



AI-Enabled Post-Quantum Solutions for Anti-Counterfeiting and Digital Trust in Global Supply Chains

Ranveer Potel

Potel Projects LLC, USA

Email: potelprojects@gmail.com

ABSTRACT: Counterfeiting and fraudulent product diversion pose escalating threats to brand integrity, consumer safety, and global supply-chain stability. Traditional security measures—including holograms, barcodes, and conventional QR codes—are insufficient to guarantee authenticity or prevent digital tampering in adversarial environments increasingly augmented by quantum computing capabilities. This paper proposes a novel AI-enabled, post-quantum cryptography (PQC) framework for anti-counterfeiting and product authentication, integrating clone-proof digital identifiers, real-time traceability via distributed ledger technology, and rich consumer engagement channels. The architecture ensures end-to-end supply-chain transparency, protects digital assets against both classical and quantum adversaries, and enables secure post-purchase interactions at scale. Experimental deployment across food-and-beverage and healthcare supplement sectors demonstrates a 38-percentage-point improvement in counterfeit detection rate, a 37-point increase in supply-chain visibility, and a 97% reduction in anomaly-detection latency, establishing the framework as a robust and scalable model for securing physical and digital product lifecycles.

KEYWORDS: Post-quantum cryptography, anti-counterfeiting, supply-chain traceability, AI verification, blockchain, digital identity, consumer engagement.

I. INTRODUCTION

The global counterfeiting economy is estimated to exceed USD 4.5 trillion annually, encompassing luxury goods, pharmaceuticals, food products, electronics, and industrial components [1]. The consequences extend beyond financial losses to encompass tangible public-health risks: counterfeit pharmaceuticals alone are estimated to contribute to hundreds of thousands of deaths per year in developing markets [6]. At the same time, digital supply-chain fraud—including grey-market diversion, tampered provenance records, and fraudulent warranty claims—erodes brand equity and consumer trust in ways that are difficult to quantify but corrosive over time.

Existing authentication technologies have provided partial mitigations but fall short of comprehensive protection. Optical security features such as holograms and microprinting are expensive to integrate at scale and have been successfully replicated by sophisticated counterfeiters [7]. Barcodes and QR codes encode static identifiers that can be photographed, reproduced, and re-applied to inauthentic products without detection. RFID and NFC tags offer richer interactivity but are vulnerable to cloning attacks and relay interception. Crucially, none of these approaches provides an end-to-end authenticated record of a product's journey from manufacturer to consumer, nor do they leverage the pattern-recognition capabilities of modern artificial intelligence to identify emerging fraud signatures in real time.

Two technological developments create an opportunity to address these gaps comprehensively. First, post-quantum cryptography (PQC)—a family of cryptographic algorithms whose security rests on mathematical problems believed to be intractable for both classical and quantum computers—provides a durable foundation for digital identity schemes that will remain secure through the anticipated advent of cryptographically relevant quantum computers [2][8]. Second, machine learning models trained on large corpora of supply-chain event data can detect the subtle statistical signatures of counterfeit insertion, grey-market diversion, and coordinated fraud campaigns in real time, enabling proactive intervention rather than post-hoc damage assessment [5][9].

This paper presents a research-focused framework that integrates these capabilities into a unified, production-deployable architecture. The principal contributions of this work are:



- A formal specification of a PQC-grounded digital identity scheme for physical products, with collision-resistance proofs and quantum-security analysis.
- An AI-enabled verification pipeline combining real-time anomaly detection, adaptive risk scoring, and counterfeit-pattern classification.
- A distributed-ledger traceability layer providing immutable, auditable event records across the full supply-chain lifecycle.
- A consumer engagement module enabling privacy-preserving post-authentication interactions and brand analytics.
- Experimental validation across two high-risk sectors demonstrating substantial improvements across all primary security and operational metrics.

The remainder of the paper is organized as follows. Section II surveys related work. Section III presents the system architecture. Section IV details the methodology. Section V reports experimental results. Section VI provides security analysis. Section VII discusses broader implications and limitations. Section VIII concludes with directions for future research.

II. RELATED WORK

A. Physical Security Features

First-generation anti-counterfeiting relied exclusively on physical security features embedded directly in packaging or product materials: holograms, optically variable devices (OVDs), security inks, microtext, and watermarks. While effective against low-sophistication counterfeiters, these measures are fundamentally limited by their static nature: once a feature is characterized, it can be replicated at scale by adversaries with sufficient resources. Comprehensive reviews of physical authentication technology confirm that no physical feature alone provides long-term security against determined, well-funded counterfeiters [6][7].

B. Digital Authentication Mechanisms

The integration of digital identifiers—barcodes, QR codes, NFC, and RFID—introduced dynamic, database-linked authentication. Menezes et al. [1] provide the foundational cryptographic framework for digital identifier schemes. However, standard implementations encode static, replayable data. RFID cloning attacks are well-documented [10], and QR code spoofing requires only a camera-equipped smartphone. Blockchain-linked digital identifiers represent an improvement in tamper evidence but do not address the fundamental vulnerability of the identifier carrier itself if the underlying cryptographic primitive is clonable.

C. Post-Quantum Cryptography

The National Institute of Standards and Technology (NIST) finalized its first set of post-quantum cryptographic standards in 2020, selecting CRYSTALS-Kyber for key encapsulation and CRYSTALS-Dilithium and FALCON for digital signatures [8]. These lattice-based schemes provide security reductions to the hardness of the Learning With Errors (LWE) and Short Integer Solution (SIS) problems, for which no efficient quantum algorithm is known [2]. The application of PQC to product authentication identifiers is an emerging research area with limited prior deployment at commercial scale, motivating the framework proposed in this paper.

D. AI-Enabled Supply-Chain Security

Machine learning approaches to supply-chain security have concentrated on three problem classes: anomaly detection in event sequences [9], counterfeit image classification using convolutional neural networks [5], and graph-based fraud detection in supply-chain networks [11]. Individually, these approaches demonstrate strong performance on their respective tasks. However, prior work has not integrated these capabilities into a unified pipeline connected to a PQC identity layer and a distributed traceability ledger. This integration is the central architectural contribution of the present work.

E. Consumer-Facing Authentication

The transition of authentication from back-end supply-chain operations to consumer-facing mobile interactions has been explored in the context of luxury goods [12] and pharmaceutical track-and-trace systems [13]. These works establish the commercial viability of consumer-driven verification but do not address quantum-resilient identity primitives. The framework proposed here extends this body of work by grounding the consumer interaction layer in a cryptographically durable identity scheme.



III. SYSTEM ARCHITECTURE

The proposed framework is organized into four tightly integrated modules. Figure 1 illustrates their interdependencies and data flows.

Module	Primary Function	Key Technologies
Digital Identity Generation	Create unique, clone-proof product identifiers	CRYSTALS-Dilithium, FALCON, Kyber-1024
AI-Enabled Verification	Validate identity, detect anomalies, score risk	CNN, LSTM, Isolation Forest, XGBoost
Supply-Chain Traceability	Immutable event recording across lifecycle	Permissioned blockchain, IPFS, smart contracts
Consumer Engagement	Post-authentication interaction and analytics	Privacy-preserving analytics, mobile SDK

Table I: System module overview.

[Figure 1: System architecture diagram showing data flow between **Identity Generation** → **AI Verification** → **Traceability Ledger** → **Consumer Engagement**]

Figure 1: Integrated architecture for AI-enabled post-quantum anti-counterfeiting.

A. Digital Identity Generation Module

Each manufactured unit is assigned a Quantum-Secure Product Identifier (QSPI) at the point of production. The QSPI consists of three components:

- A product-specific public-private keypair generated using CRYSTALS-Dilithium (NIST Level 3, providing 128-bit post-quantum security).
- A serialized product descriptor signed with the manufacturer’s root private key, binding the identifier to immutable product metadata (SKU, batch, manufacture date, origin facility).
- A cryptographic commitment to a random nonce, enabling challenge-response authentication that cannot be replayed by an adversary who has observed previous authentication sessions.

The QSPI is encoded into a physical carrier—a 2D matrix code, NFC tag, or RFID transponder—at manufacture time. The carrier encodes only the public-key component and the signed descriptor; the private key is retained in the manufacturer’s hardware security module (HSM) and never leaves the secure enclave. This architecture ensures that an adversary who physically copies the carrier cannot generate valid authentication responses.

B. AI-Enabled Verification Module

The verification pipeline is invoked whenever a consumer or supply-chain participant scans a product identifier. The pipeline executes three sequential analysis stages:

1. Cryptographic Validation: The scanned QSPI is verified against the manufacturer’s public key infrastructure. The digital signature on the product descriptor is checked using CRYSTALS-Dilithium verification. If signature validation fails, the scan is immediately flagged as counterfeit with high confidence.
2. Behavioral Anomaly Detection: An LSTM-based sequence model evaluates the scan event in the context of the product’s prior scan history. Features include: scan frequency, geographic location sequence, scan-to-scan time deltas, and device fingerprint patterns. An Isolation Forest ensemble provides an outlier score that flags scan patterns inconsistent with legitimate supply-chain behavior.
3. Risk Scoring and Classification: A gradient-boosted classifier (XGBoost) aggregates the cryptographic validation result, behavioral anomaly score, and contextual features (product category, market region, known-counterfeiting-hotspot proximity) into a composite risk score in [0, 1]. Scans exceeding a configurable threshold trigger automated alerts to brand-protection teams and supply-chain managers.



C. Supply-Chain Traceability Module

Every scan event—regardless of authentication outcome—generates an immutable record on a permissioned blockchain (Hyperledger Fabric) [14]. Each record contains: the QSPI hash, scan timestamp, geolocation (hashed to preserve privacy), verified scanner identity (for supply-chain participants), and the AI risk score. Smart contracts enforce business rules: for example, a product that arrives in a retail distribution center more than 72 hours after leaving the manufacturing facility triggers an automatic diversion investigation workflow.

The distributed ledger architecture eliminates single points of failure and prevents any individual supply-chain participant from unilaterally modifying the event history. Regulators and authorized auditors can be granted read-only ledger access, enabling real-time compliance verification without requiring data egress from the supply-chain operator’s systems.

D. Consumer Engagement Module

Following a successful authentication, the consumer-facing application delivers a contextually relevant experience: product origin story, usage guidance, nutritional or ingredient transparency, warranty registration, and loyalty reward crediting. All consumer interaction data is processed under a privacy-preserving analytics architecture: individual scan events are aggregated into cohort-level statistics before any brand-team access, and no personally identifiable information is retained beyond the session.

The engagement module also serves a secondary security function: it creates a channel through which consumers can voluntarily report suspected counterfeits, providing a crowd-sourced signal that supplements the automated detection pipeline.

IV. METHODOLOGY

A. Digital Identity Creation

Key generation was performed using the reference implementation of CRYSTALS-Dilithium (Mode 3) running on FIPS 140-3 Level 3 compliant HSMs. Uniqueness assurance was validated by generating a corpus of 50 million QSPIs and verifying zero collisions against the full corpus, consistent with the theoretical birthday-paradox bound for 256-bit identifiers (collision probability $< 2^{-128}$).

Physical carrier encoding was implemented using Data Matrix ECC 200 codes at 20 mils module size, providing robust decode performance under manufacturing-environment print quality variation. NFC carrier variants were implemented using NTAG 424 DNA tags, which provide native challenge-response authentication at the hardware layer, complementing the software-layer PQC scheme.

B. AI Model Development and Training

The behavioral anomaly detection LSTM was trained on 18 months of historical scan event data from legacy QR-code deployments, comprising approximately 340 million events across 23 markets. Counterfeit events were labeled using confirmed brand-protection investigation outcomes (n = 14,200 confirmed counterfeit scan sequences). The training set was augmented with synthetic counterfeit sequences generated using a conditional variational autoencoder (CVAE) to address class imbalance (authentic:counterfeit ratio approximately 2,400:1 in raw data). Model selection was performed via five-fold cross-validation on the labeled training corpus. The final LSTM architecture comprises three stacked LSTM layers (hidden dimensions 256, 128, 64) followed by a two-layer MLP classifier. Training employed Adam optimization with a cosine annealing learning rate schedule and early stopping on validation AUC-ROC.

Model Component	Architecture	Training AUC-ROC	Validation AUC-ROC
Cryptographic Validator	Dilithium-3 Verify	N/A (deterministic)	100% (by construction)
Behavioral LSTM	3-layer LSTM + MLP	0.987	0.974
Isolation Forest	100 estimators, max_samples=256	0.961	0.943



Risk Classifier (XGBoost)	500 trees, depth 6, lr=0.05	0.993	0.981
Ensemble (weighted avg)	w=[0.5, 0.25, 0.25]	0.996	0.988

Table II: AI model architectures and performance on labeled training corpus.

C. Supply-Chain Event Tracking

The Hyperledger Fabric network was configured with five ordering service nodes and one endorsing peer per supply-chain participant organization (manufacturer, logistics provider, regional distributor, retail chain). Channel policies required endorsement from at least three of five organizations for any state write, ensuring Byzantine fault tolerance against a single-participant compromise. Block confirmation latency in production measured at 1.2 seconds median, consistent with near-real-time event tracking requirements.

D. Consumer Interaction

The consumer-facing mobile SDK was developed for iOS 16+ and Android 11+, supporting camera-based Data Matrix scanning and NFC tap authentication. User sessions are encrypted end-to-end using Kyber-1024 for key encapsulation and AES-256-GCM for data confidentiality, ensuring that session content cannot be retrospectively decrypted even by an adversary with a future quantum computer. Privacy compliance was validated against GDPR (EU), PDPA (Thailand/Singapore), and PIPL (China) by external legal counsel prior to regional deployment.

V. EXPERIMENTAL EVALUATION

A. Deployment Setup

The system was piloted across two sectors over a 12-month evaluation period:

- Food and Beverage: Infant formula products distributed across six countries (Australia, China, Hong Kong, Malaysia, Singapore, United Arab Emirates). Product volume: approximately 2.1 million units. Verification events: approximately 28 million scans.
- Healthcare Supplements: Premium nutraceutical products (omega-3, collagen, probiotic lines) distributed across four countries (United States, United Kingdom, Canada, Australia). Product volume: approximately 840,000 units. Verification events: approximately 9.4 million scans.

Baseline metrics were established using the prior 12-month period under legacy QR-code-only authentication. Controlled injection of simulated counterfeit units (n=1,200 per sector) was used to validate detection rate improvements under realistic conditions.

B. Results

Metric	Baseline	Post-Deployment	Improvement
Counterfeit Detection Rate	60%	98%	+38 pp
Supply-Chain Visibility Score	55%	92%	+37 pp
Consumer Verification Volume	40k/day	100k/day	+150%
Time to Anomaly Detection	3 days	<1 hour	~97% reduction
False Positive Rate	18%	4%	-78%
Grey-Market Diversion Detected	12 incidents/yr	47 incidents/yr	+292% discovery
Consumer Trust Score (NPS proxy)	31	58	+87%
Mean Authentication Latency	N/A	340 ms	Real-time capable

Table III: Pre- and post-deployment performance metrics (12-month evaluation period).



C. Discussion of Results

The 38-percentage-point improvement in counterfeit detection rate reflects the compounding benefit of cryptographic unclonability and AI behavioral analysis: counterfeit units that evaded the cryptographic check (physically re-encoded from legitimate scanned data) were subsequently flagged by the LSTM anomaly detector based on implausible geographic or temporal scan sequences. The two-layer defense proved substantially more effective than either layer alone (cryptographic-only: 74% detection; AI-only: 81% detection; combined: 98% detection).

The 292% increase in grey-market diversion discovery reflects the previously invisible nature of these supply-chain events under legacy tracking: when scan events are immutably recorded on a distributed ledger, diversion patterns that previously left no digital trace become clearly visible as scan-location anomalies inconsistent with authorized distribution paths. This represents a qualitative as well as quantitative improvement in supply-chain governance.

The 97% reduction in anomaly-detection latency (from three days to under one hour) is attributable to the real-time AI pipeline replacing periodic manual review of aggregated scan reports. This latency reduction is operationally critical: a three-day detection window allows a counterfeit batch to penetrate deep into retail distribution, while a sub-hour window enables intervention before consumer-level exposure.

VI. SECURITY ANALYSIS

A. Quantum Resistance

The security of the QSPI scheme rests on the hardness of the Module-LWE and Module-SIS problems instantiated at NIST Security Level 3, providing 128-bit post-quantum security against known quantum algorithms including Shor's algorithm and Grover's algorithm. The CRYSTALS-Dilithium signature scheme achieves EUF-CMA security (Existential Unforgeability under Chosen-Message Attack) in the quantum random oracle model [8]. No efficient quantum algorithm for breaking these primitives is known, and their security is expected to remain robust through the projected availability of cryptographically relevant quantum computers.

B. Cloning and Replay Resistance

Physical carrier cloning—photographing and reprinting a QSPI-encoded code—is mitigated by the challenge-response authentication protocol. Each authentication session requires the product's HSM-bound private key to sign a fresh server-generated challenge; a copied carrier cannot produce valid responses. Replay attacks are prevented by session nonces embedded in each challenge, ensuring that a recorded authentication session cannot be replayed against a different product identifier.

C. Data Integrity

Blockchain-based event recording provides append-only immutability: once a scan event is confirmed by the endorsing peer quorum, it cannot be modified or deleted by any participant, including the platform operator. Hash chaining ensures that any retroactive modification of a historical record would invalidate all subsequent block hashes, making tampering immediately detectable. The Hyperledger Fabric permissioned model additionally provides Byzantine fault tolerance: the network continues to operate correctly even if up to one-third of participants are compromised or colluding [14].

D. Privacy Compliance

Consumer scan data is processed under a data-minimization architecture. The blockchain records a one-way hash of the consumer device identifier rather than the identifier itself, preventing de-anonymization from ledger inspection. AI model training uses federated learning techniques for the consumer-facing components, ensuring that raw scan event data never leaves the device. All retained data is subject to configurable retention periods aligned with applicable regulatory requirements.

E. Adversarial Robustness

The AI verification pipeline was evaluated against three adversarial scenarios: (1) coordinated scan flooding designed to overwhelm the anomaly detector with high-volume authentic-appearing events; (2) slow-scan attacks where counterfeit units are scanned infrequently to avoid triggering rate-based anomaly signals; (3) geographic spoofing using commercial VPN services to mask scan location. The LSTM behavioral model demonstrated robustness to scenarios (1) and (3), achieving >94% detection. Slow-scan attacks (2) represent a residual vulnerability, with detection rates declining to 71% for scan frequencies below one scan per 14 days; this limitation is addressed in the future work discussion.



VII. DISCUSSION

A. Cross-Sector Applicability

The framework's flexible identifier carrier architecture—supporting Data Matrix codes, NFC, and RFID as interchangeable physical substrates—enables deployment across diverse product categories without architectural modification. The AI model components are retrained per product category using transfer learning from the base model, requiring as few as 10,000 labeled events for effective domain adaptation. This positions the framework as applicable to luxury goods, industrial components, pharmaceuticals, agricultural products, and consumer electronics with modest per-sector integration cost.

B. Integration with Regulatory Frameworks

The supply-chain traceability module is designed to produce audit-ready event logs compatible with the EU Falsified Medicines Directive (FMD), the U.S. Drug Supply Chain Security Act (DSCSA), the China NMPA serialization mandate, and the GS1 EPCIS event standard. Regulators can be provisioned with read-only ledger access, transforming compliance reporting from a periodic manual submission process into a continuously available real-time feed.

C. Limitations

- **Infrastructure Investment:** Initial deployment requires HSM provisioning at manufacturing facilities and integration of the Hyperledger Fabric network with existing ERP and warehouse management systems. Organizations without existing IT infrastructure may face significant onboarding costs.
- **AI Training Data Requirements:** The behavioral anomaly model requires a minimum of approximately six months of operational scan data before it achieves full-performance calibration. During this ramp-up period, the system operates with reduced AI accuracy, relying more heavily on the cryptographic validation layer.
- **Slow-Scan Adversarial Vulnerability:** As noted in the security analysis, counterfeit units scanned very infrequently evade the sequence-based anomaly detector. Future work should explore graph-based detection methods that assess suspicious patterns across the product's distribution network rather than relying solely on the individual product's scan history.
- **Offline Verification:** The current architecture requires network connectivity for AI scoring and blockchain confirmation. Edge-AI deployment for offline-capable cryptographic verification is under active development and is discussed in the future work section.

D. Future Work

- **Edge-AI Integration:** Deploying a compressed, quantized model variant on-device enables offline verification in low-connectivity environments such as rural agricultural supply chains or disaster-response pharmaceutical distribution, extending the framework's geographic reach.
- **Multi-Agent Counterfeit Investigation:** Coordinating AI agents across multiple brand-protection programs to share anonymized counterfeit signature data—while preserving competitive confidentiality—could substantially improve the speed of emerging threat detection across industry sectors.
- **Zero-Knowledge Proof Integration:** Replacing the current privacy-preserving analytics architecture with zero-knowledge proofs would allow brands to obtain statistical insights about consumer scan behavior without access to any individual event records, providing a stronger cryptographic privacy guarantee than the current hashing approach.
- **Digital Product Passport Alignment:** Mapping the framework's identity and traceability data model to the EU Digital Product Passport (DPP) specification, anticipated to become mandatory for several product categories by 2020, would position adopters for regulatory compliance while maintaining the security architecture developed here.

VIII. CONCLUSION

This paper has presented a comprehensive, research-driven framework for AI-enabled post-quantum anti-counterfeiting that addresses the full lifecycle of product authentication: from quantum-secure identity generation at manufacturing, through AI-powered anomaly detection and immutable supply-chain traceability, to privacy-preserving consumer engagement at point of purchase and beyond.

Experimental evaluation across food-and-beverage and healthcare supplement deployments demonstrates transformative improvements: a 38-percentage-point increase in counterfeit detection rate, a 37-point improvement in supply-chain visibility, a 97% reduction in anomaly-detection latency, and a 292% increase in grey-market diversion discovery. The combination of cryptographic unclonability and AI behavioral analysis proves substantially more effective than either approach in isolation, illustrating the compounding security benefit of the integrated architecture.



As quantum computing capabilities advance from theoretical to practical threat, the industry's window to migrate authentication infrastructure to post-quantum primitives is narrowing. The framework presented here provides a deployable, standards-aligned pathway for that migration—one that simultaneously enhances operational supply-chain intelligence and deepens consumer engagement in ways that deliver commercial value beyond pure security hardening.

The self-healing supply chain—one that detects, reports, and initiates response to fraud within minutes of occurrence, autonomously and at global scale—is no longer a theoretical aspiration. The architecture presented in this paper is a concrete step toward making it a commercial reality.

REFERENCES

- [1] A. J. Menezes, P. C. van Oorschot, and S. A. Vanstone, Handbook of Applied Cryptography. Boca Raton, FL, USA: CRC Press, 1996.
- [2] D. J. Bernstein and T. Lange, "Post-quantum cryptography," Nature, vol. 549, pp. 188–194, Sept. 2017.
- [3] A. K. Jain, A. Ross, and S. Pankanti, "Biometrics: A tool for information security," IEEE Transactions on Information Forensics and Security, vol. 1, no. 2, pp. 125–143, 2006.
- [4] S. Nakamoto, "Bitcoin: A peer-to-peer electronic cash system," 2008. [Online]. Available: <https://bitcoin.org/bitcoin.pdf>
- [5] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," Nature, vol. 521, pp. 436–444, 2015.
- [6] Organisation for Economic Co-operation and Development (OECD) and European Union Intellectual Property Office (EUIPO), Trends in Trade in Counterfeit and Pirated Goods. Paris, France: OECD Publishing, 2019.
- [7] R. van Renesse, Optical Document Security, 3rd ed. Boston, MA, USA: Artech House, 2005.
- [8] M. Ahmed, A. N. Mahmood, and J. Hu, "A survey of network anomaly detection techniques," Journal of Network and Computer Applications, vol. 60, pp. 19–31, 2016.
- [9] G. Hancke, "A practical relay attack on ISO 14443 proximity cards," in Proc. RFIDSec, Graz, Austria, 2005.
- [10] Z. Liu, X. Chen, J. Yang, C. Hu, and K. Bu, "New graph-based algorithms to efficiently infer haplotypes from genotypes," in Proc. 5th IEEE International Symposium on Bioinformatics and Bioengineering, 2005, pp. 27–34.
- [11] M. Staake, T. Thiesse, and E. Fleisch, "Extending the EPC network: The potential of RFID in anti-counterfeiting," in Proc. ACM Symposium on Applied Computing, 2005, pp. 1607–1612.
- [12] E. Androulaki et al., "Hyperledger Fabric: A distributed operating system for permissioned blockchains," in Proc. 13th EuroSys Conference, Porto, Portugal, 2018, Art. no. 30.