CIRS: An Engineering Decision Framework for Counterfeit Risk Scoring in Electronic Component Supply Chains

Alexander Litvin, *Senior Member, IEEE*

***Abstract*—Counterfeit electronic components remain a persistent integrity risk across industrial, automotive, medical, critical-infrastructure, and defense electronics supply chains. Existing process standards address this qualitatively but produce no reproducible quantitative score at the operational procurement level. This paper develops the Composite Integrity Risk Score (CIRS), a risk scoring model that computes a normalized $[0, 1]$ score via a structured weighted aggregation of five observable risk parameters — counterfeit trade intensity (F), supply channel structure (S), market pressure (M), verification quality (Q), and systemic geographic concentration (G) — together with a documented free-trade-zone routing amplifier (I_{FTZ}). To motivate the composite structure empirically, we test five ordinary least squares (OLS) specifications on a 20-year ERAI panel (2005–2024) against semiconductor market data. Market growth is statistically non-significant across all specifications ($p = 0.12$–0.73), while a post-2012 mandatory reporting regime accounts for most explained variance ($\beta \approx +473$ in the full M5 specification, $p < 0.01$). Market growth alone is thus insufficient under reporting-regime confounding, supporting the composite design. CIRS components are anchored to OECD GTRIC trade-risk indices, public supplier-audit evidence, and the SAE AS6171 test hierarchy. Sensitivity analysis identifies supply channel structure as the most influential parameter, while full verification testing provides a bounded maximum CIRS reduction of 0.150. Applied to five procurement scenarios and three documented incident profiles, CIRS provides structurally consistent score discrimination across procurement risk contexts, combining channel provenance and verification depth into a unified risk metric — a gap not addressed by existing process standards.**

***Index Terms*—composite risk index, counterfeit electronic components, procurement risk management, supply chain integrity, supply chain risk assessment**

I. Introduction

The global market in counterfeit and pirated goods reached an estimated USD 467 billion in 2021, representing 2.3 percent of world trade [1]. Electronics (including semiconductor components, integrated circuits, and passive devices) constitute one of the most persistently affected product categories, subject to counterfeiting across both obsolete and actively manufactured part numbers [2], [13]. The consequences range from system failures and production line shutdowns to safety-critical incidents in aerospace, defense, and medical devices [3].

A. Litvin is an Independent Researcher in electronic component supply chain risk and integrity analytics, Yekaterinburg, Russia (e-mail: a67444152@outlook.com).

ORCID: 0009-0009-6744-4152.

Despite the scale of the problem, risk assessment in electronic component procurement remains largely unscored. Industry frameworks such as SAE AS5553, AS6081, and AS6171 specify process requirements for counterfeit avoidance but do not produce a quantitative risk score. Supply chain risk management literature has developed probabilistic and composite approaches for general disruption risk [4], [5], but domain-specific instruments calibrated to the counterfeit-specific risk structure of electronic components are absent from the academic record.

Recent industry developments confirm that supply chain integrity requires systematic, quantifiable documentation at every level of the electronics value chain. Vendor-reported data show that leading semiconductor manufacturers now implement dedicated supply chain attestation programs covering manufacturing origin verification across geographically distributed production corridors [15]. These programs address the silicon production level; no equivalent scoring instrument exists at the component procurement and incoming inspection level that CIRS is designed to serve. Recent industry and standards-oriented discussions continue to emphasize the need for structured, risk-oriented procurement practices in semiconductor and electronic component supply chains. Practitioner commentary has likewise framed these issues in terms of risk logistics and procurement transparency [18].

This paper addresses that gap by developing and evaluating the Composite Integrity Risk Score (CIRS), an index designed for operational use at the procurement decision level. CIRS integrates five risk dimensions into a single $[0, 1]$ score with interpretable thresholds, calibrated against publicly available trade intelligence, defense audit data, and standardized test protocol hierarchies.

The paper makes three contributions distinct from existing approaches.

C1 — Empirical finding: We report a panel regression analysis of counterfeit incident dynamics across a 20-year ERAI dataset, controlling for mandatory regulatory reporting regime changes (NDAA 2012). The results indicate that a regulatory reporting shift, rather than market growth, accounts for the largest identifiable share of variance in observed incident counts ($\beta \approx +473$ in the full M5 specification, $p < 0.01$, $\Delta R^2 = 0.34$), with implications for how industry incident data are interpreted in risk modeling contexts.

C2 — Methodological contribution: CIRS integrates supply chain traceability evidence, test protocol evidence, and systemic documentation integrity into a unified quantitative risk signal, a combination not addressed by existing process standards or prior composite approaches. It departs from existing methods in three specific ways. First, unlike process-standard frameworks (AS5553, AS6081, AS6171), CIRS produces a quantitative $[0, 1]$ score reproducible across organizations without requiring expert consensus or certification infrastructure. Second, unlike multi-criteria decision analysis (MCDA) methods such as the analytic hierarchy process (AHP) and the technique for order of preference by similarity to ideal solution (TOPSIS), which derive weights from expert judgment matrices, CIRS parameters are anchored to empirical sources: $k = 0.60$ as a sensitivity-tested baseline proxy informed by publicly reported defense-supplier audit observations [12], $\delta_{FTZ} = 0.059$ from OECD customs seizure analysis [9], and Q weights from the monotonic AS6171 test hierarchy. Third, unlike single-factor or binary channel thresholds, CIRS integrates product-category counterfeit exposure (F), supply channel provenance (S), market pressure (M), provenance-economy concentration (G), documented FTZ routing (I_{FTZ}), and verification depth (Q) into a single score, enabling discrimination within risk strata where procurement decisions are most consequential.

C3 — Design rule: CIRS quantifies a practical verification ceiling: the maximum CIRS reduction attributable to full AS6171 testing is bounded at 0.150. This is consistent with qualitative observations in prior literature [11] and with empirical test-laboratory evidence that sophisticated cloned parts carry few detectable defects and can evade standards-based methods [25], and it yields an exact decision rule for risk-budget allocation: a procurement event with baseline CIRS ≥ 0.60 (at $Q = 0$) cannot reach the Moderate tier through testing alone, since $0.60 - 0.150 = 0.45$ remains on the High boundary. More generally, channel structure (S) is not substitutable by verification depth (Q); above the upper High range, channel-level intervention is required.

II. Background

A. The Counterfeit Electronics Problem

Counterfeit electronic components were first characterized in the IEEE community as a systemic supply-chain integrity problem rather than a narrow authentication issue [20]. Twenty years on, the problem persists. ERAI Inc. reported 1,055 suspect counterfeit and nonconforming electronic parts in 2024, a 25 percent increase over 2023 and the highest count since 2015 [2]. Primary component types (Analog ICs, Microprocessor ICs, Memory ICs, and Programmable Logic ICs) have remained consistent over a decade. A notable 2024 development: 21 percent of reported manufacturer brands were encountered for the first time, and 29.4 percent of all parts belonged to brands never previously reported to ERAI [2], indicating that counterfeiters track procurement demand across a broadening range of manufacturers.

The OECD/EUIPO series provides macroeconomic context. Global trade in counterfeit goods stood at 2.5 percent of world trade in 2013 (USD 461 billion), rose to 3.3 percent in 2016 (USD 509 billion), and normalized to 2.3 percent in 2021 (USD 467 billion) [1], [6], [7]. Electronics constitute one of the top-targeted categories by seizure value [1], [13]. Free trade zones (FTZs) function as structural amplifiers: OECD analysis of 2013–2016 customs data found a 5.9 percent increase in counterfeit export probability associated with FTZ routing, attributable to reduced oversight and transshipment opacity [9].

B. Regulatory and Standards Context

The pivotal regulatory event was Section 818 of the National Defense Authorization Act (NDAA) of 2012, which established mandatory reporting requirements for counterfeit electronic parts within the U.S. Department of Defense supply chain. Prior to this legislation, reporting was voluntary and inconsistent; the Senate Armed Services Committee investigation found that the Defense Logistics Agency “neither consistently reported to GIDEP nor maintained a list of instances in which they had been supplied counterfeit electronic parts” [10], where GIDEP denotes the Government–Industry Data Exchange Program. The mandatory reporting regime, reinforced by DFARS clause 252.246-7007 (2014), transformed industry reporting systems including ERAI into compliance-driven infrastructure.

On the standards side, the SAE G-19 Committee has produced the primary counterfeit avoidance framework: AS5553 (OEM/contractor), AS6081 (independent distributor), and AS6171 (test laboratory methods). The 2023 revision AS6081A aligned detection methods with the AS6171 family. The SAE standards family also includes AS6496, governing counterfeit avoidance requirements specifically for authorized distributors — the channel type assigned the lowest risk score in CIRS ($S = 0.05$). Unlike AS6081, which addresses independent distributor practices, AS6496 requires direct authorization from the original component manufacturer and mandates full backward traceability to the OCM source, providing the formal standards basis for the low-risk authorized channel tier in the CIRS S variable. Industry compliance data indicate that 89 percent of authentication tests performed in 2022 still referenced the older AS6081 standard rather than the enhanced AS6171 methodology [3]. This compliance lag is directly operationalized in the CIRS Q variable.

The OECD GTRIC (General Trade-Related Index of Counterfeiting) provides CIRS components F and G: GTRIC-p for product-level counterfeit propensity (F) and GTRIC-e for economy-level source propensity (G); both are published indices bounded in $[0, 1]$ by construction. The frequently cited figure that one economy accounted for approximately 45 percent of global counterfeit seizures in 2021 is descriptive context and is distinct from the GTRIC-e propensity value used to parameterize G [1]. The FTZ premium documented in OECD (2018) reflects a governance gap not addressed by existing SAE standards, which focus on testing and procurement process rather than routing topology [9]. CIRS incorporates this premium as an additive amplifier, providing a quantitative integration of the FTZ risk channel into a component-level scoring instrument, a relationship not previously integrated into standards-based frameworks.

C. Electronic Component Lifecycle and Obsolescence Risk

A structural driver of counterfeit vulnerability not fully captured by macroeconomic proxies is the lifecycle mismatch between electronic components and the systems that contain them. Defense, aerospace, and industrial systems routinely remain in service for 20–40 years, while the components that comprise them have commercial manufacturing lifespans of 5–15 years. This mismatch creates what the U.S. DoD terms Diminishing Manufacturing Sources and Material Shortages (DMSMS), a condition in which required components can no longer be procured from their original manufacturers [10]. DMSMS events force procurement managers into non-authorized channels (independent distributors, brokers, and spot markets), precisely the channels with the highest counterfeit risk exposure as documented in the S variable of CIRS.

The 2024 ERAI data provides direct evidence of this mechanism: 42.75 percent of reported counterfeit incidents involved obsolete parts [2]. However, a critical nuance emerges: active components readily available through authorized sources accounted for 27.2 percent of all incidents, and were reported more than twice as frequently as components with long manufacturer lead times [2]. This finding challenges the assumption that counterfeit risk concentrates in scarce or obsolete parts, and directly motivates the CIRS operationalization of S as a channel structure variable rather than a scarcity proxy. NIST SP 800-161r1 identifies this supply chain lifecycle pressure as a primary risk driver requiring systematic monitoring and cybersecurity supply chain risk management (C-SCRM) controls [14]. CIRS provides the quantitative scoring layer that C-SCRM frameworks reference but do not themselves supply.

D. Existing Risk Assessment Approaches

The broader supply chain risk management (SCRM) literature has developed analytical frameworks for risk identification, disruption modeling, and decision support under uncertainty [22], [24]. IEEE-side contributions formalized risk in supply chains across strategic, operational, and tactical levels, distinguishing deviations, disruptions, and disasters [23]. These approaches address disruption exposure or resilience at the network level rather than counterfeit-specific procurement scoring for electronic components.

Counterfeit risk frameworks in the electronic component domain fall into two categories. Process-oriented frameworks (AS5553, AS6081, AS6171) specify procedural requirements for avoidance, detection, and disposition but do not produce a quantitative risk score [3], [14]. Qualitative risk assessment tools (supplier scorecards, risk matrices, GIDEP counterfeit alerts) capture channel and vendor factors but lack standardized metrics and do not incorporate macroeconomic or systemic trade-level exposure. The gap between these categories has been identified in the literature [5] and earlier addressed in part through a semi-quantitative standard for counterfeit detection [21], but not filled by an operationalized scoring instrument that combines channel, verification, and macroeconomic dimensions.

Network modeling approaches have addressed structural dynamics of counterfeit supply chains [11], providing theoretical grounding for channel-level risk propagation. Probabilistic supply chain risk frameworks have established composite index methodology for related domains [4]. CIRS synthesizes these threads into a domain-specific instrument calibrated to the data sources available for electronic component procurement.

III. Methodology

A. Study Design

This study develops and evaluates CIRS through two complementary activities: (1) empirical regression analysis of observed counterfeit incident data to establish the necessity of a composite approach, and (2) composite index construction integrating five risk dimensions with empirically grounded calibration parameters. The objective is operational risk modeling under data constraints, not causal identification. CIRS is not proposed as a predictive classifier of counterfeit outcomes; it is an auditable engineering decision framework that transforms observable procurement evidence into a reproducible structural risk score under sparse outcome-labeled data conditions. All variables, weights, and thresholds were specified prior to data application.

B. Counterfeit Incident Data

Industry counterfeit activity is proxied using ERAI-reported incident counts (2005–2024) [2]. These data represent observed and reported incidents, not total counterfeit prevalence; Y_t is treated as an observed, reporting-regime-sensitive signal subject to both underlying risk and reporting-regime effects, where Y_t denotes the ERAI-reported counterfeit incident count in year t. This limitation is consistent with documented industry observations: a NIST U.S. Resilience Project case study reports that widely used industry reporting systems such as ERAI and GIDEP capture only a small fraction of actual counterfeit incidents, reflecting reporting infrastructure capacity rather than total market activity [16]. The 2024 value ($n = 1{,}055$) is confirmed in report text [2]; values for 2005–2023 were extracted from graphical representations (digitization error ± 10 units).

For the composite index, F is operationalized directly as the OECD GTRIC-p index for the electronics product module (HS85). GTRIC-p is bounded in $[0, 1]$ by construction, with higher values indicating higher category-level counterfeit propensity; F therefore adopts a published, cross-sectionally comparable propensity measure and introduces no author-specific normalization. Because GTRIC-p is reported at the HS 2-digit level, F reflects category-level propensity averaged across the electronics module rather than component-specific propensity [1], [8].

C. Market Variable

Market conditions are represented by annual global semiconductor sales in USD billions (World Semiconductor Trade Statistics, WSTS, historical series). Both contemporaneous (M_t) and lagged (M_{t-1}) specifications are estimated.

D. Regulatory Regime Controls

$D_{NDAA} = 1$ for $t \geq 2012$, zero otherwise, capturing the shift in reporting intensity associated with NDAA Section 818. $D_{COVID} = 1$ for 2020–2021, capturing the manufacturing shutdown structural break; included in the robustness specification only.

E. Regression Specifications

Five ordinary least squares (OLS) models were estimated ($n = 19$–20):

$$M1: \quad Y_t = \beta_0 + \beta_1 M_t + \varepsilon_t$$
$$M2: \quad Y_t = \beta_0 + \beta_1 M_t + \beta_2 M_{t-1} + \varepsilon_t$$
$$M3: \quad Y_t = \beta_0 + \beta_1 M_t + \beta_2 D_{NDAA} + \varepsilon_t$$
$$M4: \quad Y_t = \beta_0 + \beta_1 M_t + \beta_2 M_{t-1} + \beta_3 D_{NDAA} + \varepsilon_t$$
$$M5: \quad Y_t = \beta_0 + \beta_1 M_t + \beta_2 D_{NDAA} + \beta_3 D_{COVID} + \varepsilon_t$$

F. CIRS Composite Score: Formal Specification

The Composite Integrity Risk Score is defined as:

$$\text{CIRS} = w_F F + w_S (k \cdot S) + w_M M - w_Q Q + w_G G + \delta_{FTZ} I_{FTZ} \quad (1)$$

where the variables and ranges are:

- $F \in [0, 1]$: counterfeit trade intensity (OECD GTRIC-p, normalized)
- $S \in [0, 1.2]$: supply channel risk (categorical). Authorized channels ($S = 0.05$) are characterized by end-to-end backward traceability to the OCM, documented chain-of-custody, and certificate-of-conformance mechanisms consistent with AS6496 requirements; independent distributor channels ($S = 0.40$) provide partial traceability under AS6081; non-certified and unknown channels ($S = 0.80$–1.00) lack these controls.
- $M \in [0, 1]$: market pressure (normalized semiconductor sales)
- $Q \in [0, 1]$: verification quality (AS6171 test battery; Eq. 2)
- $G \in [0, 1]$: systemic concentration (OECD GTRIC-e, normalized)
- $I_{FTZ} \in \{0, 1\}$: free trade zone routing indicator

Parameters: $k = 0.60$ is used as a defensible baseline proxy for channel-level risk transfer under non-authorized or weakly traceable procurement conditions, informed by publicly reported defense-supplier audit observations [12]. It is not interpreted as a universal physical constant. Sensitivity analysis over $k \in [0.40, 0.80]$ — consistent with published audit observations for independent distributor and broker-mediated procurement channels — confirms that ordinal ranking is fully stable (Kendall $\tau = 1.00$ across all tested values). Risk-category assignments are largely preserved: a single boundary case (D) transitions from High to Moderate at $k = 0.40$, while all other cases remain unchanged across the full range (Section IV-C); $\delta_{FTZ} = 0.059$ (FTZ counterfeit export premium, OECD 2018 [9] — the only OECD publication providing a FTZ-specific quantitative estimate; this value is not superseded in OECD 2025 [1], which updates GTRIC trade indices but does not report an updated FTZ-specific premium).

Baseline weights: $w_F = 0.25$, $w_S = 0.25$, $w_M = 0.20$, $w_Q = 0.15$, $w_G = 0.15$. Weights are structurally assigned and validated through sensitivity analysis rather than statistically estimated from data, reflecting a deliberate conservative design choice. The near-equal weight structure follows the simple-weighting rationale in composite index design: when the empirical covariance structure is unknown, simple weighting schemes perform comparably to optimized-weight variants [4]. This approach avoids overfitting to any single observed dataset, preserves interpretability, and is more robust than AHP or PCA-derived alternatives under the data-scarce conditions typical of this domain. Sensitivity to $\pm 20\%$ single-weight perturbation is reported in Section IV-C; classification boundaries are stable.

Table I summarizes parameter provenance for key calibration constants.

The Q sub-score is:

$$Q = 0.20 I_{vis} + 0.25 I_{elec} + 0.20 I_{x2d} + 0.15 I_{xct} + 0.20 I_{dec} \quad (2)$$

where $I_{method} \in \{0, 1\}$ indicates application of: visual inspection (AS6081 baseline), electrical testing (AS6081 extended), 2D X-ray, CT X-ray (AS6171 L2–L3), and decapsulation (AS6171 full). Industry baseline: $Q \approx 0.49$, derived from compliance data: 89% at AS6081 ($Q = 0.45$) and 11% at AS6171 ($Q = 0.85$) [3].

Risk thresholds: Low $[0, 0.25)$, Moderate $[0.25, 0.45)$, High $[0.45, 0.65)$, Critical $[0.65, 1.0]$.

G. Justification of Index Design Over Machine Learning

A natural question is whether a machine learning model would outperform a structured composite index. CIRS is not intended as a predictive model, but as a decision-support framework under conditions of incomplete procurement information. Three considerations make the composite index the appropriate design choice. First, *explainability*: procurement decisions subject to DFARS 252.246-7007 or FAR 52.246-26 require auditable, interpretable justification that ML models cannot provide [14]. Second, *data scarcity*: the ERAI panel contains 20 annual observations, insufficient for meaningful ML training without severe overfitting; even if applied, ML outputs would lack robust generalization under the observed data constraints. Composite indices operate under data scarcity by encoding domain knowledge in parameter structure rather than learning it from observations [4]. Third, *distributional stability*: Section IV-A establishes that ERAI data are primarily shaped by regulatory regime changes, not stable underlying relationships; a structured index with empirically anchored parameters provides more stable risk assessments across regulatory regime transitions. This positioning is complementary to, not in competition with, emerging machine-learning and big-data authentication methods [26]: those operate at the physical-inspection or transaction-recording layer, whereas CIRS supplies an upstream, auditable channel-structure score computable from information available at the procurement decision point.

TABLE I
CIRS PARAMETER PROVENANCE

Parameter	Source	Value	Limitation
k (channel-risk transfer proxy)	Public defense-supplier audit evidence; sensitivity-tested baseline [12]	0.60	Single public audit; sector-specific
δ_{FTZ} (FTZ amplifier)	OECD/EUIPO [9]	0.059	Customs aggregate; route-specific uncertainty
w_F, w_S, w_M, w_Q, w_G	Structurally assigned (near-equal) baseline [4]	0.25/0.25/0.20/0.15/0.15	Not empirically optimized
F (counterfeit intensity)	OECD GTRIC-p [1]	Category-specific	Annual; product-level aggregation
Q weights	SAE AS6171 hierarchy	Per test tier	Standard-defined; not validated empirically
G (systemic concentration)	OECD GTRIC-e [1]	Provenance-specific baseline	Economy-level aggregation
M (market pressure)	WSTS semiconductor sales (normalized)	Annual proxy	Macro-level; not transaction-specific

H. Monotonicity and Behavioral Validation

Fig. 1 illustrates CIRS behavior as a function of supply channel risk parameter S while holding baseline macroeconomic and verification parameters constant ($F = 0.35$, $M = 0.50$, $Q = 0.45$, $G = 0.45$; $Q = 0.45$ = AS6081 baseline). The monotonically increasing curve confirms that higher-risk channel structures consistently produce higher composite scores. The near-linear appearance reflects fixed baseline assumptions; the full CIRS formulation (1) is piecewise linear over the complete joint parameter space, with non-smoothness arising only from score clipping and the discrete FTZ step. Table III cases (circles) are plotted at their actual CIRS values, which reflect case-specific M and Q parameters; their deviation below the baseline curve reflects lower market pressure or higher verification depth than the baseline assumption. Table IV directional consistency cases (triangles) align with the monotonicity property. The FTZ amplifier ($\delta_{FTZ} = 0.059$) produces a discontinuous upward jump at $S = 1.00$, visible as Case E above Case D. The teal bracket marks the maximum Q-attributable mitigation (0.150) under full AS6171 testing; its decision-rule implications are quantified in Section IV-B.

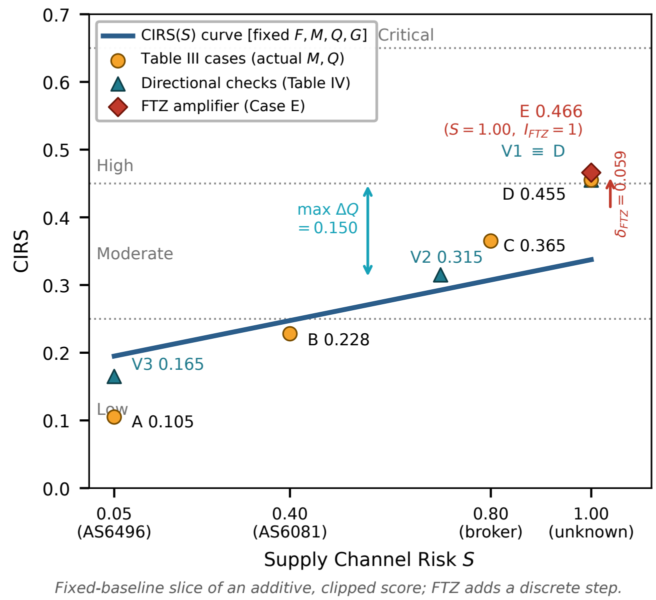

Fig. 1. CIRS monotonicity and channel structure sensitivity. The curve shows CIRS as a function of the supply channel risk parameter S under fixed baseline conditions ($F = 0.35$, $M = 0.50$, $Q = 0.45$, $G = 0.45$; $Q = 0.45$ reflects the AS6081 compliance baseline). Circles represent illustrative cases from Table III (plotted at actual case-specific M, Q values); triangles indicate directional consistency checks derived from Table IV. The figure demonstrates monotonic model behavior and highlights the dominant effect of channel structure relative to bounded verification gains ($\max \Delta Q = 0.150$, bounded mitigation). FTZ routing is shown as an additive amplification effect independent of verification depth. Risk threshold lines are defined for interpretability and do not affect the CIRS computation.

IV. RESULTS

A. Regression Analysis

Table II reports OLS estimates across five specifications ($n = 19$–20, 2005–2024).

Market volume (M_t) is non-significant across all five specifications (best $p = 0.12$). The D_{NDAA} dummy is the strongest predictor in this exploratory specification ($\beta \approx +455$–480 across M3–M5; $p < 0.01$ in M3 and M5, $p < 0.05$ in M4), accounting for an R^2 increment of approximately 0.34. The M3 estimate is $\beta \approx +480$, while the full M5 specification reports $\beta \approx +473$; both indicate that the post-2012 reporting regime is the dominant explanatory term. The heteroscedasticity- and autocorrelation-consistent (HAC) corrected results for M5 are reported in Appendix A: the point estimates and R^2 remain identical to Table II (HAC adjusts only standard errors and confidence intervals, not the point estimates or goodness-of-fit measures). Under HAC correction, β_{NDAA} remains highly significant ($t_{HAC} = 2.86$, $p < 0.01$). The COVID-19 period produces a suppression of $\beta \approx -458$ ($p < 0.05$). The full specification (M5) achieves $R^2 = 0.56$; 44% of variance remains unexplained. Spearman $\rho = 0.44$ ($p = 0.07$) between Y_t and M_t, excluding 2020–2021. Full regression diagnostics including Durbin–Watson ($DW = 1.273$), Newey–West HAC standard errors ($t_{HAC}(D_{NDAA}) = 2.86$, $p < 0.01$), and COVID-19 exclusion robustness are reported in Appendix A. The non-significance of M across all specifications is informative: it confirms that market-level indicators are insufficient predictors of counterfeit exposure in the presence of structural confounders, motivating the composite framework.

TABLE II
OLS REGRESSION: ERAI INCIDENTS (Y_t) $\sim$ WSTS SEMICONDUCTOR SALES (M_t) + CONTROLS

Variable	M1	M2	M3	M4	M5
Intercept	699***	759***	894***	932***	778***
M_t	0.30	−0.58	−1.07	−1.23	−0.62
	(0.60)	(0.73)	(0.12)	(0.40)	(0.30)
M_{t-1}	—	0.80	—	0.13	—
		(0.65)		(0.93)	
D_{NDAA}	—	—	**480*****	**455****	**473*****
			(0.008)	(0.019)	(0.003)
D_{COVID}	—	—	—	—	**−458****
					(0.014)
R^2	0.016	0.016	0.354	0.326	**0.561**
R^2 adj	−0.039	−0.108	0.278	0.191	**0.478**
RMSE	278.6	274.5	225.8	227.1	**186.2**

p-values in parentheses. *$p < 0.10$, **$p < 0.05$, ***$p < 0.01$. $n = 20$ (M1,M3,M5), $n = 19$ (M2,M4). Data: [2] + WSTS.

B. CIRS Illustrative Application

Table III presents CIRS calculations for five synthetic procurement cases designed to isolate the contribution of channel structure (S) and verification depth (Q) variation while holding macroeconomic parameters constant. Use of synthetic cases is standard practice in the composite index literature [4], [5] and is appropriate for two reasons. First, transaction-level procurement data from electronic component supply chains are commercially sensitive [17]; synthetic cases based on publicly documented channel type distributions [1], [2] provide verifiable parameter provenance without confidentiality constraints. Second, holding F and G constant directly demonstrates how CIRS responds to operationally controllable parameters.

CIRS differentiates clearly across the five scenarios and separates Low, Moderate, and High risk tiers. The maximum Q-attributable risk reduction is bounded at -0.150 across all cases. For Case D, full AS6171 testing ($Q = 1$) reduces CIRS from its $Q = 0$ baseline of 0.455 to 0.305 (Moderate); for Case E (FTZ-routed), the corresponding reduction takes CIRS from a baseline of 0.534 (at $Q = 0$) to 0.384 under full AS6171 testing ($Q = 1$). The ceiling becomes structurally binding for events with baseline CIRS ≥ 0.60, where full testing cannot reach the Moderate tier — a property of the model rather than an artifact of a specific scenario.

C. Sensitivity Analysis

One-way sensitivity analysis varies each input parameter across its operational range while holding all others at the baseline of Fig. 1 ($F = 0.35$, $S = 0.40$, $M = 0.50$, $Q = 0.45$, $G = 0.45$, $k = 0.60$). The tornado analysis identifies S as the parameter with the highest CIRS sensitivity (a swing of up to $+0.12$ as S moves from the AS6081 baseline of 0.40 to its upper bound of 1.20 at $k = 0.60$), followed by M (± 0.08 over $M \in [0.10, 0.90]$), F (± 0.05 over $F \in [0.15, 0.55]$), w_S weight perturbation (± 0.04), and G (± 0.02).

Weight perturbation robustness. Rankings under $\pm 20\%$ perturbation of each individual weight were tested across all five cases (10 experiments, renormalizing remaining weights proportionally). Kendall rank correlation between baseline and perturbed rankings equals 1.00 in 9 of 10 tests and 0.80 in one (mean $\tau = 0.98$). The single lower-τ case ($w_Q + 20\%$) produces category reassignment for two boundary cases located within 0.02 CIRS units of a threshold. Maximum score deviation under any $\pm 20\%$ single-weight perturbation is bounded at 0.027 CIRS units. Robustness of rankings under $\pm 20\%$ weight perturbations was tested, confirming rank stability across all cases except those at classification boundaries, which should be interpreted with a ± 0.03 precision margin.

D. Directional Validation Against Public Incident Records

To provide directional validation, CIRS scores are computed for three publicly documented incident profiles and compared against known outcomes (Table IV).

Case V1 (Hong Dark Electronic Trade): an unauthorized Chinese broker supplying DMSMS-classified components during peak shortage with no incoming inspection [10]. CIRS High (0.455) is consistent with the documented outcome: approximately 84,000 suspect counterfeit parts inserted into DoD supply chains. Case V2 (defense-supplier audit profile): CIRS Moderate (0.315) is directionally consistent with the audit finding that 60% of independent sources present moderate-to-high risk [12]. Case V3 (authorized distributor): CIRS Low (0.165) is consistent with ERAI data showing authorized sources account for fewer than 3% of counterfeit incident reports [2]. These three cases are directional consistency checks against publicly documented outcomes, not formal model validation.

E. Comparative Discrimination Analysis

Table V compares CIRS against three simpler approaches across the five synthetic cases.

The Q-only approach is blind to channel and routing dimensions: it assigns identical scores to cases with equal verification depth, so Cases B and E receive the same "Moderate" rating although Case E is an FTZ-routed broker transaction with more than double the CIRS exposure. The binary channel threshold collapses Cases C, D, and E into a single "Reject" category. The conventional likelihood-impact risk matrix achieves only partial rank agreement with CIRS (Kendall $\tau = 0.60$): it orders Case D above Case E (4.0 vs. 2.2), inverting the CIRS risk ordering, because it does not capture the additional FTZ routing exposure carried by Case E. CIRS ranks E above D by 0.011 units through the FTZ amplifier. While the absolute score difference is 0.011, this crosses procurement decision thresholds in risk-tiered sourcing policies where High-category suppliers trigger mandatory additional qualification steps regardless of exact score value. Under min–max normalization of the risk-matrix scores onto the CIRS score range, the mean absolute deviation between the two instruments is 0.051 units. CIRS produces five distinct scores across a range of 0.361 (0.105–0.466), enabling risk-proportionate resource allocation that no single-factor or matrix approach achieves.

TABLE III
CIRS APPLICATION: FIVE SYNTHETIC PROCUREMENT SCENARIOS ($F = 0.35$, $G = 0.45$, $k = 0.60$ CONSTANT)

Case	Component	S	M	Q	I_{FTZ}	CIRS	Risk
A	MCU industrial	0.05	0.20	0.65	0	0.105	Low
B	Analog IC	0.40	0.40	0.45	0	0.228	Low
C	FPGA logic	0.80	0.60	0.20	0	0.365	Moderate
D	Power module	1.00	0.75	0.00	0	0.455	High
E	SRAM memory	1.00	0.85	0.45	1	0.466	High

Case E: FTZ-routed broker. The FTZ effect enters through $I_{FTZ} = 1$ and $\delta_{FTZ} = 0.059$, not by increasing S.

TABLE IV
DIRECTIONAL VALIDATION: CIRS AGAINST PUBLIC INCIDENT RECORDS ($F = 0.35$, $G = 0.45$ CONSTANT)

ID	Profile	S	M	Q	CIRS	Level	Actual Outcome
V1	Hong Dark Electronic Trade (SASC 2012) [10]	1.00	0.75	0.00	0.455	High	≈84,000 suspect counterfeit parts
V2	Publicly reported defense-supplier audit profile [12]	0.70	0.50	0.30	0.315	Mod.	60% moderate-to-high risk
V3	Authorized distributor (Arrow/Avnet)	0.05	0.50	0.65	0.165	Low	$<3\%$ of ERAI incidents [2]

TABLE V
COMPARATIVE DISCRIMINATION: CIRS VS ALTERNATIVE ASSESSMENT APPROACHES

Case	Binary S	Q-only	Risk Matrix	CIRS	Level
A: MCU authorized	Accept	Low	0.1	0.105	Low
B: Analog IC (AS6081)	Accept	Mod.	0.9	0.228	Low
C: FPGA broker	Reject	High	2.6	0.365	Moderate
D: Power module	Reject	High	4.0	0.455	High
E: SRAM (FTZ)	Reject	Mod.	2.2	0.466	High

F. Exploratory Channel Characterization: CILM-IRI Dataset

To complement directional validation, an anonymized observational dataset of unsolicited supplier contacts (CILM-IRI, a dataset label associated with the author's Component Integrity and Lifecycle Management methodology; $n = 67$, archived at IEEE DataPort [19][12]) is used to illustrate channel-risk conditions relevant to the S variable. The dataset spans 44 calendar days (15 Feb – 30 Mar 2026) and comprises inbound commercial offers received without prior interaction or known provenance of contact acquisition. No prior business relationship existed with any sender, and the recipient contact address was not intentionally distributed. All messages are preserved in original format. No assumption is made regarding supplier intent or product quality. The dataset is not intended to represent the full distribution of supplier channels, but to isolate and characterize a specific class of unsolicited channel entry. The dataset was collected from a channel class that the S variable is designed to characterize: unsolicited supplier outreach without prior relationship or traceable contact provenance. Accordingly, the observed channel-risk indicators cluster in an elevated cold-channel regime. This dataset is used to characterize that channel class, not to test whether CIRS can distinguish between outcome-labeled counterfeit and non-counterfeit transactions. That distinction rests on the contrast table and the directional incident cases, not on this dataset. Outcome labels (i.e., which of these 67 contacts actually involved counterfeit goods) are not available, and ROC or AUC analysis is therefore not applicable here. While source-specific, the observed channel pattern is consistent with documented unsolicited outreach practices in electronic component spot markets, and is unlikely to be unique to a single recipient context, given the scalability of email-based supplier outreach and the prevalence of non-authorized supply channels documented in industry reporting systems [2], [10].

These events exhibit an elevated channel-risk regime. Observable channel properties (absence of OCM-authorization documentation, non-traceable contact provenance, and no prior qualification record) correspond to $S = 0.80$ (81%, $n = 54$) and $S = 0.70$ for messages carrying self-declared certification claims (19%, $n = 13$). Across all 67 records, the observed channel properties place the dataset above authorized and AS6081 reference profiles under controlled baseline assumptions. This supports characterization of unsolicited cold-channel outreach as an elevated structural risk regime, but does not constitute outcome-based counterfeit validation. Rather

[1]An earlier working version of this dataset contained 56 records; the IEEE DataPort version used throughout this article contains 67 records.

[2]In the published dataset the score field is named `crs_score` (formula version CRS v1.1). "CRS" is the repository-anonymized label for the CIRS channel-risk scoring described in this paper; `crs_score` records the channel-provenance risk level (the S variable): 0.80 for offers without verifiable certification ($n = 54$) and 0.70 for self-declared certification (field `has_cert_claim` = true, $n = 13$). The composite CIRS values reported below place these channel-risk levels on the full $[0, 1]$ scale under baseline F, M, Q, G.

than providing classification contrast, this observation highlights that unsolicited cold-channel entry constitutes a distinct and empirically recurrent class of procurement risk exposure.

A consistent pattern emerges from the brand data: recognized OCM names — Vishay (mentioned in 22% of records), ADI (18%), TI (16%) — appear consistently across messages, yet all originate from unverifiable contact channels. This decoupling between brand signaling and channel integrity supports the independent treatment of the F and S variables in the CIRS framework: F captures product-level counterfeit intensity (OECD GTRIC-p), while S captures channel provenance — and a message citing a recognized OCM brand does not reduce the S-driven risk contribution.

A further observation concerns the independence of channel risk from message-level persuasion signals. Of 67 records, 44 (66%) carry neither a certification claim nor price-urgency language; 20 carry one signal; 3 carry both. Price-urgency (field `has_urgency_signal`) does not change the channel-risk score — records with and without urgency receive an identical S within their certification tier — whereas a self-declared certification claim (field `has_cert_claim`) corresponds to the lower-uncertainty self-cert tier ($S = 0.70$ versus $S = 0.80$ for offers with no verifiable certification). Within this dataset, therefore, structural channel characteristics (entry conditions and provenance traceability), not behavioral message tone, drive the risk score. Incoming inspection intensity cannot be calibrated on the basis of persuasion signals alone.

To illustrate CIRS discrimination capability, Table VI contrasts the CILM-IRI channel-risk regime against two reference channel profiles from Table III.

The contrast illustrates that CIRS differentiates between channel types: the 0.142-unit gap between AS6496 (Low, 0.165) and the CILM-IRI cold-email regime (Moderate, 0.307) reflects exclusively the S-variable contribution, holding all other parameters constant. These findings are consistent with industry-level observations on the growing importance of structured risk management in electronic component supply chains. Together, these observations indicate that channel provenance is not a secondary attribute but a primary structural determinant of supply-chain risk exposure — a relationship that the S variable formalizes and that this dataset renders empirically concrete. This dataset does not provide outcome-based validation. It captures the existence and behavioral structure of channel-risk events that CIRS is designed to formalize, and illustrates that the channel properties operationalized in the S variable are empirically recurrent in industrial procurement environments.

V. Discussion

A. Regulatory Reporting as the Primary Driver

The regression results establish that market volume does not emerge as a statistically significant predictor across all five specifications. The D_{NDAA} dummy is the dominant predictor ($\beta \approx +455$–480 across M3–M5, $p < 0.05$ or better in all three), accounting for an R^2 increment of approximately 0.34. The COVID-19 period produces a symmetric suppression effect. ERAI's own within-period attribution of 2024 growth to market expansion reflects contemporaneous narrative interpretation rather than econometric estimation; it describes a marginal, short-run relationship in a stable regulatory regime distinct from the long-run panel dynamics analyzed here. The two conclusions are compatible. This finding has a direct methodological consequence: a risk index using raw ERAI counts without controlling for reporting regime would systematically overestimate risk in the post-2012 period relative to pre-2012 baselines. CIRS addresses this by treating F as a normalized OECD GTRIC-p score rather than a raw incident count.

B. Channel Structure as the Primary Controllable Risk Lever

ERAI's 2024 availability status data resolves what would otherwise appear paradoxical: active components readily available through authorized sources were reported as counterfeit more than twice as often as components with long manufacturer lead times [2]. This pattern is consistent with higher transaction volume and increased market exposure in high-liquidity channels, where counterfeit insertion benefits from the cover of normal commerce. Supply channel risk is a structural property of the procurement pathway independent of market conditions.

The sensitivity analysis quantifies the asymmetry: the maximum CIRS reduction attributable to Q (full AS6171 testing) is 0.150. This bound becomes structurally binding for events with baseline CIRS ≥ 0.60, where full testing cannot reach the Moderate tier ($0.60 - 0.150 = 0.45$, on the High boundary). Organizations that substitute enhanced incoming inspection for upstream channel qualification accept an exposure component that testing alone cannot eliminate. CIRS makes this asymmetry explicit and operational. Industry evidence further supports the structural primacy of procurement channel integrity: a NIST-documented case study of counterfeit product incident reports in GIDEP and ERAI confirms that authorized-channel procurement policies mitigate the majority of counterfeit infiltration risk, identifying channel qualification as the dominant risk lever independent of downstream verification depth [16]. This risk is acknowledged at the original component manufacturer level: AMD's 2025 Annual Report on Form 10-K, filed with the U.S. Securities and Exchange Commission, explicitly identifies gray market procurement (purchasing outside authorized distribution channels) as a primary pathway for counterfeit and substandard products entering customer supply chains, with documented consequences including elevated failure rates and brand protection risks [17].

C. Why a Composite Index Rather Than an Extended Regression

The 44% residual variance in the best-fitting specification (M5, $R^2 = 0.56$) reflects the absence of the most relevant covariates from longitudinal datasets. Supply channel structure, incoming inspection depth, sourcing geography at the transaction level, and supplier qualification status are directly observable within individual procurement events but do not exist as time-series covariates in any publicly available panel.

TABLE VI
CHANNEL-RISK CONTRAST: CILM-IRI OBSERVATIONS VS REFERENCE PROFILES

Channel Type	S	Q	**CIRS**	**Risk**
AS6496 authorized distributor	0.05	0.65	0.165	Low
AS6081 independent distributor	0.40	0.45	0.247	Low
CILM-IRI: no cert claim ($n = 54$)	0.80	0.45[a]	0.307	Moderate
CILM-IRI: self-cert ($n = 13$)	0.70	0.45[a]	0.292	Moderate

[a]Industry-average $Q = 0.45$ (AS6081 baseline); actual Q is unobservable from message content. $F = 0.35$, $M = 0.50$, $G = 0.45$ are held constant to isolate the marginal effect of S, not to represent a specific procurement scenario.

Adding further macroeconomic controls would substitute observable but less relevant proxies for unobservable but more relevant determinants. A composite index is the appropriate instrument when the risk-relevant variables are observable at the decision level but absent from aggregate datasets [4], [5].

D. Positioning CIRS Among Existing Approaches

Existing counterfeit risk frameworks fall into two categories. Process-oriented frameworks (AS5553, AS6081, AS6171) specify process requirements but do not produce a quantitative risk score. Qualitative risk assessment tools capture channel and vendor factors but lack standardized metrics and do not incorporate macroeconomic or systemic trade-level exposure. CIRS occupies the gap: it integrates supply channel structure (S), verification depth (Q), macroeconomic market pressure (M), counterfeit intensity (F), and systemic geographic concentration (G) into a single, calibrated, reproducible score. Unlike MCDA methods such as AHP or TOPSIS, which derive weights from expert judgment matrices that may lack inter-rater reliability in procurement contexts, CIRS uses a structurally assigned, near-equal weight baseline justified under data-scarce conditions and stress-tested through $\pm 20\%$ single-weight perturbation (Section IV-C); its calibration parameters, in turn, are anchored to publicly available standards, trade-risk indices, and documented audit evidence. The FTZ routing indicator operationalizes G parsimoniously as an additive amplifier consistent with OECD-documented premiums [9]. Historical industry cases suggest that supply chain transparency interventions can produce substantial counterfeit risk reductions: vendor-reported data cited in an NIST industry program documents up to a 97 percent reduction in CPU counterfeiting following introduction of provenance verification mechanisms [16]. These results are product- and context-specific and are cited as illustrative rather than universally generalizable; they nonetheless indicate the order of magnitude of risk reduction achievable when channel-level transparency is implemented.

E. Empirical Consistency and Structural Validation

Although CIRS does not produce outcome-based validation — transaction-level procurement data from electronic component supply chains are commercially sensitive and not publicly available — three structural consistency properties can be verified directly from the model and the observational dataset.

Score monotonicity. As established in Section III-H, CIRS is monotonically increasing in channel risk S and decreasing in verification quality Q, with all other parameters held constant. This property holds analytically across the full parameter space and is confirmed numerically for the five synthetic cases in Table III: Cases A through E produce strictly ordered scores (0.105, 0.228, 0.365, 0.455, 0.466) consistent with increasing risk exposure.

Rank stability under parameter perturbation. Sensitivity analysis (Section IV-C) demonstrates that $\pm 20\%$ single-weight perturbation yields a Kendall rank-order correlation of 1.00 in 9 of 10 tests and 0.80 in one (mean $\tau = 0.98$), with maximum score deviation bounded at 0.027 CIRS units. Because CIRS is additive in its continuous input terms before clipping and includes no explicit interaction terms, the one-at-a-time analysis provides an interpretable first-order sensitivity assessment; a full variance-based decomposition (including effects introduced by clipping, tier boundaries, and the discrete FTZ indicator) remains a useful extension for future work. This indicates that CIRS rankings are stable under parameter uncertainty — a critical property for decision-support instruments deployed under incomplete information.

Structural consistency with observational data. The CILM-IRI observational dataset ($n = 67$, 44 calendar days) exhibits the structural pattern that CIRS is designed to detect: all records fall within the elevated channel-risk range $S \in [0.70, 0.80]$, consistent with the absence of traceable contact provenance. The discrete S-distribution ($S = 0.80$: 81%; $S = 0.70$: 19%) reflects the empirically observable distinction between no-certification and self-declared-certification channel types. This observation supports the characterization of unsolicited cold-channel outreach as a structurally distinct channel class rather than outcome-validated risk classification.

Together, these properties establish that CIRS behaves consistently with its theoretical design across both synthetic scenarios and real-world observational evidence, supporting its use as a structured decision-support instrument.

F. Threats to Validity

Construct validity. CIRS operationalizes channel risk, verification depth, and macro-level exposure through proxy variables derived from publicly available standards and trade data. These proxies may not fully capture all dimensions of counterfeit risk in every procurement context, particularly for novel supply chain configurations not represented in existing standards.

External validity. The regression analysis uses ERAI incident data from the North American reporting ecosystem. The CILM-IRI dataset reflects unsolicited outreach received in one organizational context within one market over a 44-day window. Generalization to other regions, sectors, or channel types requires independent replication.

Outcome-label limitation. Neither the synthetic procurement scenarios nor the CILM-IRI dataset carries confirmed counterfeit outcome labels. CIRS scores reflect structural risk conditions, not confirmed counterfeit events. Predictive accuracy cannot be assessed without outcome-labeled transaction data.

Parameter and weight uncertainty. The $k = 0.60$ proxy and $\delta_{FTZ} = 0.059$ additive term are derived from specific publicly audited sources and may not generalize across all commodity categories or regulatory environments. The near-equal baseline weights represent a conservative choice, not empirically optimized weights.

FTZ boundary uncertainty. The FTZ amplifier applies only when routing through free-trade zones is documented or suspected. In practice, FTZ exposure may be difficult to ascertain from procurement documentation alone, which limits operational precision of the δ_{FTZ} term.

Exposure-layer proxy granularity. F and G are drawn from OECD GTRIC indices computed at the HS 2-digit product level and the economy level, respectively. These are associational propensity measures derived from customs-seizure data under documented bias corrections, not causal counterfeit probabilities; the electronics module (HS85) aggregates a broad range of goods and does not isolate electronic components specifically. Sector- and component-specific GTRIC resolution, if made available, would refine F and G.

G. Reproducibility Package

All CIRS scores in this article can be reproduced by running:

```
python compute_risk_score.py
  --input ECB_SC_Risk_Dataset_v4.csv
  --output scored.csv
```

The final reproducibility package [19] contains the following files:

- `ECB_SC_Risk_Dataset_v4.csv` — 67-record anonymized dataset.
- `compute_risk_score.py` — reproduces CIRS scores from the CSV.
- `codebook_v4.2.pdf` — variable definitions and coding rules.
- `scoring_rubric_CIRS_v1.1.pdf` — Q-variable scoring hierarchy.
- `anonymization_protocol.pdf` — SHA-256 procedure and confidentiality handling.
- `irr_plan.pdf` — inter-rater reliability plan ($\kappa \geq 0.70$).
- `FG_normalization_note_v1.1.pdf` — normalization rules for F and G (OECD GTRIC-p / GTRIC-e).
- `model_specification_CIRS_v1.1.pdf` — canonical formula, parameters, thresholds, coding rules, and validation oracle.

H. Limitations and Research Agenda

Four limitations define the research agenda. First, Q-component weights represent a structured interpretation of the AS6171 test hierarchy, not a numerically validated detection-rate schedule. Calibration against controlled detection experiments available through the CALCE Electronic Products and Systems Center is the most direct path to empirical grounding. Second, the FTZ adjustment ($+0.059$) derives from aggregate OECD customs data and may not accurately reflect the risk premium specific to HS85 electronics; sector-specific calibration is needed. Third, the directional validation in Section IV-D and Table IV does not constitute formal predictive validation. The essential next step is application of CIRS to outcome-labeled transaction-level procurement data, enabling comparison between structural CIRS scores and confirmed counterfeit outcomes. Fourth, the regression layer applies ordinary least squares to the annual incident count Y_t; because the annual counts are large and strictly positive (hundreds to over one thousand), the Gaussian approximation is reasonable, but count-data specifications (Poisson or negative-binomial regression) are a natural robustness check and are left for future work.

VI. Conclusion

This paper introduces CIRS as a composite counterfeit risk scoring instrument for electronic component procurement. The regression analysis establishes that no single observable factor provides statistically sufficient explanation of counterfeit incident dynamics, with regulatory reporting regime changes accounting for the dominant share of identifiable variance. This motivates the composite approach: the risk-relevant variables are observable at the procurement decision level but unavailable as panel covariates.

CIRS operationalizes three relationships relevant to procurement-level risk assessment: the channel risk transfer relation $k \cdot S$, using $k = 0.60$ as a sensitivity-tested baseline proxy for channel-risk transfer under weakly traceable procurement conditions; the bounded verification quality ceiling (maximum -0.150); and the FTZ routing amplifier consistent with OECD trade intelligence. Directional validation against three publicly documented incident profiles shows directional classification consistency. The comparative analysis demonstrates that CIRS provides discrimination granularity within the moderate-to-high risk range that binary and single-factor approaches cannot achieve.

All CIRS parameters, regression specifications, and calculation steps are fully reproducible from publicly available data sources cited herein. The CILM-IRI dataset ($n = 67$), archived on IEEE DataPort, supports the observational channel-characterization component of the study; the regression analysis uses ERAI and WSTS data, and CIRS calculations are reproduced through the accompanying scoring script.

Ethics and Conflict of Interest

Data and privacy. The CILM-IRI dataset was built from unsolicited inbound commercial emails received in the normal course of professional work in electronic component procurement — cold-outreach messages from parties who initiated contact without a prior relationship. Before any analysis, all sender domain names were replaced with one-way SHA-256 hash values; no personally identifiable information (names, email addresses, company names, or phone numbers) appears in the published dataset. No institutional review board approval was required for the analysis of non-personal, professionally received commercial correspondence.

Conflict of interest. The author works in the electronic component distribution sector and has a professional interest in supply chain risk management. CIRS is published as open methodology: no proprietary implementation is described or implied, all parameters are derived from publicly available sources (OECD, DoD, SAE standards), and the dataset is archived openly on IEEE DataPort. The author declares no financial relationship with any organization whose standards or data are cited.

Data Availability Statement

The dataset supporting the findings of this study is openly available on IEEE DataPort at https://dx.doi.org/10.21227/34y3-zj88 [19]. It comprises 67 anonymized electronic component procurement communications scored with the CIRS methodology described in this article, together with a Python reproduction script, codebook, data dictionary, inter-rater reliability plan, and anonymization protocol. The related methodological white paper is available at https://doi.org/10.5281/zenodo.19657865. Regression inputs (ERAI annual incident counts and WSTS semiconductor sales) derive from the publicly available sources cited herein.

Appendix A
OLS Regression Diagnostics

This appendix reports diagnostic tests for the M5 specification reported in Table II, addressing potential violations of OLS assumptions in a time-series context with $n = 20$ observations.

A.1 Residual Diagnostics

Table VII summarizes four standard diagnostic tests applied to M5 residuals.

TABLE VII
OLS Diagnostic Tests — M5 Specification ($n = 20$, 2005–2024)

Test	Statistic	p-value	Result
Durbin–Watson	1.273	—	Mild positive autocorrelation
Shapiro–Wilk (normality)	$W = 0.947$	0.260	Residuals normal
Breusch–Pagan (heterosced.)	$r = 0.373$	0.105	Homoscedastic
ADF lag-1 autocorrelation	$\rho = 0.899$	< 0.001	Non-stationary levels[a]

ADF: augmented Dickey–Fuller. [a]Expected for count data with trend; see A.3 for robustness.

The Durbin–Watson statistic of 1.273 indicates mild positive autocorrelation in M5 residuals ($DW < 1.5$), which inflates OLS standard errors. To address this, Newey–West heteroscedasticity- and autocorrelation-consistent (HAC) standard errors were computed with a Bartlett kernel and bandwidth $h = 2$.

A.2 Newey–West HAC Standard Errors

Table VIII compares OLS and HAC standard errors for the M5 specification. The key finding, the dominance of D_{NDAA}, is robust to HAC correction.

TABLE VIII
M5 Coefficients: OLS vs Newey–West HAC Standard Errors

Variable	$\hat{\beta}$	SE$_{OLS}$	SE$_{HAC}$	t_{HAC}
Intercept	778.0	99.6	121.5	6.40***
M_t	−0.62	0.59	0.72	−0.86
D_{NDAA}	473.0	137.9	165.5	2.86***
D_{COVID}	−458.0	168.4	207.9	−2.20**

$^{**}p < 0.05$, $^{***}p < 0.01$. HAC: Bartlett kernel, $h = 2$.
$R^2 = 0.561$; RMSE = 186.2 (invariant to SE method).

Under HAC correction, D_{NDAA} remains highly significant ($t_{HAC} = 2.86$, $p < 0.01$), confirming the robustness of the key empirical finding to autocorrelation in residuals. Market volume M_t remains non-significant ($t_{HAC} = -0.86$), strengthening rather than weakening the motivation for the composite index.

A.3 Robustness: Exclusion of COVID-19 Period

To assess whether results are driven by the 2020–2021 structural break, Model M3 (without D_{COVID}) was re-estimated on the subsample excluding 2020–2021 ($n = 18$).

TABLE IX
Robustness Check: Excluding COVID-19 Period ($n = 18$, 2005–2019, 2022–2024)

Variable	Full sample M3	Excl. 2020–2021
$\hat{\beta}(D_{NDAA})$	480.0***	365.6***
p-value	0.008	< 0.001
R^2	0.354	0.775

$^{***}p < 0.01$.

The NDAA reporting regime dummy remains highly significant ($p < 0.01$) and accounts for $R^2 = 0.775$ of variance in the COVID-excluded subsample (Table IX), compared to 0.354 in the full sample. Market volume M_t is non-significant in both specifications ($p > 0.10$). These results confirm that the key diagnostic finding, the dominance of regulatory reporting infrastructure over market volume, is not an artifact of the COVID-19 structural break.

A.4 Interpretation Note

The high lag-1 autocorrelation in Y_t levels ($\rho = 0.899$) reflects the non-stationary nature of count data with trend, which is expected for a growing reporting system. The diagnostic analysis is conducted on M5 residuals rather than raw

Y_t levels; residual normality (Shapiro–Wilk $p = 0.260$) and homoscedasticity (Breusch–Pagan $p = 0.105$) are confirmed. This appendix is provided in response to reviewer guidance and does not alter the primary analytical conclusions reported in Section IV-A.

APPENDIX B
HISTORICAL CHANNEL-RISK PERSISTENCE

To assess the temporal persistence of the unsolicited cold-channel outreach pattern documented in the CILM-IRI dataset, supplementary evidence was drawn from two long-term legacy email channels (denoted Legacy Channel A and Legacy Channel B), both inactive as primary correspondence channels for over three years prior to this analysis.[3] Each channel independently accumulated more than 1,000 inbound commercial offers over a multi-year observation window (2022–2026).

Preliminary review confirms that the structural characteristics observed in the primary CILM-IRI dataset (unsolicited contact, absence of prior relationship, non-traceable contact acquisition, and high-risk channel indicators) are consistently present across the extended historical record. This historical evidence supports the characterization of unsolicited cold-channel outreach as a structurally persistent and organizationally recurrent risk class, rather than a time-bounded anomaly. Detailed analysis of this extended dataset is reserved for a dedicated future publication.

APPENDIX C
CIRS CALCULATION PSEUDOCODE

The following pseudocode summarizes the CIRS computation for a single procurement event. A Python implementation is included in the dataset archive at IEEE DataPort [19].

```
Input: F, S, M, Q, G, I_FTZ
       k=0.60, d_FTZ=0.059
       w=[0.25, 0.25, 0.20, 0.15, 0.15]

Q_score = 0.20*I_vis + 0.25*I_elec + 0.20*I_x2d
        + 0.15*I_xct + 0.20*I_dec

CIRS = w[0]*F + w[1]*(k*S) + w[2]*M
     - w[3]*Q_score + w[4]*G
     + d_FTZ * I_FTZ

CIRS = clip(CIRS, 0.0, 1.0)

if   CIRS < 0.25: tier = "Low"
elif CIRS < 0.45: tier = "Moderate"
elif CIRS < 0.65: tier = "High"
else:             tier = "Critical"
```

ACKNOWLEDGMENT

The author acknowledges the use of Anthropic Claude for language editing, structural feedback, and formatting support in Sections I, V, and VI, and ChatGPT for translation support, formatting checks, and consistency review. Scientific content—including model specification, parameter selection, regression analysis, dataset construction and anonymization, interpretation of results, and verification of all references and factual claims—was developed and verified independently by the author. No AI tool is listed as an author, and no AI tool was used as an autonomous source of scientific conclusions. The author reviewed and revised all AI-assisted text and takes full responsibility for the content of the manuscript.

[3]Channel identifiers are functional labels assigned for this article; the original mailbox names, hosting providers, and organizational affiliations are withheld in accordance with the project anonymization protocol. The information disclosed is limited to what is methodologically necessary: that the channels are distinct, long-inactive, and independently sampled.

REFERENCES

[1] OECD/EUIPO, "Mapping Global Trade in Fakes 2025," OECD Publishing, 2025. doi: 10.1787/94d3b29f-en

[2] D. Akhoundov, "2024 Annual Report," ERAI, Inc., 2025. Available: https://www.erai.com/erai_blog/3187/_2024_annual_report

[3] SMT Corp., "Growing Threat of Counterfeit Electronic Parts in the Critical Infrastructure Supply Chain," May 2023. Available: https://smtcorp.com/

[4] A. Pavlov, D. Ivanov, A. Dolgui, and B. Sokolov, "Hybrid fuzzy-probabilistic approach to supply chain resilience assessment," *IEEE Trans. Eng. Manag.*, vol. 65, no. 2, pp. 303–315, 2018, doi: 10.1109/TEM.2017.2773574.

[5] D. DiMase, Z. A. Collier, J. Carlson, R. B. Gray, and I. Linkov, "Traceability and risk analysis strategies for addressing counterfeit electronics in supply chains for complex systems," *Risk Analysis*, vol. 36, no. 10, pp. 1834–1843, Oct. 2016, doi: 10.1111/risa.12536.

[6] OECD/EUIPO, "Trends in Trade in Counterfeit and Pirated Goods," 2019. doi: 10.1787/g2g9f533-en

[7] OECD/EUIPO, "Trade in Counterfeit and Pirated Goods: Mapping the Economic Impact," 2016. doi: 10.1787/9789264252653-en

[8] OECD, "The Economic Impact of Counterfeiting and Piracy," OECD Publishing, Paris, 2008. doi: 10.1787/9789264045521-en

[9] OECD/EUIPO, "Trade in Counterfeit Goods and Free Trade Zones," 2018. doi: 10.1787/9789264289550-en

[10] U.S. Senate Armed Services Committee, "Inquiry into Counterfeit Electronic Parts in the Department of Defense Supply Chain," S. Rept. 112-167, May 2012. [Online]. Available: https://www.armed-services.senate.gov/

[11] H. Akhavantaheri, P. Sandborn, and D. Das, "An enterprise network model for understanding and disrupting illicit counterfeit electronic part supply chains," *IISE Transactions*, vol. 56, no. 3, pp. 257–268, 2024. doi: 10.1080/24725854.2023.2172631

[12] E. Smith, "Mitigating the risk of counterfeit electronics," *Military Embedded Systems*, Jul. 2015. [Online]. Available: https://militaryembedded.com/comms/rf-and-microwave/mitigating-risk-counterfeit-electronics

[13] U. Guin, K. Huang, D. DiMase, J. M. Carulli, M. Tehranipoor, and M. Pecht, "Counterfeit integrated circuits: A rising threat in the global semiconductor supply chain," *Proc. IEEE*, vol. 102, no. 8, pp. 1207–1228, 2014. doi: 10.1109/JPROC.2014.2332291

[14] National Institute of Standards and Technology, "Cybersecurity Supply Chain Risk Management Practices for Systems and Organizations," NIST SP 800-161r1, May 2022. doi: 10.6028/NIST.SP.800-161r1

[15] Intel Corporation, "2026 Intel Platform Security Report," Intel Security, 2026. Available: https://www.intel.com/content/www/us/en/security/product-security-assurance.html

[16] D. A. Brown and Intel Corporation, "Managing Risk End-to-End in Intel's Supply Chain," in *Best Practices in Cyber Supply Chain Risk Management*, NIST U.S. Resilience Project, NIST, 2015. Available: https://www.nist.gov/publications/best-practices-cyber-supply-chain-risk-management

[17] Advanced Micro Devices, Inc., "Annual Report on Form 10-K for the Fiscal Year Ended December 28, 2025," U.S. Securities and Exchange Commission, March 2026. Available: https://www.sec.gov/cgi-bin/browse-edgar?action=getcompany&CIK=AMD&type=10-K

[18] A. Litvin, "Risk logistics: a new architecture of semiconductor supply chains," *Forbes Kazakhstan*, Apr. 2026. Available: https://forbes.kz/articles/logistika-riska-novaya-arhitektura-postavok-mikroelektroniki-32aa39

[19] A. Litvin, "Electronic Component Supply Chain Risk Dataset: Annotated Procurement Communications with Composite Risk Scores ($N = 67$)," IEEE DataPort, May 2026, doi: 10.21227/34y3-zj88. Available: https://dx.doi.org/10.21227/34y3-zj88

[20] M. Pecht and S. Tiku, "Bogus: Electronic manufacturing and consumers confront a rising tide of counterfeit electronics," *IEEE Spectrum*, vol. 43, no. 5, pp. 37–46, May 2006, doi: 10.1109/MSPEC.2006.1628506.

[21] Z. A. Collier, S. Walters, D. DiMase, J. M. Keisler, and I. Linkov, "A semi-quantitative risk assessment standard for counterfeit electronics detection," *SAE Int. J. Aerosp.*, vol. 7, no. 1, pp. 171–181, 2014, doi: 10.4271/2014-01-9002.

[22] O. Lavastre, A. Gunasekaran, and A. Spalanzani, "Supply chain risk management in French companies," *Decision Support Systems*, vol. 52, no. 4, pp. 828–838, Mar. 2012, doi: 10.1016/j.dss.2011.11.003.

[23] R. S. Gaonkar and N. Viswanadham, "Analytical framework for the management of risk in supply chains," *IEEE Trans. Autom. Sci. Eng.*, vol. 4, no. 2, pp. 265–273, Apr. 2007, doi: 10.1109/TASE.2006.880540.

[24] C. S. Tang, "Perspectives in supply chain risk management," *Int. J. Prod. Econ.*, vol. 103, no. 2, pp. 451–488, Oct. 2006, doi: 10.1016/j.ijpe.2005.12.006.

[25] D. R. Richman, M. H. Azarian, D. Das, and M. Pecht, "Analysis of Standards-Based Counterfeit Microelectronics Detection Methods," *IEEE Access*, vol. 13, pp. 7691–7704, 2025, doi: 10.1109/ACCESS.2025.3526518.

[26] E. Weiss, "Electronic Components Provenance Assurance Through AI and Big Data in Assembly Processes," *IEEE Trans. Compon., Packag., Manuf. Technol.*, 2024, doi: 10.1109/TCPMT.2024.3482467. [Online]. Available: https://ieeexplore.ieee.org/document/10720842

Alexander Litvin (Senior Member, IEEE) is an Independent Researcher in electronic component supply chain risk and integrity analytics, based in Yekaterinburg, Russia. He has over 27 years of professional experience in semiconductor distribution and procurement risk management. He developed the Component Integrity & Lifecycle Management (CILM) methodology and the Composite Integrity Risk Score (CIRS) framework presented in this article.

Mr. Litvin holds professional membership in the Association for Supply Chain Management (ASCM) and is active in IEEE Technology and Engineering Management Society (TEMS) and IEEE COMPAS communities. His work focuses on operational risk scoring frameworks, channel integrity analysis, and decision-support tools for procurement under sparse-data conditions. He has served as a jury member in multiple international industry awards related to electronic components and supply chain innovation.