

Article

A Quantum Risk Index for Cryptographic CVEs: Empirical Evidence from the National Vulnerability Database, 2016–2026

Evgeniya Ishchukova ^{1,2,*} , Faezeh Sadat Sajadi ² , Sergei Petrenko ¹ , Alexey Petrenko ¹
and Alexey Nekrasov ² 

¹ Scientific Center for Information Technologies and Artificial Intelligence, Sirius University of Science and Technology, Sirius Federal Territory, Sochi 354340, Russia; petrenko.sa@talantiuspeh.ru (S.P.); a.petrenko1999@rambler.ru (A.P.)

² Institute for Computer Technologies and Information Security, Southern Federal University, Chekhova 2, Taganrog 347922, Russia; a2r.sahar@gmail.com (F.S.S.); alexei-nekrassov@mail.ru (A.N.)

* Correspondence: jekky82@mail.ru; Tel.: +7-(928)-1435898

Abstract

The harvest now, decrypt later (HNDL) attack is an attack that collects encrypted information now and decrypts it later after the arrival of a quantum computer that can perform some cryptographic operations. Time of exposure is data retention and not Q-Day, so the threat is imminent, but it is not reflected in the Common Vulnerability Scoring System (CVSS). We present the Quantum Risk Index (QRI), which is based on CVSS base severity, the Quantum Factor (QF, vulnerability to Shor's or Grover's algorithm), and the HNDL Score (HS, susceptibility to a harvest-and-store adversary). We computed the QRI for 78,587 CVEs that were cataloged in the NIST National Vulnerability Database between January 2016 and the first quarter of 2026 and cross-checked it with the CISA Known Exploited Vulnerabilities (KEV) list. The number of CVEs related to cryptography increased by a CAGR of 11.0%, while the number of Shor-vulnerable CVEs increased at a CAGR of 8.8%. The 437 KEV-matched CVEs carry a mean QRI of 13.05, against 10.31 for the non-KEV remainder—a 26.6% separation ($p = 2.17 \times 10^{-63}$)—and Shor-vulnerable CVEs appear in the KEV list at 1.52 times the baseline rate ($p = 0.002$). This separation is valid for each of the tested settings of QF and HS. This is the first study to apply a quantum-adjusted vulnerability score to a government's in-the-wild vulnerability reporting database at the CVE scale, providing a repeatable foundation for quantum-aware vulnerability triage.

Keywords: harvest now, decrypt later; quantum risk; vulnerability scoring; CVSS; post-quantum cryptography; CISA KEV; National Vulnerability Database; CVE analysis



Academic Editor: Leandros Maglaras

Received: 10 June 2026

Revised: 31 July 2026

Accepted: 31 July 2026

Published: 2 August 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and

conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

1. Introduction

Our assumption of the security of every transport-layer session [1], every signed software update, and every blockchain settlement in flight today is based on two computational assumptions: the difficulty of integer factorization [2] and the difficulty of the discrete logarithm [3]. Shor's algorithm breaks these in polynomial time on a powerful enough quantum computer [4]. For three decades, it was a future problem, as the design of such a computer seemed to be decades off. That assessment has changed. Gidney's 2025 estimate of the requirement to factor RSA-2048 is less than one million noisy physical qubits—a factor of ten smaller than the estimate five years ago [5,6].

The HNDL strategy eliminates any remaining timing margin. If an adversary collects and stores a ciphertext today, they suffer no time penalty; it can be read at any time a

cryptographically relevant quantum computer (CRQC) becomes available [1]. Thus, the sensitive duration is not Q-Day but the lifetime of the data. Governments and intelligence reports warn of HNDL collection campaigns targeting high-value targets [7], while the U.S. Federal Reserve has issued warnings that distributed ledgers represent a sector where collected data may remain valuable for decades [3].

Policy has responded. In August 2024, the NIST finalized three post-quantum cryptography standards—FIPS 203, 204, and 205—based on lattice and hash-based cryptosystems that are resistant to Shor’s and Grover’s algorithms [8–11]. The NSA’s CNSA 2.0 requires national security systems to adopt PQC by 2035, with targets starting in 2027. NIST IR 8547 foresees a five- to 10-year turnaround for critical infrastructure. Nevertheless, uptake in critical sectors is sparse, and PQC in resource-constrained settings encounters hardware and protocol limitations not reflected in classical scoring [12,13].

What the policy momentum cannot tell us is how large the exposed surface is or whether attackers are already distinguishing cryptographically vulnerable CVEs from other CVEs when targeting. The Exploit Prediction Scoring System (EPSS) focuses on near-term exploitation but is not quantum-sensitive [14]. The QARS model applies Mosca’s inequality to a multi-dimensional formula but is not empirically validated [15]. Neither of these has been applied to question whether the population of cryptographic CVEs is increasing, whether the most quantum-risky CVEs are the ones being exploited, or whether HNDL-attractive CVEs are over-represented among those proven to have been exploited in practice [16].

This paper addresses the question empirically. We collect all cryptographically relevant CVEs from NVD between 2016 and the first quarter of 2026, calculate a composite Quantum Risk Index (QRI) that scales the CVSS by a Quantum Factor (QF) and the HNDL Score (HS), and crossmatch that scored corpus against the CISA Known Exploited Vulnerabilities catalog. Four a priori hypotheses are tested: (i) whether the volume of crypto-CVEs is increasing, (ii) whether CVEs listed in the KEV list are significantly more risky than other CVEs according to the QRI, (iii) whether CVEs vulnerable to Shor’s algorithm are over-represented among KEV CVEs, and (iv) whether the QRI risk is concentrated in a Pareto-like minority of records.

The rest of this paper is structured as follows. Section 2 reviews related work. Section 3 describes the data, scoring formulas, and hypothesis tests. Section 4 presents the four hypothesis tests and the sensitivity analysis. Section 5 provides the implications for vulnerability management for an HNDL threat model, as well as limitations. Section 6 concludes this paper.

2. Related Work

2.1. The Harvest-Now-Decrypt-Later Threat

The HNDL threat is premised on one idea: the security of public-key cryptography is contingent on the absence of a quantum adversary. Bernstein and Lange make this clear—when powerful enough quantum computers are available, the public-key cryptography that secures the Internet, as well as financial and government communications, breaks down [17]. Shor’s algorithm makes this possible; Grover’s algorithm doubles down on symmetric primitives by cutting their effective key size in half [18]. Thus, any information encrypted using today’s protocols is already “broken” in a retrospective sense, if not immediately decrypted. What is new is the evidence that adversaries are not sitting still. Kagai et al. formalized the threat in the telecommunication sector and demonstrated that high-retention industries are exposed to such threats for decades due to the delayed post-quantum cryptography (PQC) transition [1]. Mascelli and Rodden applies this to distributed ledger networks that are regulated by the Federal Reserve [16]. Gidney’s

updated qubit estimate for RSA-2048 reduces the time frame significantly [5]. This body of work presents a qualitative picture of rising concern. These studies, however, have not quantified the surface in terms of disclosed vulnerabilities.

2.2. Quantum Risk Assessment Frameworks

Grigaliūnas and Brūzgienė's QARS model extends Mosca's inequality to a multi-parameter scoring formula that is calibrated by expert opinion; this is the most similar conceptual antecedent of the QRI but is not validated by behavioral data [15].

There are several previous works that modify STRIDE-based threat modeling for blockchain systems and cloud systems. Hasan et al. suggested a security dependency analysis post-quantum cryptography (PQC) migration process [19]. Aggarwal et al. described one of the first systematic investigations of quantum attacks on cryptographic schemes in use in the real world—in this case, elliptic-curve signatures in Bitcoin [20]. Meanwhile, Taherdoost's latest survey explores quantum-resistant blockchain systems and assesses their maturity for adoption [2,21].

These approaches do have one problem, though—they are more abstract and not measured against real-life attacker activity.

2.3. Empirical Vulnerability Scoring and Prioritization

The most widely used prioritization tool is the CVSS base score, but FIRST describes it as a severity score rather than a risk score [22], and Spring et al. showed its deficiencies in risk-based triage [23]. Wunder et al. demonstrated scoring inconsistencies that compromise its further use [24,25]. Jacobs et al.'s EPSS is the current best practice for immediate exploitation probability, with an ROC-AUC of 0.838 [14]. Two insights emerge from this work: the CVSS alone is not suitable for risk-aware prioritization, and composite empirical measures are more effective than single measures.

3. Materials and Methods

Here, we describe the process of data collection, scoring, and verification. The study is observational and longitudinal: all CVE records published in NVD from 2016 to the first quarter of 2026 were acquired, filtered for cryptographic relevance, scored along three axes (CVSS, QF, and HS), and compared to both the CISA KEV list and the EPSS exploitation-probability feed. Figure 1 outlines the entire process.

3.1. Data Sources

3.1.1. NIST National Vulnerability Database

We obtained CVE data by interacting with the official NVD CVE API 2.0 [26], which provides structured JSON feeds by year. The records include the CVE ID, date of publication, description (in English), CVSS metrics (versions 2.0, 3.0, and 4.0), and CWE tags, as well as a vulnerability status indicator. Our retrieval period of 2016–2026 produced 260,253 records.

These records constitute the INPUT and FILTER phases shown in Figure 1.

3.1.2. CISA Known Exploited Vulnerabilities Catalog

The CISA KEV catalog is the official U.S. government list of CVEs that have been confirmed as exploited in the wild [27]. On the date of access, the catalog had 1558 items. KEV is the independent behavioral ground-truth source for the validation stage. If the QRI captures HNDL-relevant risk, KEV-listed CVEs should, on average, score higher than the rest of the cryptographic CVE population—the case-control setting of Allodi and Massacci [28].




Figure 1. QRI computation workflow. The pipeline proceeds in four numbered phases: (1) INPUT—three reference data streams (NVD CVE API 2.0, CISA KEV catalog, and EPSS bulk feed) are retrieved; (2) FILTER—the raw NVD corpus is reduced to the keyword-filtered crypto-relevant corpus of 78,587 CVEs; (3) SCORE—QF, CVSS base severity, and HS are computed per record and combined into the composite QRI score; and (4) OUTPUT—the scored dataset is emitted, with CISA KEV and EPSS joined to the final records by CVE identifier for independent validation.

3.1.3. Exploit Prediction Scoring System (EPSS)

The Exploit Prediction Scoring System (EPSS) [14] is a machine-learned probability of near-term exploitation, published daily by the FIRST.org Special Interest Group. The EPSS is included as a third reference data stream alongside NVD and CISA KEV. We obtained EPSS scores from the public bulk feed at first.org/epss and joined them to our corpus by CVE identifier; coverage was 78,150 of 78,587 records (99.4%). The EPSS serves two roles in the analysis. First, it acts as an independent predictive baseline against which the QRI is benchmarked (Section 4.5.3.). Second, the empirical EPSS performance on our dataset informs the positioning of the QRI relative to existing prioritization tools (Section 5.3).

3.2. Crypto-Relevance Filtering

The dataset was filtered to include records whose English description included at least one of the keywords in a three-tiered lexicon: (i) Shor-relevant primitives (RSA, ECDSA, Diffie-Hellman, X.509, PKI, and their synonyms), (ii) Grover-relevant primitives (SHA-256, SHA-1, AES-128, HMAC, and hash function), and (iii) generic cryptographic context terms (TLS, SSL, cipher, encrypt, certificate, and key exchange).

3.3. Quantum Factor (QF) Assignment

The QF quantifies the quantum vulnerability of cryptographic vulnerabilities. The assignment is based on deterministic keyword matching on the CVE description (Table 1). The justification is simple: Shor’s algorithm completely compromises asymmetric schemes [4], whereas Grover’s algorithm reduces the effective bit security of symmetric schemes by half [18]. Entries matching generic cryptographic keywords but not identifying the primi-

tive type are given QF = 1.1, accounting for any remaining quantum vulnerability where the primitive cannot be clearly identified.

Table 1. Quantum Factor (QF) assignment rules and dataset distribution.

QF	Representative Keywords	Quantum Threat Basis	<i>n</i> CVEs (%)
1.5	RSA, ECDSA, DSA, DH, ECDH, Ed25519, and X.509/PKI	Shor’s algorithm—complete break	8092 (10.3%)
1.2	SHA-1/256/512, MD5, HMAC, AES-128, and hash function	Grover’s algorithm—halved key length	498 (0.6%)
1.1	TLS, SSL, OpenSSL, cipher, encrypt, and certificate	Generic crypto—algorithm undetermined	6233 (7.9%)
1.0	No cryptographic keyword match	No quantum exposure	63,764 (81.1%)

Convergent Validity of the Keyword Classifier

Because the keyword classifier is the gateway to all downstream scoring, we assessed its convergent validity against an independent annotation signal: the National Vulnerability Database’s CWE taxonomy. We treated a curated list of cryptography-related CWE identifiers as a reference label—CWE-310, 311, 312, 321–340, and 345–347 (cryptographic primitives, key management, randomness, and signature verification), as well as CWE-295–299 (improper certificate validation). The CWE list is not a perfect ground truth: NVD CWE tagging is well known to be sparse, with many genuine cryptographic CVEs carrying no CWE annotation at all. Concordance metrics against this proxy reference therefore characterize overlap, not absolute accuracy.

Two analyses were performed. The first one computed corpus-wide concordance across all 78,587 filtered records. The second one is a stratified random sample of 200 CVEs (50 per QF tier, seed = 42) that recomputed the same metrics to confirm that corpus-level statistics generalize. Corpus-wide, 4571 of 78,587 keyword-positive CVEs (5.8%) carry a cryptography-related CWE tag. Concordance is strongly stratified by QF tier: 39.9% at QF = 1.1 (generic-cryptographic matches) and 33.1% at QF = 1.2 (Grover-relevant primitives), against 2.1% at QF = 1.0. The QF = 1.5 tier (Shor-relevant primitives) shows lower CWE concordance (6.9%) because Shor-affected CVEs are frequently tagged under product-implementation CWEs (e.g., CWE-22 path traversal, CWE-203 timing) rather than cryptographic CWEs, even when the affected primitive is RSA, ECDSA, or Diffie–Hellman. The stratified sample reproduces this pattern.

Results for the crypto-CWE list on the fixed *n* = 200 stratified sample (50 per Quantum Factor (QF) tier, archived as Supplementary File S1 to guarantee exact reproducibility independent of software version) yield the following scores: precision = 0.247, recall = 0.949, *F1* = 0.392, and Cohen’s κ = 0.119. The low precision and low κ are consistent with the well-documented incompleteness of NVD CWE coverage of cryptographic vulnerabilities and not a reflection of error by the classifier. Recall > 0.94 is more informative—the keyword classifier recovers CWE-confirmed cryptographic CVEs, as well as crypto-relevant CVEs not covered by the CWE taxonomy.

The dominant discrepancy in this validation is false positives (CVEs flagged by the keyword filter but not confirmed by the CWE proxy) rather than false negatives (CVEs confirmed by the CWE proxy but missed by the keyword filter)—consistent with the high recall (0.949) and low precision (0.247) reported above. As argued, this pattern reflects

CWE-label sparsity rather than genuine over-inclusion, which bounds its practical impact on the QRI distributions reported in Section 4.

Two practical points follow. Taken at face value, the precision figure would imply that part of the keyword-filtered corpus is not strictly cryptographic in the protocol-level sense. The empirical findings in Section 4 are nevertheless unaffected because KEV labels are produced by CISA independently of our keyword classifier; the separation between KEV-confirmed and non-KEV CVEs in the QRI (Section 4.2) is therefore a property of records that exist in CISA’s catalog and not of how we filtered the surrounding population. Any noise introduced at the keyword stage would attenuate the observed effect sizes, not inflate them, and the Mann–Whitney separation remains significant at $p = 2.17 \times 10^{-63}$ (where p denotes the probability of obtaining the observed separation or one more extreme under the null hypothesis of no difference) despite this attenuation—evidence that the signal is robust. Section 5.5 returns to this point.

3.4. HNDL Score (HS) Computation

The HS reflects the appeal of a vulnerability to an adversary following the harvest-and-store model. It is the sum of four independent criteria (Table 2) identified in the CVE description, capped at 6. We present the frequencies of each level in the data in Table 3.

Table 2. HNDL Score (HS) additive scoring rubric.

Pts	Criterion	Representative Keywords	HNDL Rationale
+2	Long-term data storage	Database, archive, backup, healthcare, finance, blockchain, and government	Data available for future harvest
+2	Interceptable channel	TLS, SSL, HTTPS, SSH, VPN, IPSec, LDAP, SMTPS, key exchange	Traffic capturable in the network layer
+1	High-value target sector	blockchain, financial, bank, healthcare, government, military, and SCADA	Elevated attacker motivation
+1	Key material exposure	Private key, session key, master key, key disclosure, and key exposure	Compound risk from key compromise

Table 3. HNDL score level interpretations and frequency in the study dataset ($n = 78,587$).

HS	Interpretation	Typical CVE Profile	Frequency
0	No HNDL exposure	No storage, channel, or sector match	9.1%
1	Minor channel exposure	Weak channel match only	0.7%
2	Network channel exposure	Interceptable protocol identified	85.2%
3	Channel + sector context	Protocol and high-value domain	1.5%
4	High-value target with channel	Finance/health/government + protocol	3.4%
5	Critical: channel + storage + key	Full HNDL profile—maximal exposure	0.2%

The four HNDL component categories were chosen to align with the prioritization principles outlined in NIST IR 8547 [29], which recommends prioritizing systems that process long-lived sensitive data and high-value transactional or sector-critical assets. NIST

IR 8547, however, does not specify numerical weights; the point values used here (+2 for storage/persistence and communication-channel characteristics and +1 for high-value sector and key-exposure characteristics) reflect the authors' own calibration of which characteristics most strongly determine HNDL exposure.

To assess sensitivity to this calibration, five alternative weighting schemes—equal, storage-emphasized, channel-emphasized, sector/key-emphasized, and storage + channel-only weighting—were applied to the identical crypto-relevant corpus ($n = 78,587$; 437 CISA KEV) alongside the baseline, and KEV/non-KEV separation was recomputed for each. Across all six schemes (baseline included), mean separation ranged from 23.8% to 30.6% (baseline: 26.6%), and ROC-AUC ranged from 0.700 to 0.736 (baseline: 0.732), indicating that the paper's central empirical results are not an artifact of this specific weighting scheme.

3.5. Quantum Risk Index (QRI) Formula

The quantum risk index is expressed as follows:

$$\text{QRI} = \text{CVSS}_{\text{base}} \times \text{QF} \times (1 + \text{HS}^2/10).$$

The CVSS delivers the classical anchor. QF provides a multiplicative quantum-susceptibility modifier. The HS term is quadratic so that vulnerabilities with multiple risk factors (e.g., a TLS protocol vulnerability in a backup system for a medical institution with exposed key material) will be adjusted by more than a vulnerability with a single risk factor. We chose the quadratic form over the logarithmic form because it has better KEV separation in the sensitivity analysis (Section 4.5).

The $\text{CVSS}_{\text{base}}$ and QF, together, constitute a conjunctive severity gate: a CVE can only contribute to the QRI if it has both non-negligible conventional severity and non-zero quantum relevance. The $(1 + \text{HS}^2/10)$ term then becomes a continuous exposure-amplification factor, raising this base severity as harvest-now-decrypt-later exposure grows—analogueous to how the CVSS, itself, incorporates its temporal and environmental metrics to modulate but not gate the base score.

The composition of the scored corpus by QF tier and by HS level is summarized in Tables 1 and 3, respectively.

3.6. Dataset Composition

Table 4 reports the annual evolution of the crypto-relevant CVE population alongside the Shor-vulnerable subset and the CISA KEV match counts. The dataset with a score is formed by 78,587 records over 11 years. The fraction of crypto-relevant CVEs in the yearly dataset increased from 24.4% in 2016 to 35.1% in the first quarter of 2026 (only the first three months were available). Of the 78,587 records, 8092 (10.3%) carry $\text{QF} = 1.5$, and 2777 (3.5%) carry $\text{HS} \geq 4$. Matching CISA KEV reveals 437 (0.56%) records have been exploited. See Table 4 for an annual breakdown.

3.7. Statistical Approach and Hypothesis Framework

Four hypotheses were pre-specified:

- Q1. The number of crypto-relevant CVEs and Shor-vulnerable CVEs increased significantly between 2016 and 2025.
- Q2. CVEs listed in KEV have higher QRI scores than those not listed in KEV (behavioral validation).
- Q3. Shor-vulnerable CVEs are over-represented in the KEV list compared to the population of all CVEs.
- Q4. QRI risk is highly concentrated, as with the heavy tail of exploitability in previous studies [28,30].

Q1 was tested by ordinary least-squares regression of annual counts (log-transformed) on a linear time trend (significance threshold of $\alpha = 0.05$; CAGR reported).

Table 4. Annual dataset composition.

Year	<i>n</i> Crypto CVEs (%)	Mean QRI	Shor-Vuln. (QF = 1.5)	HS ≥ 4	CISA KEV
2016	2593 (24.4%)	10.369	310 (11.9%)	96	14
2017	4231 (24.8%)	10.937	595 (14.1%)	226	30
2018	5015 (28.4%)	10.821	580 (11.6%)	154	21
2019	5036 (28.7%)	10.626	513 (10.2%)	156	34
2020	5978 (28.4%)	10.497	624 (10.4%)	201	46
2021	6818 (29.2%)	10.561	657 (9.6%)	403	54
2022	7713 (28.1%)	10.464	808 (10.5%)	287	42
2023	9833 (31.6%)	10.465	884 (9.0%)	319	58
2024	12,835 (32.9%)	10.042	1187 (9.3%)	368	60
2025	14,523 (33.1%)	9.730	1454 (10.0%)	393	63
2026 *	4012 (35.1%)	10.443	480 (12.0%)	174	15
Total	78,587 (30.2%)	9.707	8092 (10.3%)	2777	437

* The 2026 row reflects the partial first quarter and is retained in the regression series as a single additional data point. KEV = CISA Known Exploited Vulnerabilities.

The partial 2026 cohort, which includes only Q1 data, is retained in this regression and in the corresponding sensitivity analyses (Section 3.8) as a single additional data point rather than being excluded or annualized; Table 4 reflects this treatment. Q2 was tested by the Mann–Whitney *U* rank-sum statistic with a rank-biserial effect size. Q4 was tested by the Gini coefficient (a measure of distributional inequality ranging from 0 for perfect equality to 1 for maximal concentration), with a pre-specified confirmation threshold of Gini > 0.5 (considered the conventional threshold for inequality). The analyses were performed in Python 3.14 (NumPy 2.4, SciPy 1.16, pandas 3.0, scikit-learn 1.8, Matplotlib 3.10) (SciPy). Data and code are publicly available (see Data Availability Statement).

3.8. Pre-Specified Sensitivity Analysis

A pre-specified sensitivity analysis tested variations of QF multipliers (Shor: 1.3–2.0; Grover: 1.1–1.5) and the HS functional form (linear, logarithmic, cubic, binary, ablated) to see if the Q2 finding was robust. We also compared the QRI vs. CVSS by ROC-AUC to determine if the quantum-adjustment factors add incremental value.

4. Results

This section presents the results for the four tests of hypotheses and the sensitivity analysis.

4.1. Q1—Growth of the HNDL Attack Surface (2016–2026)

The number of cryptographically relevant CVEs grew from 2593 in 2016 to 14,523 in 2025—a 5.6-fold increase over the span of nine years. OLS log-count-on-time regression confirms the trend (fitted slope of the log-count trend of $\beta_1 = 0.105$; coefficient of determination of $R^2 = 0.447$; $p = 0.025$)—a compound annual growth rate of 11.0%. Shor-vulnerable CVEs (QF = 1.5) grew from 310 in 2016 to 1454 in 2025 (CAGR = 8.8%, $\beta_1 = 0.084$, $R^2 = 0.420$, and $p = 0.031$). The average QRI dropped slightly from 10.37 to 9.73 ($\beta_1 = -0.065$, $p = 0.032$)—supportive of an increased proportion of generic cryptographic CVEs in recent years driving the average down without contradicting the growth claim.

The Shor-vulnerable subset grows more slowly than the overall population (CAGR of 8.8% vs. 11.0%), consistent with Shor-relevant primitives (RSA and ECDSA) representing

a shrinking share of newly disclosed cryptographic weaknesses as legacy protocols are phased out, even as absolute disclosure volume rises.

Thus, Q1 is confirmed: both the overall crypto-relevant CVE population and the Shor-vulnerable subset grew significantly across the 2016–2026 window (the partial 2026 cohort is retained in the regression as a single additional data point; see Section 3.7).

4.2. Q2 Behavioral Validation of the QRI Against CISA KEV

The 437 CVEs confirmed as exploited in the wild carry a mean QRI of 13.05 (median: 12.74; SD: 4.17), against a mean of 10.31 (median: 10.01; SD: 3.56) for the 78,150 non-KEV records—a separation of 26.6% in means (Table 5) (Mann–Whitney $U = 24,996,992$; $p = 2.17 \times 10^{-63}$; rank-biserial correlation, i.e., an effect-size measure bounded between 0 and 1 in absolute value; $|r| = 0.46$).

Table 5. QRI descriptive statistics by CISA KEV status.

Group	<i>n</i>	Mean QRI	Median QRI	SD
CISA KEV (exploited)	437	13.048	12.740	4.169
Non-KEV (general population)	78,150	10.310	10.010	3.555
Difference (KEV – Non-KEV)	–	+2.738	+2.730	–
% Difference	–	+26.6%	+27.3%	–

Figure 2 visualizes the resulting distributional separation between the two groups.




Figure 2. QRI distribution for CISA KEV and non-KEV CVEs (Q2). Violin plots with embedded box plots; the diamond marker indicates the mean. Non-KEV CVEs: $n = 78,150$; CISA KEV CVEs: $n = 437$.

A percentile analysis (Table 6) reveals that 375 of the 437 KEV CVEs (85.8%) exceed the 50th QRI percentile ($1.7 \times$ the 50% null expectation), and 202 KEV CVEs (46.2%) are above the 90th percentile (10% null). The over-representation across the four percentiles is robust and strong.

Table 6. Percentile analysis: proportion of CISA KEV CVEs above successive QRI thresholds.

QRI Percentile Threshold	KEV CVEs Above	% of KEV	Random-Baseline Expectation
50th (QRI ≥ 10.1)	375/437	85.8%	~50% under null
75th (QRI ≥ 12.3)	250/437	57.2%	~25% under null
90th (QRI ≥ 13.7)	202/437	46.2%	~10% under null
95th (QRI ≥ 15.4)	59/437	13.5%	~5% under null

Figure 3 plots the proportion of KEV CVEs above each QRI percentile threshold against the random baseline expectation.




Figure 3. Percentile analysis of CISA KEV CVEs against QRI thresholds (Q2). Orange bars show the observed proportion of KEV CVEs exceeding each QRI percentile; gray bars show the null-model expectation under random assignment.

Thus, Q2 is confirmed: CISA KEV CVEs carry significantly higher QRI scores than non-KEV CVEs, and the separation persists across all four percentile thresholds.

4.3. Q3—Shor-Vulnerable CVEs and Current Exploitation Behavior

The KEV rate across all crypto-relevant CVEs is 0.56% (437 of 78,587). For the Shor-vulnerable subset ($n = 8092$), the KEV rate is 0.80% (65 records), a ratio of $1.52 \times$ the baseline. Chi-square confirms the over-representation: $\chi^2(1) = 9.476$, and $p = 0.002$, where χ^2 is the chi-square test statistic and the value in parentheses is its degrees of freedom. The HNDL score is independently further validated by a significant positive association between HS tier and exploit rate (HS = 0: 0.69%; HS = 1–3: 0.52%; HS ≥ 4 : 1.22%), with $\chi^2(2) = 26.622$ and $p < 0.001$.

The apparent dip at HS = 1–3 relative to HS = 0 (Figure 4b) likely reflects compositional differences rather than a genuine non-monotonic risk relationship: HS = 0 CVEs indicates that none of the four HNDL categories is matched, including many high-severity vulnerabilities in widely deployed general-purpose software whose KEV inclusion is driven by broad exploitability and deployment footprint rather than HNDL-specific attractiveness. HS = 1–3 CVEs, matching only one or two categories, may under-represent this generic

high-severity class while not yet capturing the concentrated high-value targets that dominate at $HS \geq 4$. The chi-square test nonetheless confirms a significant overall association between the HS tier and exploitation, driven primarily by the pronounced elevation at $HS \geq 4$.




Figure 4. KEV exploitation rates by quantum vulnerability tier (Q3): (a) Shor-vulnerable CVEs (QF = 1.5) versus all other crypto CVEs; (b) KEV rate by HNDL score tier (sample sizes shown for each bar). The fitted trend line is illustrative; the non-monotonic pattern between HS = 0 and HS = 2 is discussed in Section 4.3. The green arrow in panel (b) marks the overall trend across HNDL Score tiers and carries no statistical meaning of its own.

Figure 4 visualizes the elevated KEV exploitation rates for Shor-vulnerable CVEs and across HS tiers.

Thus, Q3 is confirmed: Shor-vulnerable CVEs are significantly over-represented in the CISA KEV catalog, and the HNDL score independently tracks observed exploitation rates.

We note that this association concerns current exploitation behavior under classical computing and does not presuppose the timeline or eventual feasibility of a cryptographically relevant quantum computer, which remains under active technical debate. Thus, the prioritization suggested by the QRI should be read as evidence-based triage guidance under current threat conditions, not as a claim about quantum-computing timelines.

4.4. Q4—Concentration of QRI Risk

The QRI distribution has *Gini* coefficient of 0.176—much lower than the preset threshold of $Gini > 0.5$. According to the Pareto analysis, the top 70% of CVEs cover 80% of the total aggregated QRI—the reverse of the usual Pareto situation. This negative finding is informative. The Gini-and-Pareto analysis sets up the QRI to inherit the skewness of the Allodi distribution [28]. The evidence is that the QRI does not, because the population has been filtered for cryptography: the assumptions (CVSS, QF, and HS) for this population do not cover the range that would lead to high inequality. The concentration is not in the distribution of the QRI itself but in the upper tail where exploited CVEs are concentrated, as we already found under Q2 (46.2% of KEV CVEs are above the 90th percentile, compared to 10%). We discuss this distinction again in Section 5.4.

Thus, Q4 is not confirmed under the pre-specified Gini > 0.5 threshold; instead, the operationally meaningful concentration lies in the joint QRI-KEV distribution, as established under Q2 (Section 4.2).

4.5. Sensitivity Analysis

4.5.1. Robustness of Q2 to QF Parameterization

Table 7 reports Q2 under five QF scenarios. Separations between KEV and non-KEV CVEs range from 24.6% (QF removed) to 28.3% (maximum multipliers); all scenarios have $p < 10^{-54}$. The quantum-susceptibility adjustment boosts separation but is not the driver.

Table 7. Sensitivity of Q2 to Quantum Factor (QF) parameterization.

QF Scenario	KEV Mean	Non-KEV Mean	Diff %	<i>p</i> -Value
Baseline (Shor = 1.5, Grover = 1.2)	13.048	10.310	+26.6%	2.14×10^{-63}
Scenario A (Shor = 1.3, Grover = 1.1)	12.691	10.088	+25.8%	3.50×10^{-62}
Scenario B (Shor = 1.7, Grover = 1.3)	13.406	10.532	+27.3%	8.58×10^{-62}
Scenario C (Shor = 2.0, Grover = 1.5)	13.922	10.848	+28.3%	8.66×10^{-59}
Scenario D (no QF—all equal 1.0)	12.184	9.777	+24.6%	6.37×10^{-54}

Note: All scenarios significant at $p < 10^{-50}$.

4.5.2. Robustness of Q2 to HS Formula Choice

Table 8 gives Q2 with six different functional forms for HS. Separation ranges from 22.4% (binary) to 27.2% (cubic), all with $p < 10^{-53}$. The quadratic baseline offers the best trade-off between separation and interpretability, but the binary, linear, log, and ablated forms all achieve high separation.

Table 8. Sensitivity of Q2 to HS formula choice.

HS Formula	KEV Mean	Non-KEV Mean	Diff %	<i>p</i> -Value
Quadratic—baseline: $1 + \text{HS}^2/10$	13.048	10.310	+26.6%	2.14×10^{-63}
Linear: $1 + \text{HS}/2$	17.589	14.119	+24.6%	1.33×10^{-53}
Cubic: $1 + \text{HS}^3/50$	11.114	8.737	+27.2%	1.70×10^{-64}
Binary: 2 if $\text{HS} \geq 2$ else 1	16.830	13.755	+22.4%	6.26×10^{-54}
Logarithmic: $\log(\text{HS}+2)$ v1	11.986	9.695	+23.6%	2.48×10^{-54}
No HS adjustment (HS term = 1)	9.068	7.343	+23.5%	4.95×10^{-64}

Figure 5 summarizes the sensitivity of the KEV-QRI separation across the QF and HS parameterizations of Tables 7 and 8.




Figure 5. Sensitivity analysis of the Q2 KEV–QRI separation: (a) sensitivity to QF parameterization across five scenarios; (b) sensitivity to HS functional form across six formulations. The dashed line marks the baseline 26.6% separation. All scenarios and formulas are significant at $p < 10^{-50}$. All formulas are significant at $p < 10^{-53}$.

4.5.3. QRI Versus Plain CVSS

Table 9 pits the QRI against the CVSS and EPSS using KEV as the binary outcome. The CVSS achieves 21.0% mean separation and an ROC-AUC of 0.739. QRI v2 raises mean the separation to 26.3%—5.3 percentage points above the CVSS—while its ROC-AUC (0.731) is marginally below that of the CVSS in binary classification. The EPSS, joined to the 78,150 CVEs in our dataset with EPSS coverage, reaches an ROC-AUC of 0.952 against the KEV outcome, which is substantially above both the CVSS and QRI. The EPSS is trained directly on exploitation signals and is designed for near-term binary triage, whereas the QRI is designed to encode quantum-specific risk dimensions (Shor/Grover vulnerability and HNDL attractiveness) that the EPSS does not measure. The two metrics answer different questions; their positioning is discussed in Section 5.3.

Table 9. Predictive performance: plain CVSS, QRI v2, and EPSS [14] computed on the present dataset, using CISA KEV as the outcome variable.

Metric	KEV Mean	Non-KEV Mean	Diff %	<i>p</i> -Value	ROC-AUC
Plain CVSS	8.384	6.929	+21.0%	1.48×10^{-67}	0.7394
QRI v2 (baseline)	13.048	10.328	+26.3%	1.02×10^{-62}	0.7307
EPSS [14] (empirical)	0.5955	0.0249	+2290.9%	2.89×10^{-234}	0.9523

Separation ranges from 24.6% (Scenario D, no QF adjustment) to 28.3% (Scenario C, maximum multipliers) across all five QF scenarios, all significant at $p < 10^{-54}$ (Table 7). Across all six HS functional forms (linear, cubic, binary threshold, logarithmic, and no adjustment), separation ranges from 22.4% to 27.2%, all with $p < 10^{-53}$ (Table 8). Table 9 shows that the QRI achieves better mean KEV/non-KEV separation than the plain CVSS (26.3% vs. 21.0%), with a comparable univariate ROC-AUC (QRI = 0.731; CVSS = 0.739).

These findings, together, indicate that the paper’s central empirical results are not an artifact of the particular functional form of the QRI or the particular parameter values used.

Figure 6 presents the predictive comparison and the Lorenz-curve concentration analysis.




Figure 6. Predictive performance and QRI concentration: (a) Mean KEV versus non-KEV scores for the plain CVSS and QRI v2; ROC-AUC values are annotated. Note that QRI’s ROC-AUC (0.731) is marginally below that of the plain CVSS (0.739), confirming that the QRI’s advantage is in rank-ordered triage, not binary classification. (b) Lorenz curve of the QRI distribution (Gini = 0.176), illustrating the low score concentration across the filtered corpus. Note: ROC-AUC computed using the Mann–Whitney U estimator (AUC = 0.5 = random; 1.0 = perfect). EPSS scores were obtained from the FIRST.org bulk feed and joined to the 78,150 CVEs in our dataset with EPSS coverage; the score is computed on the joined records.

4.5.4. Multivariable Logistic Regression—Does the QRI Add Information Beyond the CVSS?

The ROC-AUC comparison in Section 4.5.3 reports a univariate comparison: QRI versus CVSS, each tested in isolation as a single predictor of KEV membership. That framing leaves open a different question—whether the QRI adds incremental information beyond what the CVSS already encodes. To address this, we fitted a sequence of logistic regression models with KEV membership as the binary outcome and standardized predictors. Coefficients in standardized inputs make the contributions directly comparable on the log-odds scale.

Model 1 included CVSS alone; the standardized logistic-regression coefficient on the log-odds scale, which is distinct from the OLS trend slope (β_1) reported in Section 4.1, was $\beta = +0.970$, and AUC = 0.7394, reproducing the univariate finding. Model 2 added a standardized QRI score. The CVSS remained the dominant predictor— $\beta = +0.837$ and OR per SD = 2.31 (95% CI: 2.02–2.64)—but the QRI coefficient was significantly positive after CVSS adjustment: $\beta = +0.177$, OR per SD = 1.19 (95% CI: 1.10–1.30), and $p = 2.85 \times 10^{-5}$. A likelihood-ratio test of Model 2 against Model 1 gives $\chi^2(1) = 13.77$, $p = 0.0002$; the QRI improves model fit after the CVSS contribution is removed. Model 3 added the publication year, and the QRI coefficient was unchanged ($\beta = +0.177$, $p = 3.04 \times 10^{-5}$), indicating the QRI signal is not driven by a temporal artifact.

This is the correct response to the AUC observation. The QRI's univariate AUC (0.7487 for the additive model) is not larger than the CVSS's because the two scores share most of their variance—the QRI is, by construction, a multiplicative transformation of the CVSS. The meaningful question is whether the quantum adjustment factors carry information beyond classical severity. The logistic regression evidence is that they do, with an odds ratio per standard deviation of 1.19 (1.10–1.30), holding the CVSS fixed. For a vulnerability manager already filtering on the basis of the CVSS, the QRI provides an additional, statistically independent stratification of KEV likelihood—which is the operational claim the framework needs to support.

4.6. Summary of Hypothesis Tests

Table 10 summarizes the tests of hypotheses. Three out of four hypotheses are supported. Q1 proves that the HNDL attack surface has grown considerably. Q2 is the main result: the QRI predicts whether a vulnerability will be exploited, with a 26.6% mean separation in the exploited population. Q3 confirms the over-targeting of Shor-vulnerable CVEs and the behavioral validity of the HNDL score. Q4 is not confirmed by the pre-specified Gini criterion; we discuss this in Section 5.4.

Table 10. Summary of hypothesis tests, key statistics, and verdicts.

Q	Hypothesis	Key Statistics	Sig.	Verdict
Q1a	Annual crypto-CVE volume grows over 2016–2026	CAGR = 11.0%, $\beta_1 = 0.105$, $R^2 = 0.447$, $p = 0.025$	✓	Confirmed
Q1b	Annual Shor-vulnerable CVE count grows over 2016–2026	CAGR = 8.8%, $\beta_1 = 0.084$, $R^2 = 0.420$, $p = 0.031$	✓	Confirmed
Q2	CISA KEV CVEs carry higher QRI scores than non-KEV CVEs	$\Delta = +26.6\%$, $p = 2.17 \times 10^{-63}$, $ r = 0.46$	✓	Confirmed
Q3	Shor-vulnerable CVEs over-represented in CISA KEV	$\chi^2 = 9.476$, $df = 1$, $p = 0.002$; rate $1.52 \times$	✓	Confirmed
Q4	QRI risk concentrated in minority of CVEs	Gini = 0.176; top 70% hold 80% of QRI	✗	Not Confirmed

Note: CAGR = compound annual growth rate; $|r|$ = absolute rank-biserial correlation; χ^2 = chi-square. The Gini threshold for Q4 was pre-specified at >0.5 . In the Sig. column, ✓ indicates that the pre-specified significance criterion was met, and ✗ that it was not; the Verdict column reports the resulting decision on each hypothesis.

5. Discussion

The four hypothesis tests paint a consistent picture: the cryptographic attack surface is growing, exploitation is becoming concentrated in the upper portion of the QRI distribution, and the Shor-vulnerable subset dominates confirmed exploitation. The Pareto-like concentration test failed for reasons that can be interpreted and that do not compromise operational conclusions.

5.1. The Expanding Attack Surface and Its Urgency

The 11.0% CAGR reported in Section 4.1 is best considered in relation to the migration schedules. NIST IR 8547 estimates a 5–10-year migration for critical infrastructure [29]; CNSA 2.0 enforces a 2035 deadline for PQC adoption, with the first milestones in 2027 [31]. At the 11.0% CAGR, the number of cryptography-relevant CVEs roughly doubles every 6.6 years—comfortably within the transition time frame. The Shor-specific subset—currently 8092 CVEs—is the more immediately actionable concern. These are

vulnerabilities whereby the affected primitive (RSA private keys, ECDSA signatures, Diffie–Hellman session keys) becomes retrospectively transparent on Q-Day [32,33]. At 8.8% annual growth, the subset will grow by 130% in the final year of the Global Risk Institute’s decade-long quantum expectation window [34,35].

5.2. What the KEV Validation Says About Attacker Behavior

Q2 is the main empirical result. The CVEs currently being exploited by adversaries have mean QRI scores 26.6% higher than the rest of the cryptographic CVE population. The interpretation is that the vulnerabilities a future adversary would harvest for retrospective decryption are the same vulnerabilities being targeted today by classical adversaries. The percentile result is consistent with this—46.2% of KEV CVEs are above the 90th QRI percentile, compared to 10% null—consistent with Allodi’s observation that adversarial attention focuses on the top of the vulnerability distribution [28]. The methodological caveat is that the QRI’s ROC-AUC (0.731) is marginally below that of the plain CVSS (0.739), meaning the QRI does not improve binary classification. Its advantage lies in rank-ordered magnitude separation. Users running rank-ordered triage will benefit more from the QRI; users running thresholded filters will see limited incremental value over the CVSS.

5.3. Positioning the QRI Relative to Prior Frameworks

Table 11 positions the QRI relative to the three reference frameworks discussed above. The CVSS provides classical severity but no information about quantum exposure [24]. The EPSS provides highly precise near-term exploitability; on our dataset, its KEV ROC-AUC of 0.952 confirms it as the strongest single predictor of imminent exploitation, but the score is, by construction, quantum-neutral. QARS is the closest predecessor to the QRI in conceptual scope: it combines Mosca’s inequality with a multi-factor formula but is calibrated by expert judgment rather than by behavioral data [15]. The QRI fills a gap that none of these frameworks addresses: CVE-level quantum-risk quantification, validated against governmental exploitation data. The two metrics of most direct practical interest to the community—the EPSS and QRI—are not adversarial; they answer different questions. The EPSS quantifies the “most likely to be exploited in the next 30 days”, and the QRI quantifies which of the exploited and exploitable CVEs carry the largest forward-looking HNDL exposure. For organizations already using the EPSS for short-term triage, the QRI adds the quantum dimension that the EPSS does not encode; the overlap of high-EPSS and high-QRI records is the top target for remediation under a quantum-aware risk management framework. A natural extension discussed in Section 5.6 is a composite score that combines both signals in a single triage ranking.

Table 11. Positioning of the QRI relative to reference frameworks.

Dimension	CVSS (Base)	EPSS	QARS [15]	QRI (This Study)
Unit of analysis	CVE	CVE	System/org	CVE
Quantum exposure encoded	No	No	Theoretical	Yes (QF + HS)
HNDL attractiveness encoded	No	No	Partial	Yes (HS)
Validated vs. exploit data	Partial	Yes (AUC = 0.95)	No	Yes (AUC = 0.73)
Longitudinal attack-surface tracking	No	No	No	Yes (2016–2026)
Reproducible from open data	Yes	Yes	Partial	Yes
KEV mean separation	+21.0%	+2290.9%	—	+26.3%

Note: EPSS AUC computed empirically on our dataset [14]. QARS validation reflects absence of exploit-data validation in [15]. KEV separation figures for CVSS and QRI from Section 4.5.3.

5.4. Interpreting Q4: Why the Pareto Test Failed

Q4 is not confirmed by the study, not because of a problem with the metric but because of the design. The CVEs were filtered to identify those that are cryptographically relevant. On that subset, the ranges of the CVSS, QF, and HS are not broad enough to deliver a Gini coefficient close to 0.5. The heavy-tail character of the distribution of exploitation outcomes observed by Allodi [28] is partially absorbed by the filter that defines the QRI denominator. The practically relevant concentration is in the joint distribution of the QRI and KEV status. The version of the concentration result that is operationally meaningful is in the percentile result (46.2% of KEV CVEs above the 90th QRI percentile) reported in Section 4.2. The pre-specified Q4 hypothesis asked the wrong question of the data; we report it negatively and state what we should have asked—whether the exploitation is concentrated in the QRI upper tail—with the answer to that question being yes.

5.5. Limitations

Several limitations bound the following conclusions.

Text-based classification. QF and HS are determined by matching keywords in NVD English descriptions [36]. The convergent-validity analysis in Section Convergent Validity of the Keyword Classifier quantifies the limits of this approach: precision against the CWE reference label is 0.247 and Cohen's κ is 0.119, while recall is 0.949. This low precision reflects under-tagging of cryptographic CVEs within NVD's own CWE classifications, not an error introduced by the keyword classifier itself, and the high recall indicates the classifier successfully captures the great majority of CWE-confirmed cases. CVEs described at the level of memory corruption rather than the cryptographic-primitive level can be under-classified, and Shor-relevant CVEs are frequently tagged under product-implementation CWEs rather than cryptographic ones. Future research incorporating CPE product enumeration and a stricter Shor keyword list would increase precision; the present design errs on the side of recall to avoid missing crypto-relevant CVEs in the upstream filter.

CISA KEV as a ground-truth proxy. KEV reflects exploits that are publicly visible to U.S. authorities and is the most authoritative public CV list of its type [27], but it is not complete. The nation-state HNDL collection, which is, by design, unattributable, will not appear in KEV until (and if) a year or later after its disclosure. Thus, KEV gives a conservative estimate of the lower bound of the actual exploitation of quantum-relevant vulnerabilities.

The behavioral validity established in Section 4.2 is demonstrated in the context of currently exploited CVEs and should be read as evidence that the metric tracks realized threat activity to date, not as a guarantee of prospective validity once cryptographically relevant quantum computers become available and exploitation patterns may shift accordingly.

Partial 2026 cohort. CVEs published in the first quarter of 2026 are partial-year data; they are retained in the regression series as a single additional data point, which slightly attenuates the estimated slopes but does not change the direction or statistical significance of the trends.

Causal inference. Observational design supports predictive validation but not causal inference. The robustness checks performed in Section 4.5 show that the signal persists when the quantum adjustment is removed entirely, which suggests HNDL-specific dimensions contribute independently of classical severity, but full causal disentanglement is not feasible in this design.

Scope of disclosure. This study reflects only vulnerabilities publicly disclosed through the NVD; it does not capture undisclosed weaknesses or vulnerabilities held privately by vendors, state actors, or proprietary/closed-source systems. The distributions reported in

this study should be understood as characterizing the publicly visible cryptographic attack surface, not the entirety of it.

Generalizability. This NVD is aligned to Western disclosure practices. The QRI weighting is based on pre-quantum assumptions about which primitives are most vulnerable and will need to be adjusted as quantum hardware capability improves and the relative dominance of Shor and Grover attacks evolves. The QRI is best considered as an evidence-based prioritization tool rather than a risk metric, to be used in conjunction with EPSS scores [14], threat feeds, and exposure assessments specific to an asset.

5.6. Future Directions

There are several extensions to the current design. Substituting the keyword filter for a hybrid pipeline using CPE product enumeration, CWE product taxonomy mappings, and machine learning would improve precision. The QRI formula should be modified to capture estimated data retention times, network criticality, and industry-specific migration roadmaps as outlined by NIST IR 8547 [29]. Performing the longitudinal study on a rolling basis would permit the 11.0% CAGR to be compared with PQC migration rates, thereby empirically testing whether the attack surface slows down as classical primitives are removed.

The QRI could be correlated against patch release schedules and installed base remediation rates to see if organizations are behaving consistently with the urgency implied by the formula. The empirical complementarity established in Section 4.5.3 between the EPSS (ROC-AUC of 0.952 for near-term exploitation) and QRI (26.6% KEV separation on a quantum-aware composite) suggests a further line of work: a joint EPSS×QRI score [14], in which the two dimensions are combined into a single triage ranking that reflects both near-term and quantum-era risk. Calibrating such a composite—and validating it on the same KEV benchmark used here—is a natural extension of the present design.

A complementary direction concerns the operational cost of remediating the TLS-related CVEs that the QRI most heavily flags: empirical studies of post-quantum TLS 1.3 handshakes have shown that lattice- and hash-based signature schemes introduce measurable latency and bandwidth overhead [1], and pairing QRI-based prioritization with deployment-cost benchmarks would help organizations sequence migration not only by risk but also by feasibility.

6. Conclusions

The aim of this work was to empirically measure the harvest-now-decrypt-later attack surface, to suggest a composite scoring metric that captures the quantum-specific risk factors not encoded in CVSS, and to empirically verify said metric with respect to observed threat-actor behavior. In a dataset of 78,587 cryptographic CVEs found in the NVD from 2016 to the first quarter of 2026 and compared against the CISA KEV list, four hypotheses were tested. The principal empirical findings are summarized below.

6.1. Q1: Expanding Attack Surface

The trend analysis confirms the growth of cryptographic CVEs at a CAGR of 11.0% and of Shor-vulnerable CVEs at 8.8% annually. Given this, the cryptographic attack surface approximately doubles within the PQC migration time frame required by the NIST and NSA.

6.2. Q2: Behavioral Validation of the QRI Against CISA KEV

The mean QRI is 26.6% higher in the exploited population than in the entire cryptographic CVE sample ($p = 2.17 \times 10^{-63}$, $|r| = 0.46$) and is robust to all QF and HS model parameterizations, including complete removal of the quantum adjustment. To the best

of our knowledge, this is the first empirical evidence that a quantum-adjusted vulnerability score predicts real-world exploitation at the CVE scale, valid against an independent governmental account of in-the-wild exploitation.

6.3. Q3: Over-Targeting of Shor-Vulnerable CVEs

Shor-vulnerable CVEs are over-represented in KEV ($1.52 \times$ expected), and the HNDL score is further validated by a significant positive association between HS tier and exploit rate, concentrated at $HS \geq 4$. Remediation of high-QRI vulnerabilities will reduce short-term classical vulnerability while also reducing the vulnerability window for HNDL—a double dividend for remediation.

6.4. Q4: Pareto Concentration Hypothesis (Not Supported)

The Pareto-like concentration hypothesis was not supported. The QRI score-distribution *Gini* coefficient is 0.176, which is far less than the pre-announced criterion of >0.5 due to the corpus being filtered for cryptographic relevance before scoring and the filtered inputs not covering the range to produce high inequality. The practically relevant concentration result—46.2% of KEV CVEs above the 90th QRI percentile—is confirmed in Q2.

6.5. Synthesis and Outlook

The QRI is not intended to replace existing scoring systems. It occupies a particular niche that is currently unoccupied: CVE-level quantum-risk analysis validated using governmental exploitation data. Using the QRI in conjunction with the EPSS—which predicts short-term exploitation likelihood—provides a more comprehensive view than using either alone. The data and code for this analysis are available to the public. There are many areas of practical application, but the most urgent is in vulnerability management programs that do not yet have quantum-ready triage rules. The QRI offers a data-driven, replicable means of doing so, grounded in sources (NVD/CISA) already integrated into most enterprise vulnerability management pipelines. This research demonstrates that such ranking can be achieved and acted upon.

Supplementary Materials: The following supporting information can be downloaded at: <https://www.mdpi.com/article/10.3390/computers15080495/s1>, File S1: validation_sample.csv—the fixed stratified sample of 200 CVEs (50 per Quantum Factor tier) used for the convergent-validity analysis reported in Section Convergent Validity of the Keyword Classifier.

Author Contributions: Conceptualization, F.S.S., E.I. and S.P.; methodology, F.S.S., E.I. and S.P.; software, F.S.S., E.I. and A.P.; validation, F.S.S., E.I., S.P., A.P. and A.N.; formal analysis, F.S.S., E.I., S.P. and A.N.; investigation, F.S.S., E.I., S.P., A.P. and A.N.; resources, S.P.; data curation, F.S.S., E.I., S.P., A.P. and A.N.; writing—original draft preparation, F.S.S. and E.I.; writing—review and editing, F.S.S., E.I., S.P., A.P. and A.N.; visualization, F.S.S., E.I. and A.P.; supervision, E.I.; project administration, S.P.; funding acquisition, S.P. All authors have read and agreed to the published version of the manuscript.

Funding: The research and APC was funded by the Russian Science Foundation (project No. 25-11-20037, <https://rscf.ru/project/25-11-20037/> (accessed on 6 July 2026)).

Data Availability Statement: The data used for this study were derived from the following public-domain resources: the National Vulnerability Database (NVD) (<https://nvd.nist.gov/developers/vulnerabilities> accessed on 6 July 2026) and the CISA Known Exploited Vulnerabilities (KEV) Catalog (<https://www.cisa.gov/known-exploited-vulnerabilities-catalog> accessed on 6 July 2026).

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

AES	Advanced Encryption Standard
API	Application Programming Interface
AUC	Area Under the Curve
CAGR	Compound Annual Growth Rate
CI	Confidence Interval
CISA	Cybersecurity and Infrastructure Security Agency
CNSA	Commercial National Security Algorithm Suite
CPE	Common Platform Enumeration
CRQC	Cryptographically-Relevant Quantum Computer
CVE	Common Vulnerabilities and Exposures
CVSS	Common Vulnerability Scoring System
CWE	Common Weakness Enumeration
DH	Diffie–Hellman
DSA	Digital Signature Algorithm
ECDH	Elliptic Curve Diffie–Hellman
ECDSA	Elliptic Curve Digital Signature Algorithm
EPSS	Exploit Prediction Scoring System
FIPS	Federal Information Processing Standards
HMAC	Hash-based Message Authentication Code
HNDL	Harvest Now, Decrypt Later
HS	HNDL Score
HTTPS	Hypertext Transfer Protocol Secure
IPSec	Internet Protocol Security
KEV	Known Exploited Vulnerabilities
LDAP	Lightweight Directory Access Protocol
MD5	Message-Digest Algorithm 5
NIST	National Institute of Standards and Technology
NSA	National Security Agency
NVD	National Vulnerability Database
OLS	Ordinary Least Squares
OR	Odds Ratio
PKI	Public Key Infrastructure
PQC	Post-Quantum Cryptography
QARS	Quantum-Adjusted Risk Score
QF	Quantum Factor
QRI	Quantum Risk Index
ROC	Receiver Operating Characteristic
RSA	Rivest–Shamir–Adleman
SCADA	Supervisory Control and Data Acquisition
SD	Standard Deviation
SHA	Secure Hash Algorithm
SMTPS	Simple Mail Transfer Protocol Secure
SSH	Secure Shell
SSL	Secure Sockets Layer
STRIDE	Spoofing, Tampering, Repudiation, Information disclosure, Denial of service, Elevation of privilege
TLS	Transport Layer Security
VPN	Virtual Private Network

References

1. Sikeridis, D.; Kampanakis, P.; Devetsikiotis, M. Post-Quantum Authentication in TLS 1.3: A Performance Study. In *Proceedings of the 2020 Network and Distributed System Security Symposium, San Diego, CA, USA, 23–26 February 2020*; Internet Society: Geneva, Switzerland, 2020; pp. 1–16. Available online: <https://www.ndss-symposium.org/wp-content/uploads/2020/02/24203-paper.pdf> (accessed on 26 July 2026).
2. Taherdoost, H. A Comprehensive Review of Quantum-Resistant Architectures for Blockchain Security. *Sci* **2026**, *8*, 47. [CrossRef]
3. Udughapattuwa, P.; de Alwis, C.; Zeydan, E.; Wijewardhana, U.; Liyanage, M. Quantum-Resistant Security for Blockchain-Enabled 6G Networks: A Comprehensive Review. *IEEE Open J. Commun. Soc.* **2026**, *7*, 3253–3287. [CrossRef]
4. Bhatia, V.; Ramkumar, K.R. An Efficient Quantum Computing Technique for Cracking RSA Using Shor’s Algorithm. In *Proceedings of the 2020 IEEE 5th International Conference on Computing, Communication and Automation (ICCCA)*, Greater Noida, India, 30–31 October 2020; pp. 89–94. [CrossRef]
5. Gidney, C. How to Factor 2048 Bit RSA Integers with Less than a Million Noisy Qubits. *arXiv* **2025**, arXiv:2505.15917. [CrossRef]
6. Gidney, C.; Ekerå, M. How to factor 2048-bit RSA integers in 8 hours using 20 million noisy qubits. *Quantum* **2021**, *5*, 433. [CrossRef]
7. Kagai, F.; Branch, P.; But, J.; Allen, R. Harvest-Now, Decrypt-Later: A Temporal Cybersecurity Risk in the Quantum Transition. *Telecom* **2025**, *6*, 100. [CrossRef]
8. *FIPS 203*; Module-Lattice-Based Key-Encapsulation Mechanism Standard. Federal Information Processing Standards Publication; National Institute of Standards and Technology; U.S. Department of Commerce: Gaithersburg, MD, USA, 2024. [CrossRef]
9. *FIPS 204*; Module-Lattice-Based Digital Signature Standard. Federal Information Processing Standards Publication; National Institute of Standards and Technology; U.S. Department of Commerce: Gaithersburg, MD, USA, 2024. [CrossRef]
10. *FIPS 205*; Stateless Hash-Based Digital Signature Standard. Federal Information Processing Standards Publication; National Institute of Standards and Technology; U.S. Department of Commerce: Gaithersburg, MD, USA, 2024. [CrossRef]
11. Septien-Hernandez, J.-A.; Arellano-Vazquez, M.; Contreras-Cruz, M.A.; Ramirez-Paredes, J.-P. A Comparative Study of Post-Quantum Cryptosystems for Internet-of-Things Applications. *Sensors* **2022**, *22*, 489. [CrossRef] [PubMed]
12. Cybersecurity and Infrastructure Security Agency; National Security Agency; National Institute of Standards and Technology. *Quantum-Readiness: Migration to Post-Quantum Cryptography*; Joint Cybersecurity Factsheet; U.S. Department of Homeland Security: Washington, DC, USA, 2023. Available online: <https://www.cisa.gov/resources-tools/resources/quantum-readiness-migration-post-quantum-cryptography> (accessed on 26 July 2026).
13. Auer, R.; Dodson, D.; Dupont, A.; Haghighi, M.; Margaine, N.; Marsden, D.; McCarthy, S.; Valko, A. *Quantum-Readiness for the Financial System: A Roadmap*; BIS Papers No. 158; Bank for International Settlements: Basel, Switzerland, 2025; Available online: <https://www.bis.org/publ/bppdf/bispap158.pdf> (accessed on 26 June 2026).
14. Jacobs, J.; Romanosky, S.; Edwards, B.; Adjerid, I.; Roytman, M. Exploit Prediction Scoring System (EPSS). *Digit. Threat. Res. Pract.* **2021**, *2*, 20. [CrossRef]
15. Grigaliūnas, Š.; Brūzgienė, R. Towards a Unified Quantum Risk Assessment. *Electronics* **2025**, *14*, 3338. [CrossRef]
16. Mascelli, J.; Rodden, M. “Harvest Now Decrypt Later”: Examining Post-Quantum Cryptography and the Data Privacy Risks for Distributed Ledger Networks. *Financ. Econ. Discuss. Ser.* **2025**, 1–21. [CrossRef]
17. Bernstein, D.; Lange, T. Post-Quantum Cryptography. *Nature* **2017**, *549*, 188–194. [CrossRef] [PubMed]
18. Jaques, S.; Naehrig, M.; Roetteler, M.; Virdia, F. *Implementing Grover Oracles for Quantum Key Search on AES and LowMC*; Lecture Notes in Computer Science; Springer: Berlin/Heidelberg, Germany, 2020; Volume 12106, pp. 280–310. [CrossRef]
19. Asif, M.R.A.; Hasan, K.F.; Islam, M.Z.; Khondoker, R. STRIDE-Based Cyber Security Threat Modeling for IoT-Enabled Precision Agriculture Systems. *arXiv* **2022**, arXiv:2201.09493. [CrossRef]
20. Aggarwal, D.; Brennen, G.K.; Lee, T.; Santha, M.; Tomamichel, M. Quantum Attacks on Bitcoin, and How to Protect against Them. *arXiv* **2017**, arXiv:1710.10377. [CrossRef]
21. Taherdoost, H. A Critical Review of Blockchain Acceptance Models—Blockchain Technology Adoption Frameworks and Applications. *Computers* **2022**, *11*, 24. [CrossRef]
22. FIRST.Org, Inc. *Common Vulnerability Scoring System Version 4.0: Specification Document, Version 1.2*; Forum of Incident Response and Security Teams, CVSS Special Interest Group: Cary, NC, USA, 2024; pp. 1–40. Available online: <https://www.first.org/cvss/v4-0/cvss-v40-specification.pdf> (accessed on 26 July 2026).
23. Spring, J.; Hatleback, E.; Householder, A.; Manion, A.; Shick, D. Time to Change the CVSS? *IEEE Secur. Priv.* **2021**, *19*, 74–78. [CrossRef]
24. Wunder, J.; Kurtz, A.; Eichenmüller, C.; Gassmann, F.; Benenson, Z. Shedding Light on CVSS Scoring Inconsistencies: A User-Centric Study on Evaluating Widespread Security Vulnerabilities. *arXiv* **2024**, arXiv:2308.15259v2. [CrossRef]
25. Allodi, L.; Cremonini, M.; Massacci, F.; Shim, W. Measuring the Accuracy of Software Vulnerability Assessments: Experiments with Students and Professionals. *Empir. Softw. Eng.* **2020**, *25*, 1063–1094. [CrossRef]

26. National Institute of Standards and Technology. *Vulnerability APIs—NVD CVE API 2.0*; National Vulnerability Database, Information Technology Laboratory: Gaithersburg, MD, USA, 2023. Available online: <https://nvd.nist.gov/developers/vulnerabilities> (accessed on 26 July 2026).
27. Known Exploited Vulnerabilities Catalog | CISA. Available online: <https://www.cisa.gov/known-exploited-vulnerabilities-catalog> (accessed on 26 July 2026).
28. Allodi, L.; Massacci, F. Comparing Vulnerability Severity and Exploits Using Case-Control Studies. *ACM Trans. Inf. Syst. Secur.* **2014**, *17*, 1–20. [CrossRef]
29. Regenscheid, A. *Transition to Post-Quantum Cryptography Standards*; National Institute of Standards and Technology: Gaithersburg, MD, USA, 2024; pp. 1–29. Available online: <https://nvlpubs.nist.gov/nistpubs/ir/2024/NIST.IR.8547.ipd.pdf> (accessed on 26 July 2026).
30. Brilhante, M.d.F.; Pestana, D.; Pestana, P.; Rocha, M.L. Measuring the Risk of Vulnerabilities Exploitation. *AppliedMath* **2024**, *4*, 20–54. [CrossRef]
31. National Security Agency. *The Commercial National Security Algorithm Suite 2.0 and Quantum Computing FAQ*; Cybersecurity Information Sheet U/OO/194427-22, PP-24-4014, Ver. 2.1; National Security Agency: Fort Meade, MD, USA, 2024. Available online: https://media.defense.gov/2022/Sep/07/2003071836/-1/-1/1/CSI_CNSEA_2.0_FAQ.PDF (accessed on 26 July 2026).
32. Joseph, D.; Misoczki, R.; Manzano, M.; Tricot, J.; Pinuaga, F.D.; Lacombe, O.; Leichenauer, S.; Hidary, J.; Venables, P.; Hansen, R. Transitioning Organizations to Post-Quantum Cryptography. *Nature* **2022**, *605*, 237–243. [CrossRef] [PubMed]
33. Pasupuleti, M.K. Post-Quantum Secure Cloud Architectures: Integrating AI-Driven Cryptography, Blockchain Protocols, and Number-Theoretic Foundations for Scalable Privacy. *Int. J. Acad. Ind. Res. Innov.* **2025**, *5*, 49–73. [CrossRef]
34. Acharya, R.; Aghababaie-Beni, L.; Aleiner, I.; Andersen, T.I.; Ansmann, M.; Arute, F.; Arya, K.; Asfaw, A.; Astrakhantsev, N.; Atalaya, J.; et al. Quantum error correction below the surface code threshold. *Nature* **2025**, *638*, 920–926. [CrossRef] [PubMed]
35. Mosca, M.; Piani, M. *Quantum Threat Timeline Report 2024*; Global Risk Institute: Toronto, ON, Canada, 2024; Available online: <https://globalriskinstitute.org/publication/2024-quantum-threat-timeline-report/> (accessed on 26 July 2026).
36. Anwar, A.; Abusnaina, A.; Chen, S.; Li, F.; Mohaisen, D. Cleaning the NVD: Comprehensive Quality Assessment, Improvements, and Analyses. *IEEE Trans. Dependable Secur. Comput.* **2022**, *19*, 4255–4269. [CrossRef]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.