risk of impersonation.

An event-driven approach can preserve ordinary authentication for familiar, low-risk activity while requiring a stronger and preferably independent basis for authority where indicators of account takeover, synthetic media, or anomalous account recovery converge. The question is not when all identity obligations end, but when a change in risk warrants an escalation of assurance. CDD, fraud monitoring, and authentication each supply relevant information, yet none alone nec-

essarily resolves the payment-stage authority question.

The remaining policy question is whether materially elevated, transaction-specific impersonation risk warrants an explicit expectation of independent assurance before execution.

7.3. Verified Customer vs. Actual Person Controlling the Account

Three scenarios locate the point of failure. An impostor opening an account with synthetic evidence primarily challenges initial identity proofing; a criminal accessing a legitimate account with stolen credentials challenges authentication; and an actor using cloned voice or manipulated video to confirm a consequential payment challenges the reliability of evidence of transactional authority. These are distinct vulnerabilities, and a transaction-stage response should address the specific failure rather than repeat onboarding checks indiscriminately.

The third scenario illustrates why a familiar channel or apparently convincing identity signal may not establish genuine transactional control. Synthetic media can compromise voice, video, document, and written-message cues simultaneously; nevertheless, the existence of such techniques does not justify treating every transaction as suspect.

The legal concern is thus confined to cases in which established identity and effective control may diverge, particularly when credible indicators make the consequences of relying on ordinary authentication unusually serious.

7.4. Why Transaction Monitoring Alone May Be Insufficient

One important objection to the proposed Identity-to-Transaction Gap is that banks already have systems in place to detect suspicious transactions. This is a legitimate concern and deserves careful consideration. Under the existing Bank Secrecy Act and Anti-Money Laundering (BSA/AML) framework, banks are required to maintain appropriate systems for identifying and reporting suspicious activity. Risk-based Customer Due Diligence (CDD) helps banks understand their customers and their expected financial behavior, while transaction-monitoring systems are designed to identify activities that depart from those expectations.

In practice, these systems may detect unusual transaction amounts, unfamiliar beneficiaries, changes in transaction frequency, unexpected locations, or other patterns associated with financial crime. For example, a transaction involving a newly added beneficiary, an unfamiliar device, or the rapid transfer of funds may raise concerns about possible account takeover. Depending on the circumstances, such activity may prompt further investigation, customer contact, internal escalation, transaction restrictions, or suspicious activity reporting. Transaction monitoring therefore remains an important safeguard against fraud and money laundering.

Transaction monitoring and identity assurance address related but distinct risks: the former identifies activity requiring scrutiny, while the latter asks whether apparent authorization of a consequential payment is sufficiently trustworthy. A

monitoring alert may supply a reason to escalate assurance, but it is not itself proof of unauthorized control.

This distinction becomes particularly important in cases involving sophisticated impersonation. An unauthorized person who gains access to a customer's account may deliberately conduct transactions that resemble the customer's normal financial behavior. By keeping transaction amounts within familiar limits, using previously known beneficiaries, or imitating established payment patterns, the impersonator may avoid triggering conventional monitoring alerts. In such circumstances, a transaction may appear entirely consistent with the customer's financial profile even though the person exercising control over the account is unauthorized.

The opposite situation is equally possible. A genuine customer may initiate a large or unusual transaction for legitimate reasons, causing the monitoring system to generate an alert despite the absence of any identity compromise. These examples illustrate two related but distinct propositions: a suspicious transaction does not necessarily indicate an unauthorized transactor, and a transaction initiated by an unauthorized person will not necessarily appear suspicious.

Accordingly, a fraudulent payment can resemble normal activity and a legitimate payment can look anomalous. Behavioral signals should inform, rather than substitute for, an appropriately independent assessment of authority when the risk warrants it.

The distinction also reflects differences in regulatory and operational objectives. Suspicious activity monitoring under the BSA/AML framework is primarily directed toward identifying, investigating, and reporting potentially unlawful or unusual financial activity. Transaction-level identity assurance may serve a more immediate preventive purpose by determining whether additional confirmation of the transactor's identity or authority is necessary before a transaction is executed. These functions should not be treated as competing approaches. Instead, they can operate together as complementary elements of a risk-based banking control framework.

Indeed, existing transaction-monitoring systems may provide a practical starting point for determining when enhanced identity assurance is appropriate. Requiring banks to repeat full identity verification before every transaction would be unnecessarily burdensome and could interfere with ordinary banking operations. A more proportionate approach would distinguish between routine transactions and those presenting materially elevated risks of impersonation. Transactions conducted through familiar devices, established channels, and patterns consistent with the customer's previous activity may generally proceed through ordinary authentication procedures. However, where additional warning signs emerge, such as suspicious account recovery attempts, unusual device activity, or credible indications of AI-generated voice or video impersonation, further independent confirmation may be justified before the bank relies on the apparent authorization.

The legal distinction between suspicious activity reporting and preventive au-

thentication reinforces the need for regulatory coordination. Section 5318(g) of Title 31 and 31 C.F.R. § 1020.320 establish reporting obligations for qualifying suspicious activity, whereas authentication and access-security controls serve a different operational purpose. These mechanisms may nevertheless reinforce one another when transaction-monitoring indicators reveal possible impersonation. In such circumstances, the institution could use the identified risk to determine whether independent confirmation of transactional authority is appropriate before execution, subject to applicable payment rules and operational constraints. For example, suppose a bank customer regularly transfers funds to a small group of beneficiaries. An impersonator gains access to the customer's online banking account using compromised credentials and attempts to transfer money to a newly added beneficiary. The bank's transaction-monitoring system may flag the payment because it differs from the customer's usual transaction patterns. Rather than relying solely on the successful login, the bank could require additional confirmation through a previously verified communication channel before processing the payment. This example illustrates how transaction monitoring can identify potential risks, while enhanced identity assurance can help prevent an unauthorized transaction from being completed.

7.5. The Identity-to-Transaction Gap as a Regulatory Blind Spot

The proposed "regulatory blind spot" is not the absence of identity, authentication, fraud-prevention, or suspicious-activity duties. It is the lack of a single, clearly articulated banking-specific principle for coordinating those duties when credible impersonation indicators arise at the payment stage. The argument therefore concerns integration and timing, not an unregulated category of fraud. The normative question is whether supervisory guidance should make the escalation decision more explicit while preserving the discretion, proportionality, and legal obligations already embedded in existing controls.

The proposed TLIA framework differentiates among ordinary, elevated-risk, and high-impersonation-risk transactions. Ordinary transactions may proceed under existing authentication controls where no material risk indicators are present. Transactions presenting elevated risks may require additional authentication or verification, while those involving credible indicators of account takeover or AI-assisted impersonation should trigger enhanced, independently supported assurance measures. This graduated approach would align the intensity of institutional controls with the nature and seriousness of the identified threat.

The TLIA model is summarized in **Figure 1**. The framework begins with established customer identification and ordinary account use. Transaction-specific risk assessment then evaluates contextual anomalies and evidence of impersonation. Materially elevated risk activates a proportionate escalation decision, which may require independent assurance before the institution approves, delays, rejects, or refers the transaction. A low-risk transaction follows existing controls without additional re-identification.

The respective functions should remain separate: CIP identifies the customer; authentication protects account access; monitoring detects unusual activity; and investigation and reporting address suspected misconduct. TLIA would connect the information generated by those functions to a documented decision about transactional authority when a material risk trigger is present.

Calling it a regulatory "blind spot" does not mean regulators are unaware of account takeover, deepfakes, or identity fraud. FinCEN's recent analysis demonstrates the opposite. U.S. regulators expressly recognize identity exploitation at account creation, account access, and transaction processing, while contemporary supervisory guidance requires ongoing monitoring and risk-based controls [2] [3]. The corresponding policy question is whether regulators should clarify how banks combine transaction context, indicators of impersonation, and independent confirmation. Any clarification should be technology-neutral, compatible with existing legal duties, and sensitive to the cost of false positives.

The policy principle can be described as "Know Your Transactor" in defined high-risk circumstances. This is not a universal KYC requirement or a justification for continuous biometric surveillance. Instead, it calls for proportionate assurance of transactional authority where credible impersonation indicators make ordinary authentication unreliable.

A regulatory response could therefore focus on defined risk triggers rather than fixed transaction values alone. Relevant indicators might include unusual transaction behavior, new beneficiaries, account-recovery activity, recent credential changes, anomalous device or network information, evidence of synthetic media, repeated authentication anomalies, or combinations of factors suggesting possible account takeover. Once a defined risk threshold is reached, additional assurance could be required before the institution relies on the apparent customer's authority. The precise form of that assurance should remain technologically flexible. Regulation should not assume that facial recognition, voice verification, or any other single technology will remain reliable. The legal obligation should focus on the required outcome and level of assurance, while technical standards evolve over time. This approach also creates a useful connection with NIST's risk-based model without incorrectly treating NIST guidance as banking law.

The resulting claim is deliberately limited. Existing U.S. rules address customer identity and suspicious activity extensively, but do not articulate with equivalent specificity a common transaction-stage escalation principle for credible AI-assisted impersonation indicators. Whether a TLIA-style clarification improves fraud outcomes must be tested empirically.

8. Legal Consequences and Comparative Perspectives

The Identity-to-Transaction Gap has consequences beyond customer identification. Where an unauthorized actor uses AI-generated voice, video, documents, or other synthetic media to control an account or initiate a transaction, questions arise concerning liability, AML compliance, suspicious activity reporting, and the

standard of care expected from financial institutions.

8.1. Liability for AI-Enabled Unauthorized Transactions

AI-enabled impersonation complicates the allocation of loss because the customer may have been properly identified while the later transaction was initiated by an unauthorized actor. Existing U.S. payment law already contains rules governing unauthorized transactions, including protections applicable to certain consumer electronic fund transfers under the Electronic Fund Transfer Act and Regulation E. The use of AI does not automatically create a new category of liability [20].

The more difficult issue arises where an institution relies on apparently valid authentication that has been defeated through sophisticated impersonation. As deepfake and synthetic-media risks become increasingly recognized, disputes may focus not only on whether authentication occurred, but also on whether the authentication method was reasonable in the circumstances.

8.2. AML/CFT and Suspicious Activity Reporting

AI-enabled impersonation is also relevant to AML/CFT because compromised or synthetically created accounts may be used to receive, transfer, or layer proceeds of fraud. FinCEN has specifically advised financial institutions of suspicious activity associated with deepfake media and has identified red flags that may assist institutions in detecting such schemes [1].

Existing SAR obligations remain applicable where the regulatory requirements for reporting are satisfied. The important point for this article is that SAR reporting and identity assurance perform different functions. SAR requirements primarily support the detection and reporting of suspicious activity; transaction-level identity assurance would seek, where appropriate, to reduce impersonation risk before a high-risk transaction is completed.

8.3. Compliance and Standard of Care

The growing recognition of deepfake and injection attacks may also affect expectations concerning reasonable institutional controls. A security procedure that was considered adequate when synthetic impersonation was technically difficult may not necessarily remain adequate as the threat becomes widely known and technically accessible.

This does not mean that every bank should be legally required to deploy the newest biometric or AI-detection technology. Such a requirement would quickly become outdated. A more sustainable approach would evaluate whether an institution maintains risk-appropriate controls in light of reasonably foreseeable threats, transaction characteristics, and available security measures.

8.4. Privacy and Civil-Liberty Concerns

Stronger identity assurance can itself create risks. Continuous facial recognition,

voice monitoring, behavioral profiling, or repeated biometric verification could increase collection of sensitive information and create privacy, cybersecurity, accessibility, and financial inclusion concerns.

For this reason, the Identity-to-Transaction Gap should not be addressed through continuous surveillance. The proposed model should instead follow principles of necessity, proportionality, data minimization, and risk-based escalation. Enhanced assurance should be triggered by materially elevated risk rather than imposed on every customer and every transaction.

The United States is not alone in confronting the relationship between digital identity, authentication, and payment fraud. The UK and European Union provide useful comparative lessons, although their regulatory structures differ significantly from the U.S. system.

8.5. U.S. Approach

The U.S. framework is comparatively fragmented. Customer identification is principally addressed through CIP and CDD, suspicious activity through the BSA/AML framework, and payment authentication and fraud risk through additional statutory, regulatory, supervisory, and private-sector controls.

The strength of this model is flexibility. Its weakness is that identity assurance between onboarding and transaction execution is not expressed through a single regulatory principle. This fragmentation is particularly relevant to the Identity-to-Transaction Gap.

8.6. Lessons from the UK and EU

European payment regulation provides a useful contrast because Strong Customer Authentication (SCA) has been incorporated more explicitly into the regulation of electronic payments under the EU's payment-services framework. SCA generally relies on authentication using independent elements drawn from knowledge, possession, and inherence, subject to applicable exemptions.

The UK retained a broadly comparable SCA framework after leaving the EU. These systems do not eliminate deepfake fraud or account takeover, nor should they simply be transplanted into U.S. law. They nevertheless demonstrate that authentication requirements can be connected more directly to payment risk rather than treated solely as an account-opening issue.

The comparative lesson is therefore one of regulatory structure, not regulatory copying: transaction-stage assurance can be expressly incorporated into financial regulation while remaining proportionate and risk-based.

8.7. Technology-Neutral or Technology-Specific Regulation?

A future U.S. response should generally favor technology-neutral regulation supported by adaptable technical standards. A rule requiring a specific deepfake detector, facial-recognition system, or biometric technology could become obsolete quickly as both fraud techniques and defensive technologies evolve. Regulation

should instead specify the required outcome—for example, an appropriate level of assurance where a transaction presents materially elevated impersonation risk. Technical standards such as NIST guidance could then assist institutions in determining how that outcome should be achieved.

A technology-neutral standard could specify risk triggers, the objective of independent confirmation, and documentation of the resulting decision, while leaving institutions flexibility to select accessible and effective controls.

9. Transaction-Level Identity Assurance: Proposed Model

The analysis above suggests that the appropriate response to the Identity-to-Transaction Gap is not another universal KYC requirement. Requiring banks to re-identify customers before every transaction would be costly, intrusive, and disproportionate to the actual risk.

This article instead proposes Transaction-Level Identity Assurance (TLIA): a risk-based framework under which banks would apply enhanced identity or authority assurance when a transaction presents materially elevated indicators of impersonation or account takeover.

For clarity, TLIA is a proposed decision framework rather than an existing U.S. legal requirement or a tested fraud-prevention system. It connects established customer identity, transaction context, independent assurance where warranted, and a documented decision about execution.

Customer Identification Programs remain essential to establishing who the customer is, but a valid onboarding record cannot settle every later question of account control. The additional inquiry proposed here is whether, in a materially higher-risk transaction, the bank has sufficient independent evidence that the apparent transactor is the verified customer or another authorized actor. TLIA therefore builds on, rather than replaces, initial identification and ongoing controls.

Because identity risk changes after onboarding, TLIA would use evidence of compromised credentials, suspicious account recovery, or synthetic-media impersonation to determine when a previously established identity requires additional transaction-stage assurance. The objective is to identify a material change in authority risk, not to repeat the original CIP procedure.

9.1. Risk-Based Transaction Authentication

TLIA should not apply equally to every transaction. Its central principle should be proportionality. Routine activity consistent with the customer's normal behavior could continue under ordinary authentication procedures. Enhanced assurance would arise only when specified indicators suggest materially elevated impersonation risk. Potential triggers could include unusual transaction value, a new beneficiary, abnormal device or network characteristics, recent password or credential changes, unusual account-recovery activity, significant departure from established customer behavior, repeated authentication anomalies, or indicators sug-

gesting possible synthetic-media use.

No single indicator necessarily requires intervention. Under the proposed risk-based model, ordinary transactions may proceed using existing authentication; elevated-risk activity may require additional confirmation; and credible indicators of impersonation may justify independently supported assurance, a temporary delay where lawful, or further investigation. The selected response should be proportionate to the risk and the transaction context.

This approach would preserve convenience for ordinary customers while concentrating stronger controls where potential harm is greater.

9.2. Enhanced Verification for High-Risk Transactions

When a transaction crosses a defined risk threshold, the institution should obtain additional assurance through an independent or stronger mechanism. The precise method should not be prescribed permanently in legislation. Depending on the circumstances, measures could include stronger multi-factor authentication, confirmation through a previously trusted device or channel, cryptographically protected authentication, independent confirmation of unusual instructions, or other controls capable of providing increased assurance. The objective is not simply to add more authentication steps. Multiple weak controls do not necessarily produce strong assurance, particularly where the same compromised channel is used repeatedly.

For example, where a suspicious telephone instruction creates the initial concern, confirming the transaction through another telephone interaction based only on voice recognition may provide limited additional protection against voice cloning.

TLIA should therefore favor independent and risk-appropriate confirmation rather than repetitive use of the same potentially compromised evidence.

9.3. Deepfake and Injection-Attack Standards

Regulation should also recognize that synthetic media may attack the verification process itself. Where banks use remote facial, video, documentary, or biometric verification for high-risk activities, their systems should incorporate reasonable protections against forged media and digital injection attacks. NIST SP 800-63A-4 provides a useful technical benchmark. Its 2025 framework expressly addresses forged media, digital injection, and the need to establish confidence in the integrity and source of media used during remote identity proofing [5]. However, federal banking regulation should avoid requiring a particular commercial "deepfake detector." Such technology may rapidly become outdated and may generate false positives or false negatives.

The regulatory obligation should therefore remain technology-neutral: Banks using synthetic-media-sensitive verification should maintain controls reasonably designed to detect or mitigate manipulation and injection risks consistent with recognized technical standards. This formulation would allow defensive technol-

ogy to evolve without requiring constant amendment of the legal rule.

9.4. Integrating NIST Standards into Banking Compliance

NIST standards should not simply be transformed into mandatory banking rules. NIST's Digital Identity Guidelines were developed for a broader digital-identity environment and are not substitutes for banking regulation.

Nevertheless, their underlying principles can inform financial regulation. Three principles are particularly useful: First, identity proofing and authentication are separate assurance problems. Second, the required level of assurance should reflect the potential consequences of failure. Third, technical controls should respond to contemporary threats, including phishing, forged media, and injection attacks.

FinCEN and the federal banking agencies could incorporate these principles by reference to recognized technical standards while allowing institutions flexibility in implementation. Rather than requiring banks to adopt a specific NIST assurance level for every transaction, regulators could require institutions to demonstrate that authentication and identity-assurance controls are proportionate to defined transaction and impersonation risks. NIST would therefore function as a technical benchmark, while banking regulators would determine the legal obligation.

9.5. Proposed FinCEN Guidance or Regulatory Amendment

TLIA could initially be introduced through FinCEN guidance rather than immediate creation of a broad new statutory regime. FinCEN could clarify that where transaction monitoring, authentication systems, or other available information indicates a materially elevated risk of identity impersonation or account takeover, covered financial institutions should consider additional measures to establish the authority of the person initiating or controlling the transaction. Guidance could identify non-exclusive risk indicators and examples of proportionate controls while preserving institutional flexibility.

A possible regulatory principle could state: Where a transaction presents materially elevated indicators of identity impersonation or unauthorized account control, a covered financial institution should apply risk-based measures reasonably designed to obtain additional assurance that the person exercising transactional authority is the verified customer or another legitimately authorized person before executing the transaction, where practicable and appropriate. The important elements are material risk, proportionality, reasonable measures, and legitimate authority.

After regulatory experience develops, policymakers could determine whether more formal amendments to BSA/AML regulations or coordinated interagency guidance are necessary. This gradual approach would allow regulators to evaluate costs, effectiveness, privacy implications, and technological developments before creating prescriptive obligations.

TLIA would therefore clarify the point at which existing identity, authentication, and fraud controls should interact; it would not introduce a new biometric technology or imply that every anomalous payment is unauthorized.

9.6. Conditional Safe Harbor for Banks Following Recognized Standards

A TLIA framework could create uncertainty for financial institutions if banks face liability whenever an advanced impersonation attack succeeds despite reasonable controls. For that reason, this article proposes consideration of a conditional safe harbor. A conditional safe harbor could be considered for institutions that reasonably implement recognized standards and documented risk-based controls, subject to applicable statutory authority and safeguards against material compliance failures.

Any proposed regulatory safe harbor should be conditional rather than absolute. Eligibility could depend on an institution's reasonable implementation of recognized technical standards, maintenance of effective controls, documentation of risk-based decisions, and appropriate responses to material warning signs. A safe harbor should not excuse willful misconduct, gross negligence, or failure to comply with independently applicable legal duties. Moreover, the availability and scope of such protection would require separate legislative or regulatory consideration; compliance with TLIA should not be presented as automatically creating immunity under existing U.S. law.

First, banks would have stronger incentives to invest in modern identity-assurance controls because compliance with recognized standards would provide greater regulatory certainty. Second, regulators could encourage technological improvement without prescribing a single technology that may quickly become obsolete. The safe-harbor concept therefore supports the broader technology-neutral structure of TLIA.

9.7. Proposed TLIA Model

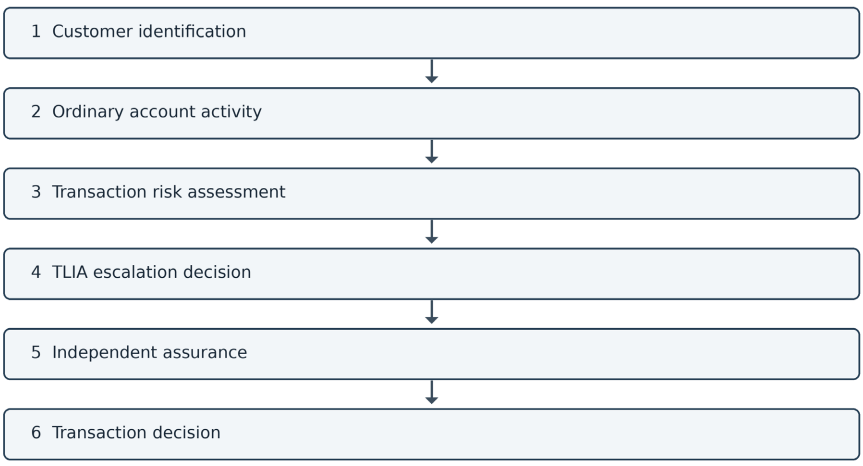

Figure 1. Proposed transaction-level identity assurance (TLIA) framework.

Figure 1 presents the proposed six-stage Transaction-Level Identity Assurance (TLIA) framework. It connects customer identification with ordinary account activity, transaction-specific risk assessment, proportionate escalation, independent assurance, and the ultimate transaction decision. The stages represent a regulatory decision protocol rather than a tested technology or an additional universal KYC obligation. The framework consists of six interconnected stages.

Stage 1—Customer Identification: The process begins with the Customer Identification Program (CIP), through which the bank establishes and verifies the customer's identity in accordance with applicable regulatory requirements. This provides the foundation for the banking relationship.

Stage 2—Normal Account Activity: Once the customer relationship is established, routine transactions proceed through existing authentication, account security, and transaction-monitoring systems. Additional identity verification is not ordinarily required where the available risk indicators do not justify it.

Stage 3—Transaction-Specific Risk Assessment: The bank evaluates relevant indicators, including unusual transaction patterns, unfamiliar devices, irregular account access, changes in customer behavior, and possible attempts at impersonation. The purpose is to determine whether the transaction presents risks that ordinary authentication measures may not adequately address.

Stage 4—TLIA Escalation Trigger: Where the assessment identifies materially elevated impersonation risk, the bank considers whether additional assurance is necessary before proceeding. This decision should be based on the nature and seriousness of the identified threat rather than on unusual transaction activity alone.

Stage 5—Enhanced Identity Assurance: When escalation is justified, the bank applies proportionate and sufficiently independent verification measures. These may include confirmation through a previously verified communication channel or other risk-appropriate procedures designed to establish whether the person exercising transactional authority is genuinely connected to the verified customer or otherwise authorized to act.

Stage 6—Transaction Decision: Following the assurance assessment, the bank determines whether to approve, delay, reject, or refer the transaction for further investigation. The decision should reflect the available evidence, the level of remaining risk, and applicable legal and regulatory obligations.

The six stages describe a proposed analytical sequence, not a validated operational protocol. Future studies should measure fraud prevention, false positives, payment delays, accessibility, and customer impact before claims of effectiveness can be made.

Importantly, TLIA is not intended to replace the existing U.S. banking regulatory framework. CIP would continue to establish customer identity, while Customer Due Diligence (CDD) would support customer-risk assessment and ongoing monitoring. Existing authentication controls would remain responsible for protecting account access, and suspicious activity reporting requirements would continue to govern the reporting of potentially illicit transactions.

The proposed framework instead seeks to improve coordination between these established controls. In particular, it addresses situations where a transaction presents credible indications of impersonation despite the customer's identity having been verified and ordinary authentication procedures having been satisfied. By linking transaction-specific risk indicators with proportionate identity assurance measures, TLIA provides a structured basis for determining when additional confirmation of transactional authority may be appropriate.

In this respect, TLIA proposes a modest but meaningful development in banking identity governance. It extends attention beyond the identity established at the beginning of the customer relationship to the question of who is exercising financial authority when a particular transaction takes place. The approach can be summarized through three principles:

The framework rests on three complementary principles: reliable customer identification at onboarding, continuing assessment of account activity throughout the banking relationship, and proportionate confirmation of transactional authority when credible risk indicators justify additional assurance. These principles are intended to integrate existing controls rather than establish continuous re-identification as a new banking requirement.

The framework also recognizes that stronger security measures may create practical difficulties for legitimate customers. Additional verification can delay urgent payments, disadvantage customers with limited access to digital services, and increase the collection of sensitive personal information. These concerns are particularly relevant where banks rely on biometric authentication or automated fraud-detection technologies.

For this reason, TLIA should be understood not merely as a technical security mechanism but as a broader approach to institutional risk governance. Its implementation would require clear escalation criteria, appropriate human oversight, documented decision-making, and safeguards for privacy and accessibility. Enhanced assurance should therefore remain proportionate to the identified risk and should not impose unnecessary burdens on customers.

Ultimately, the proposed framework seeks to balance fraud prevention with efficient and inclusive banking services. By integrating existing monitoring and authentication controls with targeted verification of transactional authority, TLIA offers a possible approach to addressing impersonation risks while preserving the practical benefits of risk-based banking regulation.

Implementation of TLIA would create practical and legal challenges. These concerns support a proportionate framework rather than continuous or universal enhanced verification.

9.8. Cost and Operational Burden

Enhanced authentication, fraud-detection systems, staff training, and technical upgrades may impose significant costs, particularly on smaller financial institutions. TLIA should therefore apply primarily to transactions presenting materially

elevated impersonation risk rather than routine activity.

9.9. False Positives and Customer Friction

Risk-based systems may incorrectly classify legitimate transactions as suspicious. Excessive intervention could delay payments and inconvenience customers. Banks should therefore use multiple risk indicators and provide appropriate alternative verification methods.

9.10. Privacy and Biometric Risks

Greater reliance on facial, voice, behavioral, or other biometric information could create privacy and cybersecurity concerns. TLIA should follow principles of necessity, proportionality, data minimization, and appropriate data protection rather than encouraging continuous biometric surveillance.

9.11. Technology Obsolescence

Deepfake and fraud techniques will continue to evolve. Regulation should therefore avoid mandating a particular detection product or technology. Legal requirements should remain technology-neutral while recognized technical standards, including relevant NIST guidance, are periodically updated.

9.12. Financial Inclusion

Stronger authentication requirements may disproportionately affect elderly customers, persons with disabilities, customers with limited digital access, and others who cannot easily use advanced authentication technologies. Financial institutions should maintain reasonable alternative verification channels.

These implementation risks support a targeted rather than universal approach. Any proposed escalation duty should be proportionate to the evidence of impersonation and should accommodate accessible alternatives for customers unable to use a particular authentication method.

Four institutional recommendations follow from this analysis. They allocate distinct responsibilities for guidance, supervision, implementation, and evaluation rather than proposing another layer of universal customer re-identification.

9.13. FinCEN

FinCEN should issue guidance recognizing AI-enabled impersonation and account takeover as transaction-stage identity risks and identify circumstances in which elevated impersonation risk may justify enhanced identity or authority assurance.

9.14. Federal Banking Regulators

Federal banking regulators should coordinate with FinCEN and relevant technical bodies to develop technology-neutral supervisory expectations for high-risk transaction authentication. NIST standards should serve as an adaptable technical

benchmark rather than a fixed technological mandate.

9.15. Financial Institutions

Banks should integrate identity risk more closely with existing fraud and transaction-monitoring systems. Where indicators suggest account takeover or AI-assisted impersonation, institutions should use independent, risk-appropriate methods to confirm transactional authority.

Banks that reasonably implement recognized technical standards and documented risk-based controls could also receive a qualified safe harbor, subject to exceptions for material control failures, ignored red flags, gross negligence, or willful misconduct.

The recommended regulatory model applies graduated assurance requirements according to transaction-specific risk. Ordinary transactions would remain subject to existing authentication procedures, while elevated-risk activities could require additional verification. Where credible evidence indicates possible impersonation or unauthorized account control, TLIA would require proportionate independent assurance before the institution determines whether to execute, delay, reject, or investigate the transaction. This approach preserves operational efficiency while strengthening safeguards against AI-enabled financial fraud.

This approach would preserve the existing strengths of CIP, CDD, transaction monitoring, and SAR reporting while creating a clearer connection between verified customer identity and the person exercising authority over a high-risk transaction.

10. Conclusions

Generative AI has not created identity fraud, but it has changed the conditions under which identity can be trusted. Synthetic documents, deepfake video, voice cloning, and AI-assisted impersonation can allow an unauthorized actor to appear convincingly connected to a legitimate customer. As a result, successful customer identification at onboarding does not necessarily establish that the same person is controlling a later transaction.

The analysis answers the principal research question in qualified terms. Existing U.S. banking regulation provides substantial safeguards against identity-related financial crime through CIP, CDD, authentication, fraud controls, and suspicious activity reporting. However, these mechanisms do not establish a single, consistently articulated banking-wide standard specifying when materially elevated transaction-specific impersonation risk should trigger independent assurance of transactional authority before execution. The resulting Identity-to-Transaction Gap is therefore best understood as a problem of regulatory coordination and articulation rather than a complete absence of legal protection. This conclusion supports consideration of risk-based TLIA guidance while acknowledging that its effectiveness, proportionality, and regulatory necessity require further empirical and institutional evaluation.

To address it, the article proposes Transaction-Level Identity Assurance (TLIA). TLIA would not require repeated KYC or continuous biometric verification. Instead, it would introduce a proportionate, risk-based mechanism under which high-risk indicators—such as unusual transaction behavior, account-access anomalies, credential changes, or evidence of AI-assisted impersonation—could trigger enhanced assurance before a consequential transaction is completed.

The proposed framework connects four complementary regulatory functions: establishing the customer's identity, monitoring subsequent account activity, assessing transaction-specific impersonation risks, and obtaining additional assurance of transactional authority when circumstances justify it. By integrating these functions, TLIA seeks to strengthen existing safeguards without replacing customer identification, due diligence, or suspicious activity reporting obligations. NIST's contemporary digital identity standards provide a useful technical foundation for developing such a proportionate, technology-neutral approach.

Future research should test the TLIA model against real-world banking data, including account-takeover cases, false-positive rates, customer impact, implementation costs, and the effectiveness of different authentication methods against AI-enabled impersonation. Further research could also examine the relationship between TLIA and payment liability, biometric privacy law, corporate-account authority, and authorized payment fraud.

Ultimately, the regulatory challenge created by generative AI is not simply whether a bank knows its customer. In an environment where identity can be convincingly imitated, the increasingly important question is whether the bank has sufficient reason to trust the person acting through that identity when significant financial authority is exercised.

The contribution of TLIA is thus a testable proposal for narrowing the distance between verified customer identity and trustworthy transaction-stage authority.

Data Availability

No original participant-level, institutional, or transaction dataset was generated or analysed. This study uses qualitative documentary case analysis of publicly available government and law-enforcement material; no empirical effectiveness estimates are reported.

Declaration of Generative AI and AI-Assisted Technologies in the Manuscript Preparation Process

During the preparation of this manuscript, the author used ChatGPT (OpenAI) for limited assistance with language editing, grammatical refinement, sentence structure, academic writing style, and formatting to improve the clarity, readability, and presentation of the manuscript. The research concept, substantive legal analysis, arguments, conclusions, and intellectual contributions are the author's own. All AI-assisted suggestions were critically reviewed, independently verified where necessary, and revised or approved by the author, who takes full responsi-

bility for the accuracy, originality, and final content of the manuscript.

Conflicts of Interest

The author declares that there are no known competing financial interests or personal relationships that could have appeared to influence the work reported in this article.

References

[1] Financial Crimes Enforcement Network (FinCEN) (2024) Fraud Schemes In-Volving Deepfake Media Targeting Financial Institutions. https://www.fincen.gov/system/files/shared/FinCEN-Alert-DeepFakes-Alert508FI-NAL.pdf

[2] Financial Crimes Enforcement Network (FinCEN) (2024) FinCEN Issues Analysis of Identity-Related Suspicious Activity. https://www.fincen.gov/news/news-releases/fincen-issues-analysis-identity-related-suspicious-activity

[3] Financial Crimes Enforcement Network (FinCEN) (2024) Financial Trend Analysis: Identity-Related Suspicious Activity.

[4] National Institute of Standards and Technology (NIST) (2025) Digital Identity Guidelines. NIST Special Publication 800-63-4.

[5] National Institute of Standards and Technology (NIST) (2025) Digital Identity Guidelines: Identity Proofing and Enrollment. NIST Special Publication 800-63A-4. https://csrc.nist.gov/pubs/sp/800/63/a/4/final

[6] National Institute of Standards and Technology (NIST) (2025) Digital Identity Guidelines: Authentication and Authenticator Management. NIST Special Publication 800-63B-4. https://csrc.nist.gov/pubs/sp/800/63/b/4/final

[7] Code of Federal Regulations (C.F.R.) (2021) Customer Identification Program. Requirements for Banks. 31 C.F.R. § 1020.220. https://www.ecfr.gov/current/title-31/section-1020.220?utm_source=chatgpt.com

[8] Federal Financial Institutions Examination Council (FFIEC) (n.d.) Customer Identification Program—Overview. BSA/AML Examination Manual, Current Online Edition. https://bsaaml.ffiec.gov/manual/AssessingComplianceWithBSARegulatoryRequirements/01?utm_source=chatgpt.com

[9] Hong Kong SAR Government (2024) LCQ9: Combating Frauds Involving Deepfake. https://www.info.gov.hk/gia/general/202406/26/P2024062600192.htm

[10] Federal Bureau of Investigation, Internet Crime Complaint Center (2024) Criminals Use Generative Artificial Intelligence to Facilitate Financial Fraud. PSA I-120324-PSA. https://www.ic3.gov/psa/2024/psa241203

[11] Financial Crimes Enforcement Network (FinCEN) (2024) FinCEN Issues Alert on Fraud Schemes Involving Deepfake Media Targeting Financial Institutions. FIN-2024-Alert004. https://www.fincen.gov/system/files/shared/FinCEN-Alert-DeepFakes-Alert508FINAL.pdf

[12] Code of Federal Regulations (C.F.R.) (n.d.) Risk-Based Ongoing Customer Due Diligence Requirements for Banks. 31 C.F.R. § 1020.210(b)(5). https://www.ecfr.gov/current/title-31/section-1020.210?utm_source=chatgpt.com

[13] Financial Crimes Enforcement Network (FinCEN) (2016) Customer Due Diligence Requirements for Financial Institutions. *Federal Register*, **81**, Article 29398. https://www.federalregister.gov/documents/2016/05/11/2016-10567/customer-due-diligence-requirements-for-financial-institutions

[14] Federal Financial Institutions Examination Council (FFIEC) (n.d.) Customer Due Diligence—Overview. BSA/AML Examination Manual, Current Online Edition. https://bsaaml.ffiec.gov/manual/AssessingComplianceWithBSARegulatoryRequirements/02?utm_source=chatgpt.com

[15] Financial Crimes Enforcement Network (FinCEN) (2026) FinCEN Issues Exceptive Relief to Streamline Customer Due Diligence Requirements. https://www.fincen.gov/news/news-releases/fincen-issues-exceptive-relief-streamline-customer-due-diligence-requirements?utm_source=chatgpt.com

[16] Financial Crimes Enforcement Network (FinCEN) (2026) CDD Rule FAQs. https://www.fincen.gov/resources/statutes-and-regulations/cdd-rule-faqs

[17] Code of Federal Regulations (C.F.R.) (n.d.) Reports by Banks of Suspicious Transactions. 31 C.F.R. § 1020.320. https://www.ecfr.gov/current/title-31/subtitle-B/chapter-X/part-1020/subpart-C/section-1020.320

[18] Simpson, A. (2024) Into the Unknown: The Need to Reframe Risk Analysis. *Journal of Cybersecurity*, **10**, tyae022. https://doi.org/10.1093/cybsec/tyae022

[19] Spencer, M., Coles-Kemp, L. and Hansen, R.R. (2024) Navigating the Landscape of Security Modelling: The MORS Grid. *Journal of Cybersecurity*, **10**, tyae024. https://doi.org/10.1093/cybsec/tyae024

[20] United States Congress (n.d.) Electronic Fund Transfer Act. 15 U.S.C. §§ 1693-1693r; Regulation E, 12 C.F.R. Part 1005. https://www.federalreserve.gov/boarddocs/caletters/2008/0807/08-07_attachment.pdf